

Linear Algebraic Techniques in Algorithms and Complexity

by

Josh Alman

S.B., Massachusetts Institute of Technology (2014)

M.S., Stanford University (2016)

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author

Department of Electrical Engineering and Computer Science

August 30, 2019

Certified by.....

R. Ryan Williams

Associate Professor of Electrical Engineering and Computer Science

Thesis Supervisor

Certified by.....

Virginia Vassilevska Williams

Steven and Renee Finn Career Development Associate Professor of

Electrical Engineering and Computer Science

Thesis Supervisor

Accepted by

Leslie A. Kolodziejcki

Professor of Electrical Engineering and Computer Science

Chair, Department Committee on Graduate Students

Linear Algebraic Techniques in Algorithms and Complexity

by
Josh Alman

Submitted to the Department of Electrical Engineering and Computer Science
on August 30, 2019, in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Abstract

We develop linear algebraic techniques in algorithms and complexity, and apply them to a variety of different problems. We focus in particular on *matrix multiplication algorithms*, which have surprisingly fast running times and can hence be used to design fast algorithms in many settings, and *matrix rank methods*, which can be used to design algorithms or prove lower bounds by analyzing the ranks of matrices corresponding to computational tasks.

First, we study the design of matrix multiplication algorithms. We define a new general method, called the Universal Method, which subsumes all the known approaches to designing these algorithms. We then design a suite of techniques for proving lower bounds on the running times which can be achieved by algorithms using many tensors and the Universal Method. Our main limitation result is that a large class of tensors generalizing the Coppersmith-Winograd tensors (the family of tensors used in all record-holding algorithms for the past 30+ years) cannot achieve a better running time for multiplying n by n matrices than $O(n^{2.168})$.

Second, we design faster algorithms for batch nearest neighbor search, the problem where one is given sets of data points and query points, and one wants to find the most similar data point to each query point, according to some distance measure. We give the first subquadratic time algorithm for the exact problem in high dimensions, and the fastest known algorithm for the approximate problem, for various distance measures including Hamming and Euclidean distance. Our algorithms make use of new probabilistic polynomial constructions to reduce the problem to the multiplication of low-rank matrices.

Third, we study rigid matrices, which cannot be written as the sum of a low rank matrix and a sparse matrix. Finding explicit rigid matrices is an important open problem in complexity theory with applications in many different areas. We show that the Walsh-Hadamard transform, previously a leading candidate rigid matrix, is in fact not rigid. We also give the first nontrivial construction of rigid matrices in a certain parameter regime with applications to communication complexity, using an efficient algorithm with access to an NP oracle.

Thesis Supervisor: R. Ryan Williams

Title: Associate Professor of Electrical Engineering and Computer Science

Thesis Supervisor: Virginia Vassilevska Williams

Title: Steven and Renee Finn Career Development Associate Professor of Electrical Engineering and Computer Science

Acknowledgments

First and foremost, I want to thank my advisors Ryan Williams and Virginia Vassilevska Williams. They've supported me in every stage of my grad school career, from helping me to find research topics that I'm excited about, to encouraging me when I'm totally stuck on research problems, to teaching me how to write and communicate and apply for jobs. Working together with them has been a great experience, not only because of their endless streams of great research ideas, but also because of their strong passion for theoretical computer science and their desire to share it with others. Outside of work, Ryan and Virginia made life a pleasure, between hosting weekly music nights, feeding me chocolate, having long sports discussions with me in the middle of meetings, and just being great friends. I doubt any of the work in this thesis would have been possible without their support, and I hope this is only the beginning of our collaborations and friendships.

I also want to thank the other mentors I have had throughout my research career so far. Piotr Indyk, my third thesis committee member, also supervised my first computer science research experience as part of a final project for sublinear algorithms class. Greg Valiant was my rotation advisor during my first few months at Stanford; he helped get me acquainted with grad student life, and introduced Chebyshev polynomials to me. Vitaly Feldman and T.S. Jayram mentored me during a fruitful summer at IBM Almaden where I learned about some practical applications of theory. Pavlo Pylyavskyy advised me in the amazing REU in combinatorics at the University of Minnesota, Twin Cities, which has probably been the most productive couple of months of my life. Lionel Levine and Lorenzo Orecchia mentored me in research projects while I was an undergraduate at MIT, and taught me me how exciting research can be. All of these mentors helped to shape who I am as a researcher.

I want to thank everyone I've collaborated with, both for their contributions and for everything they've taught me: Arkadev Chattopadhyay, Timothy M. Chan, Lijie Chen, Timothy Chu, Cesar Cuenca, Vitaly Feldman, Jiaoyang Huang, Robin Hui, T.S. Jayram, Carl Lian, Dylan McKay, Matthias Mnich, Liat Peterfreund, Aaron Schild, Zhao Song, Brandon Tran, Virginia Vassilevska Williams, Greg Valiant, Nikhil Vyas, Joshua Wang, Ryan Williams, and Huacheng Yu. Research is much more enjoyable with such great people to work with.

I've had the unique experience of spending my grad school career at two wonderful schools: Stanford for the first half, and MIT for the second half. Both the Stanford and MIT theory groups are warm and welcoming communities full of great people to work, learn, and hang out with. I especially want to thank my colleagues and friends at both schools and throughout the TCS community, who made my time in grad school so memorable, including: Amir Abboud, Greg Bodwin, Vaggos Chatziafratis, Lijie Chen, Tobias Christiani, Michael Cohen, Ofir Geri, Daniel Grier, Robin Hui, Gautam Kamath, Michael Kim, Jerry Li, Andrea Lincoln, Quanquan Liu, Alex Lombardi, Nathan Manohar, Dylan McKay, Saeed Mehraban, Yonatan Naamad, Govind Ramnarayan, Luke Schaefer, Adam Sealfon, Vatsal Sharan, Zhao Song, Warut Sukhompong, Paris Syminelakis, Joshua Wang, Nicole Wein, and Huacheng Yu. Special shoutout to Dylan McKay, my only friend who would write a paper about League of

Legends with me [\[AM17\]](#).

Finally, I want to thank my parents Zena and Ben, my sister Sophie, my girlfriend DD, and my whole family, for always supporting me, even when they don't necessarily understand what I'm doing.

Contents

1	Introduction	11
1.1	Matrix Multiplication and Matrix Rank	11
1.2	Summary of Results	13
1.3	Bibliographic Details	22
2	Preliminaries	23
2.1	Notation	23
2.2	Boolean and Arithmetic Circuits	25
2.3	Models of Communication	26
2.4	Tail Bounds and Probabilistic Tools	27
2.5	Bounds on Binomial Coefficients	27
2.5.1	Multinomial Coefficients	29
I	Limitations on Matrix Multiplication Algorithms	31
3	Background and Overview	33
3.1	Our Results	34
3.1.1	The Universal Method	34
3.1.2	Limits on the Universal Method	35
3.1.3	Additional Results	38
3.2	Other Related Work	38
3.3	Bibliographic Details	40
4	The Universal Method	41
4.1	Tensors and Tensor Rank	41
4.2	Matrix Multiplication Tensors	42
4.3	Tensor Powers and Asymptotic Rank	45
4.4	Reductions Between Tensors	46
4.5	The Universal Method	50
4.5.1	The Solar and Galactic Methods	52
4.6	The Known Approaches to Matrix Multiplication	53
4.6.1	The Laser Method	53
4.6.2	The Group-theoretic Method	55

5	Limits on the Universal Method	57
5.1	Asymptotic Slice Rank	57
5.2	Limits on the Universal Method from Asymptotic Slice Rank	58
5.3	Combinatorial Tools for Asymptotic Slice Rank Upper Bounds	60
5.3.1	Bounds from Variable-Deficient Partitions	60
5.3.2	Bounds from Block Partitions	61
5.3.3	Bounds from Parts with Low X-Rank	64
5.4	Slice Rank Lower Bounds via the Laser Method	65
5.4.1	Slice Rank Versus Asymptotic Subrank	70
5.5	Computing the Slice Ranks for Tensors of Interest	71
5.5.1	Generalized Coppersmith-Winograd Tensors	71
5.5.2	Generalized Simple Coppersmith-Winograd Tensors	73
5.5.3	Cyclic Group Tensors	74
5.5.4	Lower Triangular Tensors	75
5.6	Upper Lower Bounds for Group Tensors	76
5.6.1	Tri-Colored Sum-Free Sets	76
5.6.2	Asymptotic Slice Rank and Tri-Colored Sum-Free Set Constructions for All Finite Groups	78

II Probabilistic Polynomials and Hamming Nearest Neighbors 81

6	Background and Overview	83
6.1	Our Results	84
6.1.1	Polynomial Constructions	84
6.1.2	Algorithmic Applications	86
6.2	Other Related Work	91
6.3	Bibliographic Details	93
7	Probabilistic Polynomials	95
7.1	Multilinear Polynomials Computing Boolean Functions	95
7.2	Typically Correct Polynomials	96
7.2.1	Interpolating Polynomials for Symmetric Functions	97
7.2.2	The Razborov-Smolensky Lower Bound	99
7.3	Probabilistic Polynomials	101
7.3.1	The Probabilistic Degree of OR	101
7.3.2	The Probabilistic Degree of Biased Threshold Functions	103
7.4	Probabilistic Polynomials for Majority	104
7.5	Further Probabilistic Polynomial Constructions	108
7.5.1	Symmetric Functions	108
7.5.2	Derandomization	109
7.6	PTFs for ORs of Threshold Functions	110
7.6.1	Deterministic Construction	111
7.6.2	Probabilistic Construction	113

8	Algorithmic Applications	117
8.1	Sparse Polynomials and Rectangular Matrix Multiplication	117
8.2	Exact Batch Nearest Neighbor Search	119
8.2.1	Closest Pair in Hamming Space is Hard	122
8.2.2	Metrics Beyond Hamming Distance	123
8.3	Approximate Batch Nearest Neighbor Search	124
8.4	The Light Bulb Problem	127
8.4.1	Deterministic Algorithms	130
8.5	Faster Algorithms For MAX-SAT	132
8.6	Circuit Satisfiability Algorithms	136
8.6.1	Satisfiability on a Cartesian Product	136
8.6.2	Simulating LTFs with AC0 of MAJORITY	138
8.6.3	Satisfiability for ACC of LTF of LTF	140
8.6.4	Satisfiability for Three Layers of Majority and AC0	144
III	Probabilistic Rank and Matrix Rigidity	147
9	Background and Overview	149
9.1	Our Results	151
9.1.1	Probabilistic Rank and Matrix Rigidity	151
9.1.2	Efficient Construction of Rigid Matrices Using an NP Oracle	155
9.2	Other Related Work	158
9.3	Bibliographic Details	161
10	Probabilistic Rank and Matrix Rigidity	163
10.1	Polynomials, Rank, and Rigidity	163
10.2	Non-Rigidity of Walsh-Hadamard	165
10.2.1	Rigidity Upper Bound for Low Error	165
10.2.2	Rigidity Upper Bound for High Error	169
10.2.3	Generalization To SYM-AND circuits	170
10.3	Equivalence Between Probabilistic Rank and Rigidity	170
10.3.1	Generalization to Random Self-Reducible Functions	172
10.4	Explicit Rigid Matrices and Threshold Circuits	173
10.5	Sign-Rank Rigidity and Depth-Two Threshold Circuits	175
10.6	Equivalence Between Probabilistic Rank Modulo m and BP-MODm Communication Complexity	178
11	Efficient Construction of Rigid Matrices Using an NP Oracle	181
11.1	Construction Overview	181
11.1.1	Either $NE \not\subseteq P_{=poly}$ or a Construction of Rigid Matrices	181
11.1.2	Unconditional Construction of Rigid Matrices	187
11.2	Tools from Complexity Theory	189
11.3	Construction of Rigid Matrices Assuming an Easy Witness Lemma	191
11.4	Unconditional Construction of Rigid Matrices	198

11.5 Applications	201
11.5.1 PH^{cc} Communication Lower Bound	201
11.5.2 Depth-2 Arithmetic Circuit Lower Bound	202
11.5.3 Threshold Circuit Lower Bound for E^{NP}	204
11.6 Algorithm for Counting Orthogonal Vectors over Finite Fields	204
11.6.1 Reduction to Prime Fields	204
11.6.2 Algorithm for Prime Fields	206

Chapter 1

Introduction

In this Chapter, we give a high-level introduction to this dissertation. After this Introduction and the Preliminaries in Chapter 2, we get into the technical components of the dissertation, which are divided into three main Parts: Part I is about matrix multiplication algorithms, Part II is about the polynomial method and nearest neighbor search algorithms, and Part III is about matrix rigidity. Each of those Parts begins with its own, more technical overview of the results therein.

1.1 Matrix Multiplication and Matrix Rank

Linear algebra is used throughout computer science, including in areas like error correcting codes, signal analysis, graphics, optimization algorithms, secure encryption and secret sharing schemes, graph analysis algorithms like PageRank, and a number of important machine learning algorithms. Understanding and improving the linear algebraic ideas underlying these different applications is imperative to the theory and practice of computation.

In this dissertation, we develop novel bridges between linear algebra and computer science. We will use linear algebraic techniques which were originally designed for certain tasks, and repurpose them to solve new problems. At the same time, rather than just using known linear algebraic techniques, many of our results will require new linear algebraic concepts and constructions. Although a number of known concepts from linear algebra will make appearances, there are two in particular which will play a central role: *matrix multiplication*, and *matrix rank*.

Matrix Multiplication Matrix multiplication is one of the most basic algebraic operations, and most computational tasks in linear algebra can be performed in the same number of arithmetic operations as matrix multiplication, including computing the determinant [Str69] or the inverse [Str69, BH74] of a matrix, computing various matrix factorizations [BH74], solving systems of linear equations [Str69], and even solving linear programs [CLS18]. Thus, almost every algorithmic application of linear algebra makes use of matrix multiplication algorithms, and the ‘bottleneck’ in designing faster algorithms for these applications is frequently matrix multiplication.

It was widely believed that two $n \times n$ matrices cannot be multiplied using fewer than n^3 arithmetic operations until 1969, when Strassen [Str69] published a breakthrough algorithm which uses only $O(n^{2.81})$ arithmetic operations. We measure the running time of matrix multiplication algorithms in terms of the number of required arithmetic operations since, when multiplying matrices with very large entries, the time to multiply and add those entries can also significantly contribute to the total running time. That said, typically in applications, the entries of the matrices are small enough that this contribution is negligible.

Since Strassen's algorithm, an enormous amount of work has gone into speeding up matrix multiplication. The best known theoretical algorithm for matrix multiplication uses a multitude of clever ideas coming from algorithm design, algebra, and combinatorics (see Figure 1-1 below). Matrix multiplication in practice has received as much if not more attention, with a number of hardware and software optimizations which yield fast practical algorithms.

Reducing other algorithmic problems to matrix multiplication gives a way to apply the same ideas and algorithms to a wide variety of computational tasks. Algorithmic problems from areas as diverse as parsing, graph algorithms, cryptography, statistics, and learning theory, which a priori have nothing to do with matrix multiplication, have been sped up by clever reductions to matrix multiplication. Hence, understanding the computational complexity of matrix multiplication is one of the most central and applicable problems in computer science.

Matrix Rank The rank of a matrix M measures the 'intrinsic dimensionality' of the rows and columns of M . In today's era of 'big data' and high-dimensional datasets, it is no surprise that matrix rank can play a role in algorithm design: the observation that some high-dimensional matrix actually has low rank can often be used to perform computations faster on that matrix. Indeed, most basic operations can be performed faster on low-rank matrices (see e.g. [CKL13]).

Matrix rank can also be helpful for proving lower bounds. Roughly, the idea is to show that if a matrix M has high rank (or a variant on rank), then M is so complicated that computations related to M cannot be done efficiently. Rank methods like this can be used in the most evident way to show lower bounds for algebraic problems directly related to M . For example, if M has high rank, then computing the linear transformation defined by M applied to an input vector cannot be done with small, constant-depth linear circuits. In fact, almost every known lower bound in arithmetic complexity theory has been proved via rank methods; see e.g. [EGOW18].

Rank methods can also be used to show lower bounds for functions and models of computation that are seemingly unrelated to linear algebra. Often one can associate a matrix M_f with a function f , and show that if the rank of M_f is high, then f itself cannot be computed by efficient algorithms or protocols. For example, if one can show that the truth table matrix of a Boolean function f has high rank, then this implies f does not have efficient deterministic communication protocols [MS82].

1.2 Summary of Results

In this dissertation, we develop linear algebraic tools in algorithms and complexity, especially developing the theory of matrix multiplication algorithms, and introducing new rank methods in algorithms and complexity. We apply these tools, often in conjunction with each other, to solve a number of computer science problems, including problems where linear algebraic tools haven't been used before.

We now describe the main contributions of this dissertation. The main body of the dissertation is arranged into three Parts.

Part I: Limitations on Matrix Multiplication Algorithms

Since Strassen published his algorithm in 1969, there has been a long line of work developing many different tools for designing faster matrix multiplication algorithms, leading to the best known running time of about $O(n^{2.373})$ arithmetic operations [Wil12, LG14]. In Figure 1-1, we show the history of the best known exponent of matrix multiplication over time, i.e. the best upper bound on the constant ω such that one can multiply two $n \times n$ matrices using about $O(n^\omega)$ arithmetic operations. It is popularly conjectured that one can design an algorithm achieving $\omega = 2$, and this conjecture is very appealing since it would mean that matrix multiplication, along with many different applications, can be solved in nearly linear time in the input size.

Figure 1-1: The best known exponent of matrix multiplication over time. For instance, Strassen's breakthrough algorithm from 1969 [Str69] multiplies n matrices using $O(n^{2.81})$ arithmetic operations, so the point (1969, 2.81) is plotted above.

One striking feature of Figure 1-1 is that, although there was a flurry of improvements to ω in the first 20 years after Strassen's result, the best known value has remained nearly unchanged for more than 30 years. In Part I of this dissertation, we help to explain why progress has been mostly stagnant for so long.

The known approaches to designing matrix multiplication algorithms follow a formula which uses two main components:

1. An efficient algorithm for evaluating some order-3 tensor T . Matrix multiplication can be seen as the task of evaluating a prescribed set of bilinear polynomials on a given input, which can in turn be seen as evaluating a certain order-3 tensor on a given input; T corresponds to such a task but with a different set of bilinear polynomials.
2. A method of reducing from one tensor to another, so that algorithms for evaluating the latter can be converted into algorithms for evaluating the former.

Combined, these two components give a matrix multiplication algorithm, by using the method to reduce to T , and then applying the algorithm for T . Both major approaches to designing matrix multiplication algorithms—the 'Laser Method' spearheaded by Strassen [Str87] which is used to achieve the best current bound on ω as well as the more recent Group-theoretic Method introduced by Cohn and Umans [CU03]—follow this formula, but with restrictions on what tensor T may be used in component (1), and only a limited type of reduction used in component (2). In particular, all the record-holding algorithms for the past 30+ years have come from such an approach where T is the 'Coppersmith-Winograd tensor' CW_q introduced by [CW90].

In Chapter 4, we define a new generalization of all the known approaches to designing matrix multiplication algorithms, which we call the Universal Method. It makes use of the most general type of reduction between tensors which is known to be applicable in component (2) of the above formula, called α -generation.

Then, in Chapter 5, we prove lower bounds on the algorithms one can design using the Universal Method. We show that if the Universal Method is applied to any tensor in a big family generalizing CW_q , then the resulting upper bound on ω cannot be better than 2.168 . Our limitation result is quite general, and also applies to all other record-holding tensors in the history of matrix multiplication algorithms. Hence, in order to prove a better bound on ω , one must take a radically different approach, either by starting with a tensor T which is very different from those which have led to the best algorithms, or else by analyzing tensors to yield matrix multiplication algorithms in an entirely new way.

Our limitation result, which at first seems to be a negative one, actually leads to a number of interesting algorithmic ideas. First, in defining the Universal Method itself, we highlight steps in the current best matrix multiplication algorithms where more powerful techniques may be possible but aren't being used; while our result rules out achieving a running time of $O(n^2)$ using these techniques, it doesn't rule out an improved running time of, say, $O(n^{2.2})$. Second, in the process of proving our limitation result, we also identify a large number of fundamentally different algorithms, arising from new, different tensors, which are able to match the best-known bound on ω of 2.373 . Perhaps one of these different algorithms will help improve our matrix multiplication algorithms.

The proof of our limitation result makes use of a measure of complexity of a tensor called its slice rank. When one generalizes the notion of matrix rank to tensors, there

are a number of natural ways to do so; slice rank is one such generalization. Roughly, we observe that any tensor whose slice rank is not high enough is too simple to be used with the Universal Method to design very fast matrix multiplication algorithms. We then give slice rank upper bounds for $\mathbb{C}W_q$ and a wide variety of other tensors.

Part II: Probabilistic Polynomials and Hamming Nearest Neighbors

The polynomial method has been a powerful tool for studying Boolean functions, at least since Minsky and Papert's 1969 book [MP69]. The method concerns representing Boolean functions by 'simple' polynomials. Minsky and Papert first used the polynomial method as a way to prove limitations on different models of computation (they focused in particular on 'perceptrons'). Roughly, the idea is to show that

1. any function computed by a certain model of computation can also be computed by a simple polynomial, and
2. some particular Boolean function f cannot be computed by a simple polynomial.

Combined, this means the model of computation cannot compute the function. This method is still one of the most popular approaches today for proving lower bounds in complexity theory.

There are numerous ways to measure the 'simplicity' of a polynomial, but the most common is to use the polynomial's degree. Consider, for instance, the Boolean OR function on n inputs from $\{0, 1\}^n$. It can be computed exactly by the polynomial

$$p(x_1, x_2, \dots, x_n) = 1 - \prod_{i=1}^n (1 - x_i)$$

Indeed, if all the x_i are 0, then p evaluates to 0, but if any of them is 1, then p evaluates to 1. However, p has degree n , which is as big as possible every Boolean function on n inputs can be computed by some polynomial of degree at most

One way around this high degree is to weaken the constraints on what it means for a polynomial to 'compute' a function. For instance, instead of aiming for an exact polynomial for OR, we can instead design a probabilistic polynomial. A distribution P on n -input polynomials such that for every $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ we have

$$\Pr_P [p(x_1, x_2, \dots, x_n) = \text{OR}(x_1, x_2, \dots, x_n)] \geq 1 - \epsilon$$

for some error parameter $\epsilon > 0$. We can design such a probabilistic polynomial over F_2 (the field with two elements, i.e. the integers mod 2) as follows: to draw a polynomial from P , pick $k = \lceil \log(1/\epsilon) \rceil$ independent uniformly random subsets $I_1, I_2, \dots, I_k \subseteq \{1, 2, \dots, n\}$, and output

$$p(x_1, x_2, \dots, x_n) = \prod_{i=1}^k \sum_{j \in I_i} x_j$$

If all the x_i are 0, then p always evaluates to 0. Otherwise, for each $i \in \{1, 2, \dots, k\}$, the sum $\sum_{j \in I_i} x_j$ is odd with probability $1/2$, and so the polynomial p evaluates to 1

(over F_2) with error only 2^{-k} . P is a distribution on polynomials of degree only $O(\log(1/\epsilon))$; for constant $\epsilon > 0$, this is constant degree!

The polynomial method is concerned with trade-offs like the above: between the guarantees of a polynomial, and the degree or simplicity that can be achieved. Probabilistic polynomials, for instance, can achieve much lower degrees than exact polynomials, but they have a chance of outputting the wrong value.

While complexity theorists see these low-degree polynomial representations as a weakness of the computational model, algorithm designers can instead view them as algorithmic tools: If a critical subroutine in an algorithm can be converted into a low-degree polynomial, then a fast algorithm for manipulating polynomials can sometimes be applied to speed up that subroutine, and solve the original problem faster. This viewpoint has led to the same polynomials, which complexity theorists designed for proving lower bounds, being used in the design of faster algorithms for many problems, including in learning theory [Val15], constraint satisfaction [Wil14c], and graph algorithms [Wil14a].

In Part II of this dissertation, we apply the polynomial method in novel ways to design new algorithms and prove new lower bounds. Our results critically make use of a connection between low-degree polynomials and low-rank matrices: if the entries of a matrix can be computed by a low-degree polynomial, then that polynomial can be used to construct a low-rank representation of the matrix. This can allow us to use fast matrix multiplication as the fast algorithm for manipulating polynomials from the previous paragraph. In other words, the polynomial method can be seen as a way to design faster algorithms for many different algorithmic problems by giving reductions to matrix multiplication, allowing us to take advantage of fast matrix multiplication algorithms.

In Chapter 7, we design new low-degree polynomial representations of Boolean functions. We focus in particular on polynomial representations of threshold functions like the majority function MAJ, although our results will extend to symmetric Boolean functions as well as a number of classes of Boolean circuits. Threshold functions arise naturally in many settings, including in linear programming, in machine learning algorithms like perceptrons and neural networks, and in nearest neighbor search (as we will discuss shortly).

We first construct a probabilistic polynomial for MAJ on n inputs with error ϵ and degree $O(\frac{n}{\epsilon} \log(1/\epsilon))$. This matches a classical $(\frac{n}{\epsilon} \log(1/\epsilon))$ degree lower bound due to Razborov [Raz87] and Smolensky [Smo87]; they originally introduced probabilistic polynomials and proved this degree lower bound in order to show a circuit lower bound, that MAJ cannot be computed by AC^0 circuits.

We then show it is possible to circumvent Razborov and Smolensky's lower bound and achieve even lower degree polynomials. To do this, we consider a new generalization of a probabilistic polynomial which we call a probabilistic polynomial threshold function (probabilistic PTF). While a probabilistic polynomial for a function f must exactly compute f on any given input with high probability, a probabilistic PTF must only output a positive real number when f is true, and a negative real number when f is false, with high probability. It is easy to construct a degree 1 probabilistic PTF for MAJ (in fact, randomness isn't even needed), but we aim to design a probabilistic

PTF for an OR of many MAJ functions, which is much less straightforward. One way to do this is to sum together independent copies of our probabilistic polynomial for MAJ, resulting in a probabilistic PTF for an OR of $O(1/\epsilon)$ different MAJs, each on n inputs, with degree $O(\sqrt{n \log(1/\epsilon)})$. However, we are able to improve on this construction and achieve degree only $O(n^{1/3} \log^{2/3}(n/\epsilon))$. Our construction combines ideas from the design of randomized algorithms with 100-year-old constructions from polynomial approximation theory, especially the Chebyshev polynomials [Che99].

New Nearest Neighbor Search Algorithms Next, in Chapter 8, we apply our polynomial constructions to design new algorithms as well as prove new lower bounds. Our main algorithmic application is for nearest neighbor search problems. In the (batch) nearest neighbor search problem, one is given as input m data points, and n query points, and the goal is to find the nearest data point to each query point. Nearest neighbor search has applications in almost every domain, including computational geometry, coding theory, pattern recognition, and DNA sequencing. There are, of course, many different settings of this problem depending on the specific details.

First, one needs to specify what types of data points can be input, and how one should measure the distance between them. Some natural choices include:

- points from the d -dimensional Boolean hypercube $\{0, 1\}^d$, with distance measured by the Hamming distance, which counts the number of the entries in which two points differ,
- points from d -dimensional Euclidean space \mathbb{R}^d , with distance measured by the standard Euclidean metric,
- points which are d -dimensional vectors of real numbers, from \mathbb{R}^d , along with a different distance measure like the ℓ_1 distance (also known as Manhattan distance), and
- points which are subsets of a large universe, with distance measured by the Jaccard index, which equals the ratio of the size of the symmetric difference and the size of the union of the two sets.

Second, one needs to choose whether exact nearest neighbors are necessary, or whether approximate nearest neighbors are sufficient. In the $(1 + \epsilon)$ -approximate nearest neighbor problem for some small constant $\epsilon > 0$, it is sufficient to find, for each query point, a data point which is within a $(1 + \epsilon)$ factor of the distance to the actual nearest data point. In many applications, finding approximate nearest neighbors is sufficient, and such a relaxation can allow for substantially faster algorithms.

In all these choices of settings, there is a brute force quadratic-time algorithm, which simply iterates over all pairs of points and computes their distances. However, quadratic time can be too slow in applications with many data points! Using the polynomial method, we give the fastest known, subquadratic time algorithm for the approximate problem for all of the aforementioned distance measures, and the exact problem for most of the distance measures, in high dimensions where subquadratic time algorithms weren't previously known.

To give two examples:

For the exact batch nearest neighbor search problem with Hamming distance, in

dimension $d = c \log n$, we design an algorithm with running time $n^{2 - \Theta(\frac{1}{c})}$. For any constant c , this is a truly subquadratic running time (i.e. time $O(n^2)$ for some constant $\epsilon > 0$). Previously, no truly subquadratic time algorithm was known even in dimension, say, $d = 2 \log n$.

For the $(1 + \epsilon)$ -approximate batch nearest neighbor search problem with Hamming distance, in any dimension d , we design an algorithm with running time $dn + n^{2 - \epsilon^{1/3}}$. This algorithm runs in subquadratic time, even up to dimension $d = n^{\epsilon}$. For small enough constant $\epsilon > 0$, this improves on the previous best running time of $dn + n^{2 - \epsilon^{1/2}}$ [Val15], as well as running times $dn + n^{2 - \epsilon}$ which one can achieve using techniques like Locality-Sensitive Hashing [IM98].

We also achieve nearly identical results for Euclidean distance, Jaccard distance, and many other choices of distance measure instead of Hamming distance. We complement our algorithms with new conditional lower bounds, showing that if one can design algorithms which are significantly faster than the ones we design here, it would refute a popular conjecture from complexity theory (the 'Strong Exponential Time Hypothesis' [IPZ01]) about the time required to solve the Boolean satisfiability problem.

Because threshold functions are so versatile, we apply our polynomial constructions to other applications as well, including basic problems in data analysis and statistics, constraint satisfaction problems like MAX-SAT, and new circuit lower bounds for circuits with threshold gates.

Part III: Probabilistic Rank and Matrix Rigidity

Informally, a matrix is called rigid if it has high rank, and one must change many of its entries before it has low rank. Of course there are parameters involved: the rank-rigidity of a $N \times N$ matrix M , denoted $R_M(r)$, is the minimum number of entries of M which one must change in order to make its rank at most r .

Consider, for example, the $N \times N$ identity matrix I_N . Although I_N has full rank, it is not particularly rigid: each time a 1 on the diagonal changes to 0, the rank of I_N decreases by one. Hence, for all ranks r we have $R_{I_N}(r) = N - r$. In fact, since changing one entry of a matrix can never decrease the rank by more than one, the identity matrix I_N is as non-rigid as a full rank matrix can be!

The above example can be summarized by saying that the identity matrix is not rigid because it is sparse. We could try to get around this to find a rigid matrix by considering simple dense matrices instead, like perhaps the $N \times N$ upper triangular matrix U_N with all 1s above the diagonal. However, with some work one can show that U_N isn't very rigid either. We could then move on to even more complicated matrices like a Vandermonde or Fourier matrix, and although these seem more rigid, it's hard to prove that this is the case. In fact, it's an open problem to show that any explicit matrices, like these, are rigid (for certain parameters we describe in the next paragraph).

Finding explicit rigid matrices has been a central open challenge in complexity ever since the notion of rigidity was introduced by Leslie Valiant in 1977 [Val77]. At a high level, rigid matrices are of interest because they are inherently complicated:

following the outline of using rank methods for proving lower bounds, complexity theorists have shown that explicit rigid matrices would yield new lower bounds for several different models of computation. The two most interesting rigidity parameter regimes for a family M_N of $N \times N$ matrices, where M_N is a $N \times N$ matrix, are as follows:

A family M_N is called Valiant-rigid if there is a constant $\epsilon > 0$ such that

$$R_{M_N}(N = \log \log N) \geq (N^{1+\epsilon}):$$

Valiant [Val77] showed that the linear transformations corresponding to Valiant-rigid matrices cannot be computed by $O(N)$ -size $O(\log N)$ -depth arithmetic circuits. There are currently no known lower bounds showing that such circuits cannot compute any explicit families of matrices.

A family M_N is called Razborov-rigid if there is any super-constant function $\epsilon(N) = \omega(1)$ such that

$$R_{M_N}(2^{(\log \log N)^{\epsilon(N)}}) \geq (N^2):$$

Razborov [Raz89] (see also [Wun12]) showed that if the communication matrix M_f of a Boolean function f is Razborov-rigid, then f is not in PH^{cc} , the communication analogue of the polynomial hierarchy. There are currently no explicit Boolean functions known to be outside PH^{cc} .

In other words, M_N is Valiant-rigid if you have to change a super-linear number of its entries to make its rank drop to barely sublinear, and M_N is Razborov-rigid if reducing it to a tiny rank requires changing a constant fraction of its entries. Showing that an explicit family of matrices M_N is rigid in either regime would imply new, breakthrough lower bounds in complexity theory.

We have used the word 'explicit' a number of times; what does it mean? We say M_N is explicit if there is a deterministic algorithm which, on input N , outputs the matrix M_N in $\text{poly}(N)$ time. Aiming for a deterministic algorithm is important, since random matrices are known to be very rigid. For instance, for any rank $r = o(N)$, a random $\{0, 1\}$ matrix $M_N \in \{0, 1\}^{N \times N}$ over any field has rigidity $R_{M_N}(r) \geq (N^2)$ with high probability, and is hence both Valiant-rigid and Razborov-rigid. This is not particularly exciting, since random functions are already known to require large circuits and inefficient communication protocols. Aiming for a $\text{poly}(N)$ time construction is also important since, with enough time, an algorithm could simply brute force over, say, every $N \times N$ matrix over F_2 , and compute the rigidity of every one. Beyond these details, finding and understanding explicit rigid matrices is important, since these are matrices which are efficiently computable in one sense, but inherently inefficient in other (e.g. can't be computed by super-linear size circuits or PH^{cc} protocols).

To summarize: although rigidity was defined more than 40 years ago, and explicit rigid matrices are known to have many important applications throughout complexity

theory, there are still no known constructions of explicit rigid matrices. In Part III of this dissertation, we make new progress on this problem, both by explaining the weakness of current attempts at proving matrices are rigid, and by giving new non-trivial constructions of rigid matrices. We approach the problem by combining linear algebra with techniques from algorithm design and complexity theory which haven't been used before in this setting, including fast matrix multiplication algorithms, the polynomial method, and a new generalization of matrix rank which we call probabilistic rank.

Hadamard Matrices Are Not Rigid Among the many attempts to prove lower bounds via rigidity, perhaps the most commonly studied explicit matrix has been the Walsh-Hadamard transform [PS88, Alo90, Gri, Nis, KR98, Cod00, Lok01, LTV03, Mid05, dW06b, Ras16]. The Walsh-Hadamard transform is a family of matrices of size $2^n \times 2^n$ with entries ± 1 , defined recursively as:

$$H_1 = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}; \text{ and } H_{n+1} = \begin{pmatrix} H_n & H_n \\ H_n & -H_n \end{pmatrix} \text{ for } n \geq 1:$$

H_n is a Hadamard matrix (i.e. its rows are mutually orthogonal) corresponding to the discrete Fourier transform for the power of the cyclic group \mathbb{Z}_2^n , and it has applications in areas like quantum computing, signal processing, and data compression. These properties were believed to imply that H_n must be rigid; in fact, many of the references above were working toward proving that every Hadamard matrix is rigid. The best known rigidity lower bounds for H_n for rank r are that $R_{H_n}(r) = \Omega(2^n - r)$, which is insufficient to show they are Valiant-rigid or Razborov-rigid.

In Chapter 10, we partially explain why these rigidity lower bounds are not stronger: we refute the popular conjecture that H_n is rigid, and show that it is not Valiant-rigid, by giving a new rigidity upper bound. More precisely, letting $N = 2^n$ be the side-length of H_n , we show that for every field F , and all sufficiently small constants $\epsilon > 0$,

$$R_{H_n}(N^{1-\epsilon}) \leq N^{1+\epsilon}$$

over F . The choice of field F can sometimes make a difference in how rigid a matrix is (for instance, H_n has constant rank over \mathbb{F}_2), but our rigidity upper bound works over any field.

Our proof makes use of a new generalization of matrix rank we introduce, called probabilistic rank. Generalizing the notion of a probabilistic polynomial, we say that a matrix H has probabilistic rank r for error ϵ if there is a distribution M on matrices of the same dimensions as H and rank at most r such that, for every entry $H[i; j]$ of the matrix H , we have

$$\Pr_{M \sim M} [H[i; j] = M[i; j]] \geq 1 - \epsilon:$$

It is not hard to see that a 'typical' matrix M from the probabilistic matrix M is a rank r matrix which differs from H in only an ϵ fraction of its entries, hence giving

a rigidity upper bound for H . However, in principle, there could be better rigidity upper bounds than this, since a rigidity upper bound only requires a small total number of errors, whereas a probabilistic matrix is required to have a low probability of error on every entry of H . In fact, we show that the two notions are equivalent for many families of matrices including the Walsh-Hadamard transform H_n , so it suffices to focus on giving probabilistic rank upper bounds for H_n .

Our new probabilistic rank upper bounds are inspired by our earlier use of the polynomial method in algorithm design. In order to show probabilistic rank upper bounds, we show that probabilistic rank also generalizes probabilistic polynomial degree. When using the polynomial method to design new algorithms, we used the fact that low-degree polynomials which compute the entries of a matrix give rank upper bounds for that matrix. Similarly, if the entries of a matrix are computed by low-degree probabilistic polynomials, then the same connection gives a probabilistic rank upper bound for that matrix. Our rigidity upper bound for H_n critically makes use of a new probabilistic polynomial construction for most of the rows and columns of H_n .

Efficient Construction of Rigid Matrices in P^{NP} Finally, in Chapter 11, we give a new construction of rigid matrices. We give a family of matrices $M_N \in \mathbb{F}_q^{M_N \times N}$, with $M_N = 2^{\Omega(N)} \log^N N$, which is Razborov-rigid (in a slightly different sense), and which can be constructed in deterministic $\text{poly}(N)$ time with access to an NP oracle. More precisely, for infinitely many N , our $N \times N$ matrix M_N has the rigidity bound

$$R_{M_N} \left(2^{(\log N)^{1-\epsilon}} \right) = \Omega(N^2)$$

for any $\epsilon > 0$ over any constant-sized finite field \mathbb{F}_q .

This is the first nontrivial construction of Razborov-rigid matrices which doesn't use randomness. Although it doesn't qualify as an explicit construction in the sense we previously described (since the construction uses an NP oracle), it still implies a number of new lower bounds in complexity theory in conjunction with the various known applications of rigid matrices, including:

There is a function in $\text{TIME}[2^{(\log n)^{1-\epsilon}}]^{NP}$ which is not in PH^{CC} . Here, PH^{CC} is the communication complexity analogue of the polynomial hierarchy, consisting of functions with efficient communication protocols that can make use of alternating nondeterministic and co-nondeterministic guesses by the two players. It was previously even open whether every function in the larger class $\text{TIME}[2^{O(n)}]^{NP}$ is also in AM^{CC} , an important subclass of PH^{CC} .

Depth-2 linear circuits for computing the linear transformation defined by the $N \times N$ matrix M_N described above require size $\Omega(N^2)$. The previous best nontrivial such lower bounds for non-randomly constructed matrices were at best $(N \log^2 N = \log \log N)$.

Our construction takes a very different approach from past attempts at constructing rigid matrices. While past constructions have mainly used tools from algebra and

combinatorics, our construction is instead inspired by circuit complexity theory. The main idea is to view low-rank expressions for matrices as a special type of 'circuit class' for computing matrices. In this way, finding a rigid matrix is equivalent to proving a certain average-case lower bound against these 'circuits'.

In order to prove this circuit lower bound, we use the algorithmic approach introduced by Williams [Wil13]. Very roughly, Williams' approach shows that fast, deterministic algorithms for analyzing circuits (e.g. for counting the number of satisfying assignments to a circuit) can be used to prove lower bounds against those circuits. Tools from linear algebra, algorithm design, and complexity theory come into play as we design the appropriate circuit analysis algorithm and use it to prove our rigidity bound. For instance, we use an algorithm for counting the number of 1s in a low-rank matrix by Chan and Williams [CW16], which applies the polynomial method and fast matrix multiplication.

1.3 Bibliographic Details

This dissertation is based on the results in eight previously published papers:

'Further Limitations of the Known Approaches for Matrix Multiplication' with Virginia Vassilevska Williams [AW18a], which appeared in ITCS 2018,

'Limits on All Known (and Some Unknown) Approaches to Matrix Multiplication' with Virginia Vassilevska Williams [AW18b], which appeared in FOCS 2018,

'Limits on the Universal Method for Matrix Multiplication' [Alm19b], which appeared in CCC 2019 and won the Best Student Paper Award,

'Probabilistic Polynomials and Hamming Nearest Neighbors' with Ryan Williams [AW15], which appeared in FOCS 2015,

'Polynomial Representations of Threshold Functions and Algorithmic Applications' with Timothy M. Chan and Ryan Williams [ACW16], which appeared in FOCS 2016,

'An Illuminating Algorithm for the Light Bulb Problem' [Alm19a], which appeared in SOSA 2019,

'Probabilistic Rank and Matrix Rigidity' with Ryan Williams [AW17], which appeared in STOC 2017, and

'Efficient Construction of Rigid Matrices Using an NP Oracle' with Lijie Chen [AC19], which will appear in FOCS 2019.

See the Introduction to each Part for more specific details about which results correspond to each reference.

Chapter 2

Preliminaries

We assume familiarity with basic facts about algorithms, complexity, combinatorics, probability, and algebra (especially linear algebra and properties of polynomials). That said, in this Chapter, we will review some of the less common notions from these areas which will play important roles throughout this dissertation. We begin first by discussing the notation we will use.

2.1 Notation

Sets, Vectors, and Matrices We use the standard notation for common sets of numbers: \mathbb{Z} is the set of integers, $\mathbb{N} := \{1; 2; 3; \dots\}$ is the set of natural numbers, \mathbb{R} is the set of real numbers, and \mathbb{C} is the set of complex numbers. For $n \in \mathbb{N}$, we write $[n] := \{1; 2; \dots; n\}$, and $\mathbb{Z}_n := \mathbb{Z}/n\mathbb{Z}$ to denote the integers mod n .

For any set S , $n \in \mathbb{N}$, $i \in [n]$, and n -dimensional vector $v \in S^n$, we write $v_i \in S$ for the i th entry of v . We may sometimes write $v(i)$ or $v[i]$ instead of v_i to avoid ambiguity with other subscripts; the meaning should be clear from context. When we say $v \in S^n$ is 'indexed by X ' for a set X of size $|X| = n$, we implicitly define a bijection $m_X : X \rightarrow [n]$, and for $x \in X$ define $v_x := v_{m_X(x)}$.

Similarly, for a $n \times m$ matrix $M \in S^{n \times m}$ over set S , and for $i \in [n]$ and $j \in [m]$, we write $M(i; j)$ (or sometimes $M[i; j]$) for the $(i; j)$ th entry of M . We may also write $M[i; :]$ for the i th row of M , or $M[:, j]$ for the j th column of M . If $X; Y$ are sets of size $|X| = n$ and $|Y| = m$, then we can index the entries of M by $X; Y$ and refer to entry $M(x; y)$ for $x \in X$ and $y \in Y$ (using similar implicit bijections m_X and m_Y as above).

Logarithms and Asymptotics We write \log_b for the base b logarithm. We also write \log for \log_2 and \ln for \log_e for short.

We use the standard symbols from asymptotic analysis $\mathcal{O}; \omega; \Theta; \Omega; \sim$, and \dots . We write $f = \text{poly}(n)$ if there is a constant $c > 0$ such that $f(n) = \mathcal{O}(n^c)$, and write $\text{polylog}(n) := \text{poly}(\log(n))$. If $f; g$ are functions of many variables including n , we write $f = \mathcal{O}_n(g)$ to mean that $f = \mathcal{O}(g)$ when n grows and all variables other than

are fixed to any constant values. We write $f = \Theta(g)$ to ignore polylog factors, i.e. if there is a constant $c \in \mathbb{Z}$ such that $f = O(g \log^c(g))$. Define \sim similarly.

Boolean functions A Boolean function is a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ for some $n \in \mathbb{N}$, where we think of 0 as 'false' and 1 as 'true'. Some simple Boolean functions which will recur include:

OR, which outputs 1 unless all its inputs are 0. We sometimes also write OR_n to emphasize the number of inputs.

AND, which outputs 1 when all its inputs are 1.

XOR, which outputs 1 when an odd number of its inputs are 1. For $x, y \in \{0, 1\}$ we will write $x \oplus y := \text{XOR}(x; y)$, and for $x \in \{0, 1\}^n$ we write $\bigoplus_{i=1}^n x_i := \text{XOR}(x_1; \dots; x_n)$. (The symbol \oplus will also be used for direct sum, but it will be clear from context which of the two meanings we are using.)

MOD_m for $m \in \mathbb{N}$, which outputs 1 when a multiple of m of its inputs are 1. In particular, for any $x \in \{0, 1\}^n$, $\text{MOD}_2(x) = \bigoplus \text{XOR}(x)$.

MAJ (MAJORITY), which outputs 1 when at least half of its inputs are 1.

For $x \in \{0, 1\}^n$, we write $|x| := \sum_{i=1}^n x_i \in \mathbb{Z}$ to denote the Hamming weight of x . A Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is called symmetric if $f(x) = f(y)$ for any $x, y \in \{0, 1\}^n$ such that $|x| = |y|$. All the aforementioned Boolean functions are symmetric.

For a logical predicate P , we use Iverson bracket notation $[P]$ to denote the function which outputs 1 when P is true, and 0 when P is false. For instance, for $x \in \{0, 1\}^n$, we have $\text{MAJ}(x) = [|x| \geq n/2]$. Brackets $[\cdot]$ will also be used as parentheses for emphasis in some places; the meaning should be clear from context.

Groups, Rings, and Fields We will typically use multiplicative notation for the group operation of groups. Two particular groups which will arise frequently are, for $n \in \mathbb{N}$, the cyclic group C_n of order n , and the symmetric group S_n of permutations on n elements.

Every ring R has two distinguished elements: the additive identity 0 and the multiplicative identity 1. When using R to represent a Boolean function, we will use 0 to denote 'false' and 1 to denote 'true'. The characteristic of a ring R is the smallest $n \in \mathbb{N}$ such that

$$\underbrace{1 + 1 + \dots + 1}_{n \text{ times}} = 0$$

over R if such an n exists, and 0 if there is no such n . When $0 \neq 1$ (i.e. when R is not the trivial ring with one element) then the characteristic is never 1. In this case, when c is the characteristic of R , there is a natural (and unique) ring homomorphism from \mathbb{Z} to R whose range is \mathbb{Z}_c (or \mathbb{Z} when $c = 0$). Hence, polynomials over \mathbb{Z} can also

be viewed as polynomials over any commutative ring R where we take all outputs mod c (and in particular, do not change the values 0 and 1).

When R is a commutative ring, for $n \geq 1$ and $x, y \in R^n$, we write $hx; yi_R := \sum_{i=1}^n x_i y_i$. When the ring R is clear from context, we omit it and simply write $hx; yi$. When $q \geq 2$ is a power of a prime, we write F_q for the finite field of order q .

2.2 Boolean and Arithmetic Circuits

We study two types of circuits:

Boolean circuits, whose inputs are Boolean $\{0, 1\}$ values, and whose gates compute Boolean functions of their inputs (the default gates are AND, OR, and NOT), and

Arithmetic circuits over a field F , whose inputs are values from F , and whose gates compute F -valued functions of their inputs (the default gates are $+$ and \cdot). Arithmetic circuits may also take constant values from F as input in addition to the usual inputs. One can more generally consider arithmetic circuits over a ring.

Boolean Circuits The depth, size, fan-in, and set of allowed gates can drastically change what functions can be computed by a given class of circuits. The classes we will encounter in this dissertation include:

AC^0 : functions computable by families of constant-depth unbounded fan-in polynomial-size circuits over the basis $\{\text{AND}, \text{OR}, \text{NOT}\}$

$AC^0[m]$: functions computable by families of constant-depth unbounded fan-in polynomial-size circuits over the basis $\{\text{AND}, \text{OR}, \text{NOT}, \text{MOD}_m\}$

ACC^0 : the union of $AC^0[m]$ for all $m \geq 2$

TC^0 : functions computable by families of constant-depth unbounded fan-in polynomial-size circuits over the basis $\{\text{AND}, \text{OR}, \text{NOT}, \text{MAJ}\}$

It is known that

$$AC^0 \subset AC^0[m] \subset ACC^0 \subset TC^0:$$

It is believed that $MAJ \not\subseteq ACC^0$, and hence that $ACC^0 \subsetneq TC^0$, but this is still an open problem.

We will also study the class LTF of linear threshold functions, i.e. Boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ of the form $f(x) = [\sum_{i=1}^n a_i x_i \geq t]$ for constants $a_i \in \mathbb{R}$ and $t \in \mathbb{R}$. For instance, $MAJ \in LTF$.

For classes of circuits C and D , we write $C \circ D$ to denote the class of circuits consisting of a single circuit $C \in C$ whose inputs are the outputs of some circuits from D . That is, $C \circ D$ is simply the composition of circuits from C and D . For instance, $LTF \circ LTF$ denotes depth-two linear threshold circuits, $LTF \circ \text{MOD}_2$ denotes linear threshold function of parities, etc.

Arithmetic Circuits Since arithmetic circuits only use $+$ and \cdot gates, they always compute sets of polynomials in their inputs. It is natural to ask why arithmetic circuits do not typically allow for $/$ gates. In fact, it is known that allowing $/$ gates can only save polynomial factors in arithmetic circuit size:

Proposition 2.1 ([Str73, HY09]). If a polynomial $f \in \mathbb{F}[x_1; \dots; x_n]$ of degree d can be computed by an arithmetic circuit of size s using $+$; \cdot ; $/$, then it can be computed by an arithmetic circuit of size $\text{poly}(s; r; n)$ using only $+$; \cdot .

This answer is unsatisfying in situations where we care about precise polynomial factors, such as in the design of matrix multiplication algorithms. However, Strassen [Str73] also showed that divisions do not help, even by constant additive or multiplicative factors in the circuit size, when computing sets of quadratic forms (such as matrix multiplication).

One type of arithmetic circuit which will be of particular interest to us is a linear circuit. In such a circuit, each gate computes a linear combination of its inputs. Hence, linear circuits can only compute linear transformations of their inputs. In other words, linear circuits with n inputs and m outputs correspond to matrices $A \in \mathbb{F}^{m \times n}$, such that given as input $x \in \mathbb{F}^n$, the circuit outputs Ax . It is known that any arithmetic circuit for computing a linear transformation can be converted into a linear circuit for the same linear transformation with only constant factor increases in the size and depth (see e.g. [BCS13, Theorem 13.1]).

2.3 Models of Communication

In a communication protocol for a function $F : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}^n$, two players, each given one of $x, y \in \mathbb{F}^n$, send each other messages in order to compute $F(x; y)$. The number of bits of communication used in π is the maximum, over all $x, y \in \mathbb{F}^n$, of the sum of the lengths of the messages (as binary strings) that the players send to each other. See [KN97] for more details.

Starting with [BFS86], a growing line of work has studied the communication complexity analogues of different classical complexity classes. The most relevant communication complexity classes for us will be:

P^{cc} : Functions $F : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}^n$ which can be computed by a deterministic communication protocol using only $\text{poly}(\log(n))$ bits of communication.

NP^{cc} : Functions $F : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}^n$ which can be written as $\bigvee_{i=1}^k R_i(x; y)$, where $k = \text{poly}(n)$, and each R_i is a rectangle i.e. a function of the form $R_i(x; y) = [x \in S_x \wedge y \in S_y]$ for subsets $S_x, S_y \subseteq \mathbb{F}^n$. (This can be viewed as a communication protocol where the two players nondeterministically guess which rectangle is satisfied by their inputs.)

AM^{cc} : Functions $F : \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}^n$ such that there is a distribution D on NP^{cc} protocols π , such that for any $x, y \in \mathbb{F}^n$ we have $\Pr_D [F(x; y) = \pi(x; y)] \geq \frac{2}{3}$.

PH^{cc} : Functions $F : \{0, 1\}^n \rightarrow \{0, 1\}$ which can be computed by a formula which is a $\text{poly}(n)$ -ary tree of constant depth, where each gate computes an AND or an OR, and each leaf computes a rectangle of the inputs.

Similar to the classical complexity setting, we know that

$$P^{cc} \subseteq NP^{cc} \subseteq AM^{cc} \subseteq PH^{cc}.$$

Further communication complexity classes can similarly be defined in a natural way. For instance, $MOD_m P^{cc}$ consists of functions $F : \{0, 1\}^n \rightarrow \{0, 1\}$ which can be written as a MOD_m of $\text{poly}(n)$ many rectangles, and $BP \text{ } MOD_m P^{cc}$ consists of functions which can be computed with high probability on every input by distributions on $MOD_m P^{cc}$ protocols. See [GPW18], for instance, for more about the known relationships between these and other communication complexity classes.

2.4 Tail Bounds and Probabilistic Tools

We assume familiarity with standard tools from probability, including the union bound, Markov's inequality, Chernoff bounds, and Chebyshev's inequality. We will occasionally pay particular attention to the constants in tail bounds on Binomial distributions, by applying the following instantiation:

Lemma 2.1 (Hoeffding's Inequality for Binomial Distributions [Hoe63, Theorem 1]) If m independent random draws $x_1, \dots, x_m \in \{0, 1\}$ are made with $\Pr[x_i = 1] = p$ for all i , then for any $k \geq mp$ we have

$$\Pr\left[\sum_{i=1}^m x_i \geq k\right] \leq \exp\left(-\frac{2(mp - k)^2}{m}\right);$$

where $\exp(x) = e^x$.

When designing deterministic algorithms, we will also need a Chernoff bound for samples with limited independence:

Lemma 2.2 ([SSS95, Theorem 5 (I)(b)]) If X is the sum of k -wise independent random variables, each of which is confined to the interval $[0, 1]$, with $\mu = E[X]$, 1 , and $k = b^2 e^{-1/3} c$, then

$$\Pr[|X - \mu| \geq c] \leq e^{-c^2/3}.$$

2.5 Bounds on Binomial Coefficients

We now present some standard bounds on the growth of binomial and multinomial coefficients. We will make use of these bounds throughout this dissertation.

Proposition 2.2. For all $n; k \in \mathbb{N}$ with $1 \leq k \leq n$, we have

$$\frac{n^k}{k} \leq \binom{n}{k} < \frac{n e^k}{k} :$$

Proof. The lower bound follows since:

$$\binom{n}{k} = \sum_{i=0}^{k-1} \frac{n^i}{i!} \geq \frac{n^{k-1}}{(k-1)!} = \frac{n^k}{k} :$$

For the upper bound, first recall from the Taylor series for e^x that

$$e^k = \sum_{i=0}^{\infty} \frac{k^i}{i!} > \frac{k^k}{k!} :$$

Rearranging gives that $k! > (k/e)^k$. It thus follows that:

$$\binom{n}{k} = \frac{\sum_{i=0}^{k-1} n^i}{k!} \leq \frac{\sum_{i=0}^{k-1} n^i}{k!} = \frac{n^k}{k!} < \frac{n e^k}{k} :$$

□

The bound from Proposition 2.2, which shows that $\binom{n}{k} = (n/k)^k$, will be sufficient in most cases where we need to use bounds on binomial coefficients. However, in some cases where n and k are both large, we will need a tighter bound on the constant hidden by the \dots .

Definition 2.1. The binary entropy function $H : [0, 1] \rightarrow [0, 1]$ is given by

$$H(x) = -x \log_2 x - (1-x) \log_2 (1-x) ;$$

where we define $H(0) = H(1) = 0$. Hence, $2^{H(p)} = \frac{1}{p^p (1-p)^{1-p}}$.

Proposition 2.3. For any $n \in \mathbb{N}$ and any $p \in [0, 1]$ such that pn is an integer, we have

$$\frac{1}{n+1} 2^{n H(p)} \leq \binom{n}{pn} \leq 2^{n H(p)} :$$

Proof. We can verify that the claim is true when $p = 0$ or $p = 1$, so assume $p \in (0, 1)$. Define $T : (0, 1) \rightarrow \mathbb{R}$ by $T(k) = \frac{\binom{n}{k} p^k (1-p)^{n-k}}{2^{n H(p)}}$. Notice in particular that

$$T(pn) = \frac{\binom{n}{pn} p^{pn} (1-p)^{(1-p)n}}{2^{n H(p)}} = \frac{\binom{n}{pn}}{2^{n H(p)}} :$$

Our goal is hence to prove that

$$\frac{1}{n+1} \leq T(pn) \leq 1 :$$

We have that $T(k) > 0$ for all k , and the binomial theorem says that

$$\sum_{k=0}^n T(k) = (p + (1-p))^n = 1:$$

In particular, the upper bound $T(p/n) \leq 1$ immediately follows.

Next, to prove the lower bound, it is sufficient to show that $T(p/n) \leq T(k)$ for all k . This will imply that $T(p/n)$ is at least the average value of all the $T(k)$'s, which by the binomial theorem above, is at least $1/(n+1)$.

More precisely, we will prove that $T(k) - T(k+1)$ is nonnegative when $k < p/n$, and nonpositive when $k > p/n$. This will imply that $(T(k))_{0 \leq k \leq n}$ is a unimodal sequence with maximum at $k = p/n$. To see this, note that

$$\begin{aligned} T(k) - T(k+1) &= \binom{n}{k} p^k (1-p)^{n-k} - \binom{n}{k+1} p^{k+1} (1-p)^{n-k-1} \\ &= \binom{n}{k} p^k (1-p)^{n-k} \left(1 - \frac{n-k}{k+1} \frac{p}{1-p} \right) \end{aligned}$$

Hence, $T(k) - T(k+1) \geq 0$ if and only if $(n-k)p \leq (k+1)(1-p)$, which rearranges to $k \leq p/n - (1-p)$. Since $(1-p) \in (0, 1)$, but the argument k to T must be an integer, this means that $T(k) - T(k+1) \geq 0$ if and only if $k \leq p/n$, as desired. \square

We will also use the following estimate of $H(p)$ when p is close to $1/2$:

Proposition 2.4. For $\epsilon \in (0, 1/2)$ we have $H(1/2 + \epsilon) = 1 - O(\epsilon^2)$.

Proof. This follows from the Taylor expansion of $H(p)$ about $p = 1/2$:

$$H\left(\frac{1}{2} + \epsilon\right) = 1 - \frac{1}{2 \ln 2} \sum_{i=1}^{\infty} \frac{(2\epsilon)^{2i}}{i(2^i - 1)} = 1 - O(\epsilon^2):$$

\square

Corollary 2.1. We have

$$\binom{n}{(1/2 + \epsilon)n} = 2^{n(1 - O(\epsilon^2))}$$

for $\epsilon \in (0, 1/2)$ and $n \in \mathbb{N}$ such that $(1/2 + \epsilon)n$ is an integer.

2.5.1 Multinomial Coefficients

Definition 2.2. For $m \in \mathbb{N}$ and nonnegative integers k_1, k_2, \dots, k_m with $n = k_1 + k_2 + \dots + k_m$, the multinomial coefficient $\binom{n}{k_1, k_2, \dots, k_m}$ is given by

$$\binom{n}{k_1, k_2, \dots, k_m} = \frac{n!}{k_1! k_2! \dots k_m!} = \prod_{i=1}^m \binom{k_1 + k_2 + \dots + k_i}{k_i} :$$

Multinomial coefficients have a combinatorial interpretation similar to that of binomial coefficients: $\binom{n}{k_1, k_2, \dots, k_m}$ counts the number of ways to put n distinct objects into m distinct buckets such that k_i objects are put into the i th bucket for each $i \in [m]$. We will make use of the following approximation, which follows directly by applying Proposition 2.3:

Proposition 2.5. For any $m \in \mathbb{N}$ and any constants $p_1, p_2, \dots, p_m \in [0, 1]$ such that $p_1 + p_2 + \dots + p_m = 1$, we have

$$\binom{n}{p_1 n, p_2 n, \dots, p_m n} = \frac{1}{p_1^{p_1 n} p_2^{p_2 n} \dots p_m^{p_m n}}^{n+o(n)}$$

for all $n \in \mathbb{N}$ such that $p_i n$ is an integer for all $i \in [m]$.

Part I

Limitations on Matrix Multiplication Algorithms

Chapter 3

Background and Overview

One of the biggest open questions in computer science asks how quickly one can multiply two matrices. Progress on this problem is measured by giving bounds on ω , the exponent of matrix multiplication, defined as the smallest real number such that two $n \times n$ matrices over a field can be multiplied using $n^{\omega + \epsilon}$ field operations for any $\epsilon > 0$. Trivially, $2 \leq \omega \leq 3$. Many have conjectured over the years that $\omega = 2$, and this conjecture is very attractive: a near-linear time algorithm for MM would immediately imply near-optimal algorithms for many problems.

Almost 50 years have passed since Strassen [Str69] first showed that $2.81 < \omega < 3$. Since then, an impressive toolbox of techniques has been developed to obtain faster MM algorithms, culminating in the current best bound $\omega < 2.373$ [LG14, Wil12]. Unfortunately, this bound is far from 2, and the current methods seem to have reached a standstill. Recent research has turned to proving limitations on the two main MM algorithmic techniques: the Laser Method of Strassen [Str86] and the Group-theoretic Method of Cohn and Umans [CU03].

Both Coppersmith and Winograd [CW90] and Cohn et al. [CKSU05] proposed conjectures which, if true, would imply that $\omega = 2$. The first conjecture works in conjunction with the Laser Method, and the second with the Group-theoretic method. The first technique limitation result was by Alon, Shpilka and Umans [ASU13] who showed that both conjectures would contradict the widely believed Sunower conjecture of Erdős and Rado.

Ambainis, Filmus and Le Gall [AFLG15] formalized the specific implementation of the Laser Method proposed by Coppersmith and Winograd [CW90] which is used in the recent papers on MM. They gave limitations of this implementation, and in particular showed that the exact approach used in [CW90, DS13, LG14, Wil12] cannot achieve a bound on ω better than 2.3078. The analyzed approach, the Laser Method with Merging, is a bit more general than the approaches in [CW90, DS13, LG14, Wil12]: in a sense it corresponds to a dream implementation of the exact approach.

Blasiak et al. [BCC⁺17a] considered the Group-theoretic Method for developing MM algorithms proposed by Cohn and Umans [CU03], and showed that this approach cannot prove $\omega = 2$ using any fixed abelian group. In follow-up work, Sawin [Saw18] extended this to any fixed non-abelian group, and Blasiak et al. [BCC17b] extended it to a host of families of non-abelian groups.

All these limitations proven so far are for very specific attacks on proving $\omega = 2$. While the proofs of [AFLG15] apply directly to CW_q , they only apply to the restricted Laser Method with Merging, and no longer apply to slight changes to this. The proofs in [BCC⁺ 17a] and [BCC 17b] are tailored to the Group-theoretic Method and do not apply (for instance) to the Laser Method on non-group tensors.

3.1 Our Results

3.1.1 The Universal Method

The key to Strassen's algorithm is an algebraic identity showing how 2×2 matrix multiplication can be computed surprisingly efficiently. In particular, Strassen showed that the $2 \times 2 \times 2$ matrix multiplication tensor has tensor rank at most 7; see the beginning of Chapter 4 for precise definitions). Arguing about the ranks of larger matrix multiplication tensors has proven to be quite difficult in fact, even the rank of the $3 \times 3 \times 3$ matrix multiplication tensor isn't currently known. Progress on bounding ω since Strassen's algorithm has thus taken the following approach: Pick a tensor (trilinear form) T , typically not a matrix multiplication tensor, such that

Powers T^n of T can be efficiently computed (i.e. T has low asymptotic rank), and

T is useful for performing matrix multiplication, since large matrix multiplication tensors can be 'reduced to' powers of T .

Combined, these give an upper bound on the rank of matrix multiplication itself, and hence ω .

In Chapter 4, we define a new very general method for analyzing tensors to give MM algorithms, which we call the Universal Method. In the Universal Method, the notion of 'reduction' between tensors we use is degeneration. Degenerations are the most general type of reduction known to preserve the ranks of tensors as required for the above approach. In other words, the Universal Method captures the most general sense in which one can use T to design MM algorithms. We write $\omega_u(T)$ to denote the best bound one can prove on ω by applying the Universal Method to T .

We also define two weaker methods: the Galactic Method applied to T , in which the 'reduction' must be a more restrictive monomial degeneration resulting in the bound $\omega_g(T)$ on ω , and the Solar Method applied to T , in which the 'reduction' must be an even more restrictive zeroing out resulting in the bound $\omega_s(T)$ on ω . Since monomial degenerations and zeroing outs are successively more restrictive types of degenerations, we have that for all tensors T ,

$$\omega_u(T) \geq \omega_g(T) \geq \omega_s(T):$$

These methods are very general; there are no known methods for computing $\omega_u(T)$, $\omega_g(T)$, or $\omega_s(T)$ for a given tensor T , and these quantities are even unknown for very well-studied tensors T .

The two main approaches to designing matrix multiplication algorithms are the Laser Method of Strassen [Str87] and the Group-Theoretic Method of Cohn and Umans [CU03]. Both of these approaches show how to give upper bounds $U_q(T)$ for particular structured tensors T (and hence upper bound itself). We will conclude Chapter 4 by giving an overview of these two methods, and describing how they are special cases of the Solar Method, emphasising that the Universal Method affords an algorithm designer much more power than is taken advantage of by these methods.

3.1.2 Limits on the Universal Method

The Coppersmith-Winograd Tensor

Both the Laser Method and the Group-theoretic Method give ways to find zeroing outs of tensors into matrix multiplication tensors, but not necessarily the best zeroing outs. In fact, it is known that the Laser Method does not always give the best zeroing out for a particular tensor T , since the improvements from [CW90] to later works [DS13, Wil12, LG14] can be seen as giving slight improvements to the Laser Method to find better and better zeroing outs¹. The Group-Theoretic Method, like the Solar Method, is very general, and it is not clear how to optimally apply it to a particular group or family of groups.

All of the improvements on bounding U_q for the past 30+ years have come from studying the Coppersmith-Winograd family of tensors CW_q $q \geq 2$. The Laser Method applied to powers of CW_5 gives the bound $U_5(CW_5) = 2.3729$. The Group-Theoretic Method can also prove the best known bound $U_5(CW_5) = 2.3729$ by simulating the Laser Method analysis of CW_q (see e.g. [AW18a] for more details). Despite a long line of work on matrix multiplication, there are no known tensors² which seem to come close to achieving the bounds one can obtain using CW_q . This leads to the first main question of this Part of the dissertation:

Question 3.1. How much can we improve our bound on U_q using a more clever analysis of the Coppersmith-Winograd tensor?

To resolve Question 3.1, we prove a new lower bound for the Coppersmith-Winograd tensor in Chapter 5:

Theorem 3.1. $U_q(CW_q) \geq 2.16805$ for all q .

Thus, if one starts with the CW tensor which has led to all improvements on U_q for the last 30+ years, even if one uses the Universal method which vastly generalizes all known approaches, one cannot prove a better bound than 2.16805 on U_q . We also give stronger lower bounds for particular tensors

¹These works apply the Laser Method to higher powers of the tensor $T = CW_q$, a technique which is still captured by the Solar Method.

²The author and Vassilevska Williams [AW18b] study a generalization of CW_q which can tie the best known bound, but its analysis is identical to that of CW_q . Our lower bounds in this paper will apply equally well to this generalized class as to CW_q itself.

in the family. For instance, for the specific tensor CW_5 which yields the current best bound on $\beta_u(CW_5) = 2.21912$:

Our proof of Theorem 3.1 proceeds by upper bounding $\beta_u(CW_q)$, the asymptotic slice rank of CW_q ; we will show that for any tensor T , non-trivial upper bounds on $\beta_u(T)$ imply non-trivial lower bounds on $\beta_u(T)$. The slice rank of a tensor, denoted $\beta_u(T)$, was first introduced by Blasiak et al. [BCC⁺17a] in the context of lower bounds against the Group-Theoretic Method. In order to study degenerations of powers of tensors, rather than just tensors themselves, we need to study an asymptotic version of slice rank, β_u . This is important since the slice rank of a product of two tensors can be greater than the product of their slice ranks, and as we will show $\beta_u(CW_q^n)$ is much greater than $\beta_u(CW_q)^n$ for big enough n .

We will give three different tools for proving upper bounds on $\beta_u(T)$ for many different tensors T . They will imply our lower bound on the Universal Method for CW_q as well as many other tensors of interest, including: the same lower bound $\beta_u(CW_q) = 2.16805$ for any generalized Coppersmith-Winograd tensor CW_q ; (a new class of tensors we define which slightly modify the structure of CW_q), a similar lower bound for cw_q , the generalized 'simple' Coppersmith-Winograd tensor missing its 'corner terms', and a lower bound for $\bar{c}w_q$, the structural tensor of the cyclic group C_q , matching the lower bounds obtained by [BCC⁺17a]. In Section 5.5 we give tables of our precise lower bounds for these and other tensors.

We briefly note that our lower bound of $2.16805 > 2 + \frac{1}{6}$ in Theorem 3.1 may be significant when compared to the recent algorithm of Cohen, Lee and Song [CLS18] which solves n -variable linear programs in time about $O(n^{2+1/6})$.

The Laser Method is Complete

The second main question of this part concerns the Laser Method. The Laser Method upper bounds $\beta_s(T)$ for any tensor T with certain structure (which we describe in detail in Section 5.4), and has led to every improvement on $\beta_s(CW_q)$ since its introduction by Strassen [Str87].

Question 3.2. When the Laser Method applies to a tensor T , how close does it come to optimally analyzing T ?

We call T laser-ready if the Laser Method (as used by [CW90] on CW_q) applies to it; see Definition 5.1 for the precise definition. Tensors need certain structure to be laser-ready, but tensors \bar{T} with this structure are essentially the only ones for which successful techniques for upper bounding $\beta_u(T)$ are known. In fact, every record-holding tensor in the history of matrix multiplication algorithm design has been laser-ready.

As discussed, we know the Laser Method does not always give a tight bound on $\beta_s(T)$ for laser-ready T . For instance, Coppersmith-Winograd [CW90] applied the Laser Method to CW_q to prove $\beta_s(CW_q) = 2.376$ and then later work [DS13, Wil12, LG14] analyzed higher and higher powers of CW_q to show $\beta_s(CW_q) = 2.373$. Ambainis, Filmus and Le Gall [AFLG15] showed that analyzing higher and higher powers of CW_q itself with the Laser Method cannot yield an upper bound better than

$\beta_s(CW_q) \approx 2.3725$ What about for other tensors? Could there be a tensor such that applying the Laser Method to T yields $\beta_s(T) = c$ for some $c > 2$, but applying the Laser Method to high powers T^n of T yields $\beta_s(T) = 2$? Could applying an entirely different method to such a T , using arbitrary degenerations and not just zeroing outs, show that $\beta_u(T) = 2$?

We show that for any laser-ready tensor T , the Laser Method can be used to prove a lower bound on $\beta(T)$. Moreover, we will see that this lower bound matches the upper bound on $\beta(T)$ implied by one of our tools, Theorem 5.3. We will use this to give an intriguing answer to Question 3.2:

Theorem 3.2. If T is a laser-ready tensor, and the Laser Method applied to T yields the bound $\beta_u(T) = c$ for some $c > 2$, then $\beta_u(T) > 2$.

To reiterate: If T is any tensor to which the Laser Method applies (as in Definition 5.1), and the Laser Method does not yield $\beta = 2$ when applied to T , then in fact $\beta_u(T) > 2$, and even the substantially more general Universal Method applied to T cannot yield $\beta = 2$. Hence, the Laser Method, which was originally used as an algorithmic tool, can also be seen as a lower bounding tool. Conversely, Theorem 3.2 shows that the Laser Method is complete, in the sense that it cannot yield a bound on β worse than 2 when applied to a tensor which is able to prove $\beta = 2$.

Another consequence of our proof is that, whenever T is a laser-ready tensor, we will be able to prove matching upper and lower bounds on $\beta(T)$. As mentioned, this includes every record-holding tensor in the history of MM algorithms, including CW_q , cw_q , and all the other tensors we study in Section 5.5. Hence, for these tensors no better lower bound on $\beta_u(T)$ is possible by arguing only about $\beta(T)$.

Theorem 3.2 explains and generalizes a number of phenomena:

The fact that Coppersmith-Winograd [CW90] applied the Laser Method to the tensor CW_q and achieved an upper bound greater than 2 on β implies that $\beta_u(CW_q) > 2$, and no arbitrary degeneration of powers of CW_q can yield $\beta = 2$.

As mentioned above, it is known that applying the Laser Method to higher and higher powers of a tensor T can successively improve the resulting upper bound on β . Theorem 3.2 shows that if the Laser Method applied to the first power of any tensor T did not yield $\beta = 2$, then this sequence of Laser Method applications (which is a special case of the Universal method) must converge to a value greater than 2 as well. This generalizes the result of Ambainis, Filmus and Le Gall [AFLG15], who proved this about applying the Laser Method to higher and higher powers of the specific tensor $T = CW_q$.

Our result also generalizes the result of Kleinberg, Speyer and Sawin [Kle97], where it was shown that (what can be seen as) the Laser Method achieves a tight lower bound on $\beta(T_q^{\text{lower}})$, matching the upper bound of Blasiak et al. [BCC⁺17a]. Indeed, T_q^{lower} , the lower triangular part of T_q , is a laser-ready tensor.

3.1.3 Additional Results

In our study of tensors and slice rank, we will also prove a number of new, complementary results.

Asymptotic Subrank Equals Asymptotic Slice Rank for Laser-Ready Tensors

Our proof of Theorem 3.2 also sheds light on a notion related to the asymptotic slice rank $\mathfrak{S}(T)$ of a tensor T , called the asymptotic subrank $\mathfrak{Q}(T)$ of T . \mathfrak{Q} is a dual notion of asymptotic rank, and it is important in the definition of Strassen's asymptotic spectrum of tensors [Str87]. While the asymptotic rank of T can be thought of as the 'cost' of T , the asymptotic subrank can be thought of as its 'value'.

It is not hard to see that $\mathfrak{Q}(T) \leq \mathfrak{S}(T)$ for all tensors T . However, there are no known separations between the two notions; whether there exists a tensor such that $\mathfrak{Q}(T) < \mathfrak{S}(T)$ is an open question. As a Corollary of Theorem 3.2, we prove:

Corollary 3.1. Every laser-ready tensor T has $\mathfrak{Q}(T) = \mathfrak{S}(T)$.

Since, as discussed above, almost all of the most-studied tensors are laser-ready, this might help explain why we have been unable to separate the two notions.

The Structural Tensors of Group Algebras

We also study the relationship between the generalized CW tensors and the structural tensors of group algebras (the tensors which arise in the Group-theoretic Method). Our new results include:

1. All Finite Groups Suffice for Current Ω Bounds. We show that every finite group G has a monomial degeneration to some generalized CW tensor of parameter $q = |G|$. Thus, applying the Galactic method on T_G for every G (with sufficiently small asymptotic rank, i.e. $\mathfrak{R}(T_G) = |G|$) can yield the current best bounds on Ω .
2. New Tri-Colored Sum-Free Set Constructions. Tri-Colored Sum-Free Sets are subsets of a group G which arise in extremal combinatorics. We show that, for every finite group G , there is a constant $c_{|G|} > 2/3$ depending only on $|G|$ such that its n th tensor power G^n has a tri-colored sum-free set of size at least $|G|^{c_{|G|} n^{o(n)}}$. For moderate $|G|$, the constant $c_{|G|}$ is quite a bit larger than $2/3$. To our knowledge, such a general result was not known until now.

3.2 Other Related Work

Probabilistic Tensors and Support Rank Cohn and Umans [CU13] introduced the notion of the support rank of tensors, and showed that upper bounds on the support rank of matrix multiplication tensors can be used to design fast Boolean

matrix multiplication algorithms. Recently, Karppa and Kaski [KK19] used 'probabilistic tensors' as another way to design Boolean matrix multiplication algorithms.

In fact, our tools for proving asymptotic slice rank upper bounds can be used to prove lower bounds on these approaches as well. For instance, our results imply that finding a 'weighted' matrix multiplication tensor as a degeneration of a power $\mathcal{C}W_q$ (in order to prove a support rank upper bound) cannot result in a better exponent for Boolean matrix multiplication than 2:16805

This is because 'weighted' matrix multiplication tensors can degenerate into independent tensors just as large as their unweighted counterparts. Similarly, if a probabilistic tensor T is degenerated into a (probabilistic) matrix multiplication tensor, Karppa and Kaski show that this gives a corresponding support rank expression for matrix multiplication as well, and so upper bounds $\text{orS}(T)$ for any T in the support of T also result in lower bounds on this approach.

Rectangular Matrix Multiplication Our tools can also be used to prove lower bounds on approaches to designing rectangular matrix multiplication algorithms. For instance, the best known rectangular matrix multiplication algorithms [LGU17] show that powers of CW_q zero out into large rectangular matrix multiplication tensors. Using the fact that CW_q is variable-symmetric, this implies a corresponding upper bound on $\text{or}_u(CW_q)$, which our tools give a lower bound against; see Section 4.5 for details.

Slice Rank Upper Bounds Our limitation results critically make use of upper bounds on the asymptotic slice rank of CW_q and other tensors of interest. Slice rank was first introduced by Tao [Tao16] in a symmetric formulation of the recent proof of the capset bound [CLP17, EG17], which shows how to prove slice rank upper bounds using the 'polynomial method'. Since then, a number of papers have focused on proving slice rank upper bounds for many different tensors. Sawin and Tao [TS16, Proposition 6] show slice rank upper bounds by studying the combinatorics of the support of the power of a fixed tensor, and Naslund and Sawin [NS17] use that approach to study sunflower-free sets³; one of our slice rank bounding tools, Theorem 5.3, uses this type of approach applied to blocked tensors. Slice rank was first used in the context of matrix multiplication by Blasiak et al. [BCC⁺17a], and this line of work has led to more techniques for proving slice rank upper bounds, including connections to the notion of instability from geometric invariant theory [BCC17a], and a generalization of the polynomial method to the nonabelian setting [BCG7b].

Concurrent Work Building off of our work [AW18b], Christandl, Vrana and Zuidam [CVZ19] independently proved lower bounds oh_u , including a bound matching Theorem 3.1. Their bounds use the seemingly more complicated machinery of Strassen's asymptotic spectrum of tensors [Str91]. They thus phrase their results in

³In fact, the tensor T whose slice rank is bounded in [NS17, Section 3] can be viewed as a change of basis of a Generalized Simple Coppersmith-Winograd tensor $\mathcal{C}W_q$, which we study below in Section 5.5.2

terms of the asymptotic subrank $\mathcal{Q}(T)$ of tensors rather than the asymptotic slice rank $\mathcal{S}(T)$, and the fact that their bounds are often the same as ours is related to the fact we prove, in Corollary 3.1, that $\mathcal{Q}(T) = \mathcal{S}(T)$ for all of the tensors we study.

3.3 Bibliographic Details

This Part of the dissertation is based on the results in three previously published papers:

‘Further Limitations of the Known Approaches for Matrix Multiplication’ with Virginia Vassilevska Williams [AW18a], which appeared in ITCS 2018,

‘Limits on All Known (and Some Unknown) Approaches to Matrix Multiplication’ with Virginia Vassilevska Williams [AW18b], which appeared in FOCS 2018, and

‘Limits on the Universal Method for Matrix Multiplication’ [Alm19b], which appeared in CCC 2019 and won the Best Student Paper Award.

Chapter 4 primarily presents results from [AW18b], and Chapter 5 primarily presents results from [Alm19b], except that Subsection 5.5.4 and Section 5.6 come from [AW18b], and Theorem 5.5 follows the proof of [AW18a, Lemma 4.1].

Chapter 4

The Universal Method

In this chapter, we introduce the relevant notions and notation related to tensors and matrix multiplication (MM) algorithms. We will give an overview of the known approaches to designing MM algorithms, including the Laser Method and the Group-theoretic Method, and then we will define a new, vast generalization of these approaches which we call the Universal Method. In the next Chapter, we will prove new limitation results against the Universal Method.

4.1 Tensors and Tensor Rank

The mathematical objects of interest in the study of MM algorithms are called tensors (or 3-tensors). Recall that a $q \times r$ matrix M over a field F can be viewed in many equivalent ways, including:

a 2-dimensional grid of numbers $(M_{ij})_{i \in [q], j \in [r]} \in F^{q \times r}$,

a linear map $M : F^q \rightarrow F^r$,

a bilinear map $M : F^q \times F^r \rightarrow F$.

Analogously, a $q \times r \times s$ tensor T can be viewed in a number of different ways, including:

a hypermatrix, i.e. a 3-dimensional grid of numbers $(T_{ijk})_{i \in [q], j \in [r], k \in [s]} \in F^{q \times r \times s}$,

a bilinear map $T : F^q \times F^r \rightarrow F^s$,

a trilinear map $T : F^q \times F^r \times F^s \rightarrow F$.

In this dissertation, we will focus on the equivalent view of tensors which I find simplest: as a trilinear polynomial.

For sets $X = \{x_1, \dots, x_q\}$, $Y = \{y_1, \dots, y_r\}$, and $Z = \{z_1, \dots, z_s\}$ of formal variables, a tensor over $X; Y; Z$ is a trilinear form

$$T = \sum_{x_1 \in X; y_1 \in Y; z_1 \in Z} \dots \sum_{x_q \in X; y_q \in Y; z_q \in Z} \sum_{ijk} x_i y_j z_k;$$

as well.

Proposition 4.1 ([Str69]). For any constant $q \geq 2$, if $R(n; q; q) = r$ (over a field F), there is an algorithm which performs $n \times n$ matrix multiplication over F using $O(n^{\log_q(r)})$ field operations.

Proof. Since $R(n; q; q) = r$, we can write

$$R(n; q; q) = \sum_{i=1}^r \left(\sum_{i,j \in [q]} a_{ij}^{(i)} x_{ij} A_{ij} \right) \left(\sum_{j,k \in [q]} b_{jk}^{(i)} y_{jk} B_{jk} \right) \left(\sum_{k,i \in [q]} c_{ki}^{(i)} z_{ki} C_{ki} \right); \quad (4.3)$$

for some coefficients $a_{ij}^{(i)}; b_{jk}^{(i)}; c_{ki}^{(i)} \in F$.

We design a recursive algorithm for multiplying $A, B \in F^{n \times n}$. We may assume that n is a power of q by padding the input matrices with 0s, which increases the dimension n by less than a multiplicative factor of q .

First, we partition A into a $q \times q$ block matrix, where each block is $n/q \times n/q$ matrix; call the blocks A_{ij} for $i, j \in [q]$. We similarly partition B into a $q \times q$ block matrix, and call the blocks B_{jk} for $j, k \in [q]$. The algorithm first computes, for each $i \in [q]$, the linear combination

$$A^0 = \sum_{i,j \in [q]} a_{ij}^{(i)} A_{ij};$$

and the linear combination

$$B^0 = \sum_{j,k \in [q]} b_{jk}^{(i)} B_{jk};$$

Next, for each $i \in [q]$, the algorithm computes the $(n/q) \times (n/q)$ matrix $C^0 := A^0 B^0$, by recursively performing $(n/q) \times (n/q) \times (n/q)$ matrix multiplication. Finally, for each $i, k \in [q]$, the algorithm computes the linear combination

$$C_{ki} = \sum_{k,i \in [q]} c_{ki}^{(i)} C_{ki}^0;$$

These are the blocks of the $n \times n$ matrix C which we output; indeed, we can see from (4.3) that for all $k, i \in [q]$, the matrix C_{ki} is the coefficient of z_{ki} in $R(n; q; q)$ when the substitutions $x_{ij} = A_{ij}$ and $y_{jk} = B_{jk}$ are made, and from the definition of $R(n; q; q)$ that these are exactly the desired output blocks.

Throughout the algorithm, we performed $O(n^2)$ field operations to compute linear combinations when constructing the A^0, B^0 , and C_{ki} matrices, and we performed recursive $(n/q) \times (n/q) \times (n/q)$ matrix multiplications. Thus, the total number $T(n)$ of field operations satisfies

$$T(n) = r T(n/q) + O(n^2);$$

which solves¹ to $T(n) = O(n^{\log_q(r)})$. □

¹Here we use the known bound $r > q^2$ [CW82, BI13] to avoid additional $\log(n)$ factors.

Note that in Proposition 4.1, we gave a bound on the number of field operations performed by the algorithm, rather than on the running time of the algorithm. Over a field like F_2 where field operations can be performed in constant time, this distinction is unimportant, but over larger fields, the field operations may take super-constant time and contribute to the total running time. Throughout Part I of this dissertation, we will abstract away this issue by focusing on a model of computation, such as the arithmetic circuit model, where field operations can be performed in constant time. Later, in Part II, we will be multiplying matrices of integers, and we will need to return to this issue.

In light of Proposition 4.1, we define ω , the exponent of matrix multiplication, as

$$\omega := \liminf_{q \geq 2} \log_q(R(nq; q; q)):$$

It follows from Proposition 4.1 that, for any $\epsilon > 0$, $n \times n \times n$ matrix multiplication can be performed in $O(n^{1+\epsilon})$ field operations. In fact, it is known that in the arithmetic circuit model, any algorithm for MM (which does not necessarily come from a tensor rank upper bound, and which may use division over \mathbb{F}) can be converted into a tensor rank upper bound which yields asymptotically the same running time when combined with Proposition 4.1 (see e.g. [Blä13, Theorem 4.7]). Hence, ω exactly captures the arithmetic circuit complexity of MM, and so our goal for designing MM algorithms is to give upper bounds on the ranks of MM tensors.

Before moving on, we make two notes about the definition of ω . First, using a \liminf rather than just a \min (which would, for instance, allow us to omit the $\epsilon > 0$ in the previous paragraph) is known to be required: Coppersmith and Winograd [CW82] showed that ω is a limit point that cannot be achieved by any single algorithm. Second, our notation in defining ω is somewhat sloppy, since the $\text{rank}(nq; q; q)$ may depend on the field F of coefficients. It is known that ω only depends on the characteristic of F [Sch81], but for instance, it may be the case that ω over F_2 is different from ω over \mathbb{C} . That said, the best known upper bounds on ω hold equally well for all fields, so we will simply refer to ω without reference to the field F for simplicity.

What bounds on the ranks of MM tensors are known? Strassen [Str69] showed in 1969 that $R(7; 2; 2) = 7$, yielding $\omega \leq \log_2(7) \approx 2.81$ (see Figure 4-1 below).

The next improved bound came in 1978, when Victor Pan [Pan78] showed that $R(70; 70; 70) = 143640$ yielding $\omega \leq \log_2(70) \approx 2.80$. Since then, a long line of work has led to the best known bound $\omega \leq 2.372864$ [CW82, DS13, Wil12, LG14], which comes from a bound on $\text{rank}(nq; q; q)$ for a very large value of q .

In order to handle rank expressions for such large tensors, the known approaches to designing matrix multiplication algorithms make use of two key techniques which we describe next: tensor powers which allow us to prove properties of large tensors by arguing only about smaller tensors, and rank-preserving reductions between tensors, which allow us to argue about tensors other than MM tensors as long as we can find a reduction from MM to T.

$$\begin{aligned}
h_2; 2; 2i &= (x_{11} + x_{22})(y_{11} + y_{22})(z_{11} + z_{22}) \\
&+ (x_{21} + x_{22})y_{11}(z_{21} - z_{22}) \\
&+ x_{11}(y_{12} - y_{22})(z_{12} + z_{22}) \\
&+ x_{22}(y_{21} - y_{11})(z_{11} + z_{21}) \\
&+ (x_{11} + x_{12})y_{22}(z_{12} - z_{11}) \\
&+ (x_{21} - x_{11})(y_{11} + y_{12})z_{22} \\
&+ (x_{12} - x_{22})(y_{21} + y_{22})z_{11}
\end{aligned}$$

Figure 4-1: Strassen's algorithm as a rank expression, showing that $h_2; 2; 2i$ is a sum of 7 rank one tensors, and thus $R(h_2; 2; 2i) = 7$.

4.3 Tensor Powers and Asymptotic Rank

We first introduce the tensor product. If T_1 is a tensor over $X_1; Y_1; Z_1$, and T_2 is a tensor over $X_2; Y_2; Z_2$, then the tensor product $T_1 \otimes T_2$ is a tensor over $X_1 \times X_2; Y_1 \times Y_2; Z_1 \times Z_2$ such that, for any $(x_1; x_2) \in X_1 \times X_2$, $(y_1; y_2) \in Y_1 \times Y_2$, and $(z_1; z_2) \in Z_1 \times Z_2$, the coefficient of $(x_1; x_2)(y_1; y_2)(z_1; z_2)$ in $T_1 \otimes T_2$ is the product of the coefficient of $x_1 y_1 z_1$ in T_1 , and the coefficient of $x_2 y_2 z_2$ in T_2 . In other words, if

$$\begin{aligned}
T_1 &= \sum_{x_1 y_1 z_1} x_1 y_1 z_1; & \text{and} & & T_2 &= \sum_{x_2 y_2 z_2} x_2 y_2 z_2; \\
&\begin{matrix} X \\ x_1^2 X_1 \\ y_1^2 Y_1 \\ z_1^2 Z_1 \end{matrix} & & & & \begin{matrix} X \\ x_2^2 X_2 \\ y_2^2 Y_2 \\ z_2^2 Z_2 \end{matrix} \\
\text{then, } T_1 \otimes T_2 &= \sum_{x_1 y_1 z_1 \ x_2 y_2 z_2} (x_1; x_2)(y_1; y_2)(z_1; z_2): \\
&\begin{matrix} X \\ (x_1; x_2)^2 X_1 \times X_2 \\ (y_1; y_2)^2 Y_1 \times Y_2 \\ (z_1; z_2)^2 Z_1 \times Z_2 \end{matrix}
\end{aligned}$$

The tensor product $T_1 \otimes T_2$ is exactly the product of T_1 and T_2 as polynomials, except that instead of viewing the result as a degree 6 polynomial, we continue to view it as a degree 3 polynomial by merging variables $x_1 x_2$ and similarly for the y and z variables.

The n th tensor power of a tensor A , denoted $A^{\otimes n}$, is the result of taking the tensor product of n copies of A together, so $A^{\otimes 1} = A$, and $A^{\otimes n} = A \otimes A^{\otimes (n-1)}$. Hence, if T is over $X; Y; Z$, then $T^{\otimes n}$ is over $X^n; Y^n; Z^n$, and its variables are n -tuples of variables from T . We will use this view in some of our proofs in Chapter 5.

Tensor products interact very nicely with the notions and tensors we have already introduced. First, for $a_1; a_2; b_1; b_2; c_1; c_2 \in \mathbb{N}$, we have $h_{a_1; b_1; c_1} \otimes h_{a_2; b_2; c_2} = h_{a_1 a_2; b_1 b_2; c_1 c_2}$. This corresponds to performing a $a_1 a_2 \times b_1 b_2 \times c_1 c_2$ matrix multiplication using block matrices, reducing the problem to the multiplication of $a_1 \times b_1 \times c_1$ matrices whose entries are $a_2 \times b_2$ and $b_2 \times c_2$ matrices. In particular, it follows that for all $q; n \in \mathbb{N}$, we have $h_{q; q; q}^{\otimes n} = h_{q^n; q^n; q^n}$.

Second, for any two tensors $A; B$, we always have $R(A \otimes B) = R(A) + R(B)$.

This follows by combining the distributive property with the fact that the product of two rank one tensors also has rank one. However, this inequality is often not tight. For instance, it is known that $R(h_2; 2; 2i) = 7$ (the lower bound was shown by Winograd [Win71]), but for large n , we can see that $R(h_2; 2; 2i^n) \sim 2^{ln + o(n)}$.

This motivates defining the asymptotic rank² of a tensor T as

$$\mathbb{R}(T) := \liminf_{n \geq N} (R(T^{\otimes n}))^{1/n}$$

Because of the tensor product properties above, we can alternatively define a number of ways:

$$\mathbb{R}(T) = \liminf_{q \geq 2N} \log_q R(h_q; q; q) = \liminf_{q \geq 2N} \log_q \mathbb{R}(h_q; q; q) = \log_2(\mathbb{R}(h_2; 2; 2i)):$$

4.4 Reductions Between Tensors

We now describe four different ways to reduce between tensors. The key property we would like from a type of reduction is that, if A reduces to B , then $\mathbb{R}(A) \leq \mathbb{R}(B)$. Thus, if we can show that a MM tensor reduces to some tensor \bar{T} , then in order to prove upper bounds on \mathbb{R} , it suffices to prove upper bounds on $\mathbb{R}(\bar{T})$. The four types of reductions we will define are summarized in Figure 4-2.

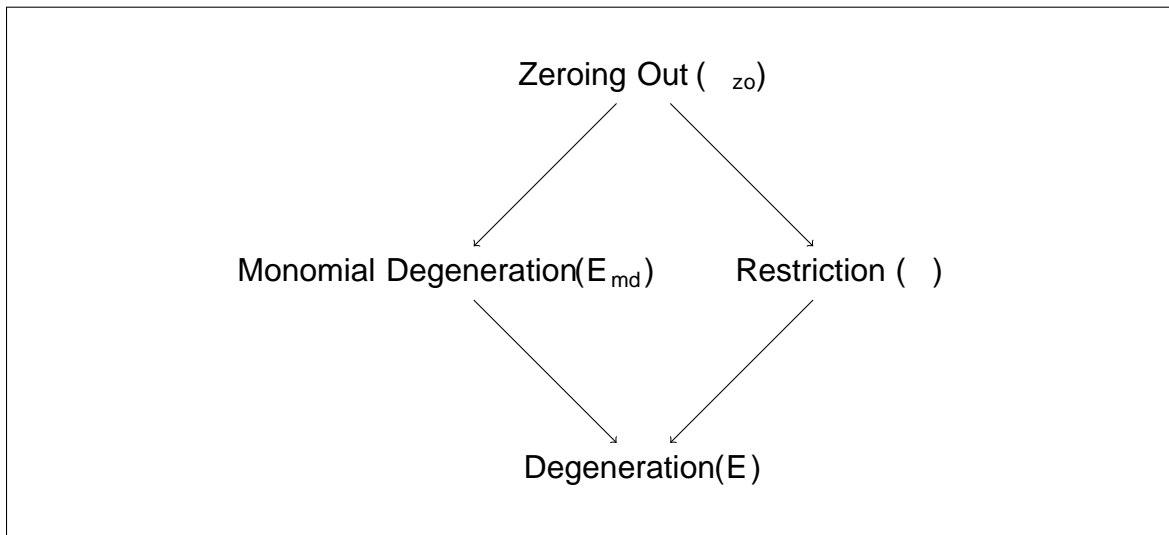


Figure 4-2: The four notions of a reduction between tensors. An arrow $\alpha \rightarrow \beta$ means that β subsumes α , meaning that for any tensors $A; B$, if $A \alpha B$, then $A \beta B$.

²The asymptotic rank is often written as \tilde{R} in the literature, but we instead write \mathbb{R} for ease of notation.

Zeroing Out The simplest type of reduction is a zeroing out (also called combinatorial restriction). Let B be a tensor over $X; Y; Z$, and A be a tensor over $X^0; Y^0; Z^0$ where $X^0 \subseteq X, Y^0 \subseteq Y$, and $Z^0 \subseteq Z$. A is a zeroing out of B , denoted $A \stackrel{zo}{\leftarrow} B$, if for all $x \in X^0, y \in Y^0$, and $z \in Z^0$, the coefficient of xyz in A equals the coefficient of xyz in B . In other words, A is obtained by setting to zero all $x \in X \setminus X^0, y \in Y \setminus Y^0$, and $z \in Z \setminus Z^0$. It is not hard to see that if $A \stackrel{zo}{\leftarrow} B$ then $R(A) \leq R(B)$ and $\mathcal{R}(A) \leq \mathcal{R}(B)$, by applying the same zeroing out to the (asymptotic) rank expression for B .

Example 4.1. For the tensors

$$B = x_0y_0z_0 + x_1y_1z_0 + x_1y_0z_1 + x_0y_1z_1;$$

$$A = x_0y_0z_0 + x_1y_1z_0;$$

we see that $A \stackrel{zo}{\leftarrow} B$ by setting $z_1 = 0$ in B , since the terms in A are exactly the terms in B that do not contain z_1 . By comparison, for the tensor

$$C = x_0y_0z_0 + x_1y_1z_0 + x_1y_0z_1;$$

there is no zeroing out from B to C , since the term $x_0y_1z_1$ from B we would like to remove shares each of its variables with a term in C that we need to keep.

Although zeroing outs are quite simple, we will see that the best known approaches to designing MM algorithms only make use of zeroing outs to reduce to MM tensors, rather than the following more powerful methods.

Monomial Degeneration Let $A; B$ be tensors over $X; Y; Z$. We say that A is a monomial degeneration of B , denoted $A \stackrel{E_{md}}{\leftarrow} B$, if the following type of transformation from B to A is possible. For a formal variable w , let $\text{Mon} := \{f^p \mid p \in \mathbb{Z}^0\}$. Pick a map $m : X \cup Y \cup Z \rightarrow \text{Mon}$, then from the tensor

$$B = \sum_{x \in X, y \in Y, z \in Z} x y z;$$

with coefficients from F , consider the transformed tensor

$$B^0 = \sum_{x \in X, y \in Y, z \in Z} m(x) m(y) m(z) x y z;$$

with coefficients from $F[w]$. B^0 can alternatively be viewed as a polynomial in w whose coefficients are tensors over $X; Y; Z$ with coefficients from F . If h is the smallest integer for which the coefficient of w^h in B^0 is nonzero, and A is the coefficient of w^h in B^0 , then this is a monomial degeneration from B to A .

If $A \stackrel{zo}{\leftarrow} B$, then $A \stackrel{E_{md}}{\leftarrow} B$ as well: if A is a zeroing out of B by setting the variables in $S \subseteq X \cup Y \cup Z$ to zero, then A is also a monomial degeneration of B by picking $m(w) = 0$ for all $w \in S$, and $m(w) = 1$ for all $w \in X \cup Y \cup Z \setminus S$, since then A will be the constant coefficient of the resulting B^0 . We will see in Example 4.2 below

that monomial degeneration strictly subsumes zeroing out. It is not evident that monomial degenerations should preserve any notion of tensor rank, but Bini [Bin80] showed that for any tensors A, B , if $A \xrightarrow{\text{md}} B$, then $\mathbb{R}(A) \geq \mathbb{R}(B)$.

Example 4.2. For the tensors

$$\begin{aligned} B &= x_0 y_0 z_0 + x_1 y_1 z_0 + x_1 y_0 z_1 + x_0 y_1 z_1; \\ C &= x_0 y_0 z_0 + x_1 y_1 z_0 + x_1 y_0 z_1; \end{aligned}$$

we see that $C \xrightarrow{\text{md}} B$ by picking $m(x_0) = m(y_1) = m(z_1) = 0$, and $m(x_1) = m(y_0) = m(z_0) = 1$. The resulting tensor B^0 is

$$B^0 = x_0 y_0 z_0 + x_1 y_1 z_0 + x_1 y_0 z_1 + 0 x_0 y_1 z_1;$$

so we have $B^0 = C + 0T$ for some tensor T . B and C are the same tensors from Example 4.1, showing that monomial degeneration is strictly more powerful than zeroing out.

Restriction Let A be a tensor over X^0, Y^0, Z^0 , and B be a tensor over X, Y, Z . We say A is a restriction of B , denoted $A \leq B$, if there are linear maps $M_X : F^X \rightarrow F^{X^0}$, $M_Y : F^Y \rightarrow F^{Y^0}$, $M_Z : F^Z \rightarrow F^{Z^0}$ such that $A = B(M_X; M_Y; M_Z)$. In other words, if

$$B = \sum_{x \in X, y \in Y, z \in Z} x y z;$$

then

$$A = \sum_{x \in X, y \in Y, z \in Z} M_X(x) M_Y(y) M_Z(z);$$

Unlike in zeroing out and monomial degenerations, we do not need the $\text{rank}(X) \leq \text{rank}(X^0)$, and it might even be the case that $\text{rank}(X^0) > \text{rank}(X)$. If $A \leq B$, then $A \leq B$, since a zeroing out corresponds to a restriction where, for each $x \in X$, we either pick $M_X(x) = x$ or $M_X(x) = 0$ (and similarly for Y and Z).

Example 4.3. For the tensors

$$\begin{aligned} D &= x_+ y_+ z_+ + x_- y_- z_-; \\ B &= x_0 y_0 z_0 + x_1 y_1 z_0 + x_1 y_0 z_1 + x_0 y_1 z_1; \end{aligned}$$

we see that $B \xrightarrow{\text{md}} D$ by picking

$$\begin{aligned} x_+ &= x_0 + x_1; & y_+ &= y_0 + y_1; & z_+ &= \frac{1}{2}(z_0 + z_1); \\ x_- &= x_0 - x_1; & y_- &= y_0 - y_1; & z_- &= \frac{1}{2}(z_0 - z_1); \end{aligned}$$

The resulting transformation of D is

$$\frac{(x_0 + x_1)(y_0 + y_1)(z_0 + z_1) + (x_0 - x_1)(y_0 - y_1)(z_0 - z_1)}{2},$$

which when expanded gives

The transformation of D in Example 4.3 is the sum of two rank one tensors, showing that B has rank at most 2. In fact, restrictions can be used in this way to exactly characterize rank.

For $q \geq 2$, write $h_{q,i} := \sum_{i=1}^q x_i y_i z_i$ for the independent tensor of size q , which has q terms which do not share any variables with each other. For instance, in Example 4.3, $D = h_{2,i}$. Since $h_{1,i}$ can be restricted to any rank one tensor, and restrictions act independently on each term of $h_{q,i}$, we see that:

Proposition 4.2. For any $q \geq 2$ and tensor T , there is a restriction $T \leq h_{q,i}$ if and only if $R(T) \leq q$.

We can see that $A \leq B \leq C$ implies that $A \leq C$ by simply composing the corresponding linear maps; it follows from Proposition 4.2 that $A \leq B$ implies $R(A) \leq R(B)$.

Proposition 4.2 illustrates how difficult it can be to determine whether there is a restriction between two given tensors, since tensor rank is hard to compute even for small explicit tensors. For instance, although $R(h_{2,2}) = 7$ is known, determining the value of $R(h_{3,3})$ is open. The best known upper bound is $R(h_{3,3}) \leq 23$, and so we do not know whether there is a restriction $h_{3,3} \leq h_{2,2}$. For more on the computational difficulty of determining whether there is a restriction between given tensors, see e.g. [GQ19].

Before continuing, we note that the independent tensor $h_{q,i}$ will appear again numerous times throughout this dissertation. For one example, it can be used to capture the disjoint sum of a tensor with itself:

Definition 4.1. For any tensors T_1 over $X_1; Y_1; Z_1$ and T_2 over $X_2; Y_2; Z_2$, their direct sum $T_1 \oplus T_2$ is a tensor over $X_1 \sqcup X_2; Y_1 \sqcup Y_2; Z_1 \sqcup Z_2$ given by the sum (as polynomials) of T_1 and T_2 . For $q \geq 2$ and tensor T , we write $q \cdot T$ for the disjoint sum of q copies of T . Note that $q \cdot T = h_{q,i} \cdot T$.

Degeneration The most powerful known asymptotic rank-preserving reduction between tensors is a degeneration³. It combines the power of the formal variable from monomial degenerations with the linear transformations from restrictions.

Let A be a tensor over $X^0; Y^0; Z^0$, and B be a tensor over $X; Y; Z$. We say A is a degeneration of B , denoted $A \leq B$, if there are linear maps $M_X : F^X \rightarrow F[X^0]$, $M_Y : F^Y \rightarrow F[Y^0]$, $M_Z : F^Z \rightarrow F[Z^0]$, whose ranges are linear combinations of the variables of A whose coefficients are polynomials in t such that: when $B^0 = B(M_X; M_Y; M_Z)$

³Slightly more powerful notions, like degenerations with multiple variables, can be captured by straightforward modifications to the notion of degeneration which preserve all the results about degenerations which we prove below.

is viewed as a polynomial in x, y, z , and h is the smallest integer such that the coefficient of $x^h y^h z^h$ in B^0 is nonzero, then the coefficient of $x^h y^h z^h$ in B^0 is A .

Example 4.4. For the tensors

$$\begin{aligned} D &= x_0 y_0 z_0 + x_1 y_1 z_1; \\ C &= x_0 y_0 z_0 + x_1 y_1 z_0 + x_1 y_0 z_1; \\ B &= x_0 y_0 z_0 + x_1 y_1 z_0 + x_1 y_0 z_1 + x_0 y_1 z_1; \end{aligned}$$

we showed in Example 4.2 that $C \in E_{\text{md}} B$, and we showed in Example 4.3 that $D \in E C$. Composing the two transformations shows that $D \in E B$. Although composing a monomial degeneration and a restriction like this is one way to give a degeneration, there are more degenerations not captured by this approach as well.

The result of Bini [Bin80] implies that if $A \in E B$ then $\mathbb{R}(A) = \mathbb{R}(B)$, just as it did for monomial degenerations. In particular, if $A \in E_{\text{hqj}}$, then $\mathbb{R}(A) = q$.

4.5 The Universal Method

We now have all the ingredients in place to define the Universal Method for designing MM algorithms, which subsumes and greatly generalizes the known approaches for designing MM algorithms. The Universal Method applied to a tensor T consists of two components:

- (1) a bound $\mathbb{R}(T) \leq r$ on the asymptotic rank of T , and
- (2) a degeneration $\text{hqj}; q; \hat{q} \in T^{\otimes n}$, which reduces MM to a tensor power of T .

Combined, the two components imply that $\mathbb{R}(\text{hqj}; q; \hat{q}) \leq r^n$, and hence that $\mathbb{R}(T) \leq n \log_q r$. We write $\mathbb{R}_q(T)$ for the \liminf over all bounds on r which can be proved in this way, including picking r to be the true asymptotic rank of T , and for each n picking the largest q such that the degeneration required for step (2) exists.

The Asymptotic Sum Inequality? Readers familiar with Schönhage's Asymptotic Sum Inequality may wonder why step (2) in the Universal Method requires a degeneration to a single MM tensor rather than a disjoint sum of many. Indeed, Schönhage shows that degenerations to disjoint sums of MM tensors are sufficient to bound $\mathbb{R}(T)$:

Theorem 4.1 (Asymptotic Sum Inequality [Sch81]) If $\mathbb{R}(f \text{ hqj}; q; \hat{q}) \leq r$, then $\mathbb{R}(T) \leq \log(r/f) = \log(q)$.⁴

⁴Schönhage's Asymptotic Sum Inequality allows for a disjoint sum of MM tensors of many different shapes (values of q) as well, but the first step of its proof shows that, by taking large tensor powers, one can assume without loss of generality that all the MM tensors have the same shape.

It may seem like one could prove a better bound on β by degenerating to a disjoint sum of MM tensors rather than a single MM tensor. However, this technique is actually captured by the Universal Method as well; one can always find degenerations from powers of T to a single MM tensor which achieve an equally good bound on β :

Proposition 4.3. If T is a tensor with $\mathcal{R}(T) = r$ and $f \leq h, q; q; q \in T$, then $\beta_u(T) \leq \log(r=f) = \log(q)$.

Proof. By definition of β , for every $\epsilon > 0$, there is an $m \in \mathbb{N}$ such that $h, a; a; a \in T^m$ and $f \leq r^{m(1+\epsilon)} \leq f^{m(1+\epsilon)}$ (where the second inequality holds because $\mathcal{R}(T) \leq f$). In particular, we have that

$$T^m \otimes D \text{ } h, q^m; q^m; q^m \in D \text{ } h, aq^m; aq^m; aq^m;$$

which yields the bound

$$\beta_u(T) \leq \frac{\log(r^m)}{\log(aq^m)} = \frac{\log(r^m)}{\log(f^{m(1+\epsilon)} q^m)} = \frac{\log(r^{1+\epsilon})}{\log(f^{1+\epsilon} q)}.$$

Rearranging yields

$$f \leq q^{1+\epsilon} r^{1+\epsilon} = r^{1+\epsilon} q;$$

and hence

$$\beta_u(T) \leq \frac{\log(r^{1+\epsilon})}{\log(q)} = \frac{\log(r=f)}{\log(q)} + \epsilon = \frac{\log(r)}{\log(q)} + \epsilon.$$

Since this holds for all sufficiently small $\epsilon > 0$, it implies that $\beta_u(T) \leq \log(r=f) = \log(q)$ as desired. \square

Rectangular MM Tensors? Step (2) in the Universal Method requires a degeneration to a square MM tensor, but degenerations to rectangular MM tensors also give bounds on β . Indeed, if $\mathcal{R}(h, a; b; c) = r$, then by the symmetry of $h, a; b; c$ we also get that $\mathcal{R}(h, b; c; a) = r$ and $\mathcal{R}(h, c; a; b) = r$, which combined mean that $\mathcal{R}(h, a, b, c; a, b, c) = r^3$, and hence $\beta = 3 \log(r) = \log(ab^2c)$.

If tensor T is 'variable-symmetric', then a similar argument shows that, in the Universal Method applied to T , allowing for degenerations to rectangular MM tensors cannot help.

Definition 4.2. If T is a tensor over $X; Y; Z$, then the rotation of T , denoted $\text{rot}(T)$, is the tensor over $Y; Z; X$ such that for any $(x_i; y_j; z_k) \in X \times Y \times Z$, the coefficient of $x_i y_j z_k$ in T is equal to the coefficient of $y_j z_k x_i$ in $\text{rot}(T)$. Tensor T is variable-symmetric if $T = \text{rot}(T)$.

The symmetrized version of T , denoted $\text{sym}(T)$, is given by $\text{sym}(T) = T + \text{rot}(T) + \text{rot}(\text{rot}(T))$. For any tensor T , the tensor $\text{sym}(T)$ is always variable-symmetric. Moreover, if T was variable-symmetric, then $\text{sym}(T) = 3T$.

For any tensor T , if $\mathcal{R}(T) = r$ and $T^3 \in D \text{ } h, a; b; c$ to yield $\beta = 3 \log(r) = \log(ab^2c)$, then it follows that $\text{sym}(T)^3 \in D \text{ } h, abc; abc; abc$ meaning

$\rho_u(\text{sym}(T)) \leq 3n \log(r) = \log(ab^3)$. In other words, any bound on ρ which we could achieve if we allowed for rectangular MM tensors in step (2) of the Universal Method applied to T , can also be achieved by applying the Universal Method as written to $\text{sym}(T)$. Notably, almost every tensor we will consider in the remainder of this chapter is variable-symmetric, for which $\text{sym}(T) = T$.

In the previous paragraph we used the fact that $\rho(\text{sym}(T)) \leq \rho(T)^3 = r^3$, but if T is such that $\rho(\text{sym}(T)) < \rho(T)^3$ then we get an even better bound of $\rho_u(\text{sym}(T)) \leq n \log(\rho(\text{sym}(T))) = \log(ab^3) < 3n \log(r) = \log(ab^3)$. To prove the best bounds on ρ when using a tensor T which is not variable-symmetric, one should always apply the Universal Method to $\text{sym}(T)$ rather than T .

Example 4.5. For any $q \geq 2$, consider the tensor $T = \sum_{i,j,k} q_i q_j q_k$. A 'attenuating' argument⁵ shows that $\rho(T) = q$. However, $\text{sym}(T) = \sum_{i,j,k} q_i q_j q_k$, and so $\rho(\text{sym}(T)) = q^3 < q^3 = \rho(T)^3$ when $q > 1$.

4.5.1 The Solar and Galactic Methods

The Universal Method is very general, and it is typically unclear how to apply it optimally to a given tensor T . In fact, all known approaches to designing MM algorithms use a substantially restricted method, which uses zeroing outs instead of degenerations, which we call the Solar Method. The Solar Method applied to a tensor T consists of two components:

- (1) a bound $\rho(T) \leq r$ on the asymptotic rank of T , and
- (2) a zeroing out $\sum_{i,j,k} q_i q_j q_k \leq T^n$.

The resulting bound on ρ is that $\rho \leq \log(r^n) = \log(q)$, and the \liminf over all bounds on ρ which can be achieved in this way is denoted $\rho_s(T)$. Here, we need to allow for a zeroing out into a disjoint sum of multiple MM tensors in step (2), unlike in the Universal Method, since Proposition 4.3 critically makes use of degenerations, and not just zeroing outs, in the first step of the proof.

One can also define an intermediate method, the Galactic Method, which is identical to the Solar Method except that the zeroing out $\sum_{i,j,k} q_i q_j q_k$ in step (2) is replaced by the more powerful monomial degeneration $\sum_{i,j,k} q_i q_j q_k$, and the best resulting bound on ρ is denoted $\rho_g(T)$. One could also consider an incomparable intermediate method which makes use of restrictions instead of monomial degenerations, but since such a method has not been studied much, and is captured by the Universal Method, it doesn't yet have a name.

Because each of degeneration, monomial degeneration, and zeroing out (strictly) subsumes the previous, we get that for all tensors T ,

$$\rho(T) \leq \rho_u(T) \leq \rho_g(T) \leq \rho_s(T):$$

⁵When T is viewed as a matrix by removing the z -variable (i.e. setting $z_1 = 1$), then T^n becomes the $q^n \times q^n$ identity matrix, which has rank q^n . If it were the case that $\rho(T^n) < q^n$, then similarly removing z_1 from the rank expression would give an upper bound $\rho(T^n)$ on the rank of the $q^n \times q^n$ identity matrix as well.

To be clear, all three of these methods are very general, and we don't know the values of $\alpha_s(T)$, $\alpha_g(T)$, or $\alpha_u(T)$ for almost any nontrivial tensors T . In fact, all the known approaches to bounding α_s proceed by giving upper bounds on $\alpha_s(T)$ for some carefully chosen tensors \mathbb{T} ; the most successful has been the Coppersmith-Winograd family of tensors $T = CW_q$, which has yielded all the best known bounds on α_s since the 80's [CW90, DS13, Wil12, LG14]. In particular, it is not generally believed that our current upper bounds on $\alpha_s(CW_q)$ are optimal, since the techniques from [DS13, Wil12, LG14] seem able to further improve on the current bounds (by a small amount) with more effort.

Finally, we remark that the tensor T which these methods apply to is very crucial. The three different methods will trivially give the same bound, $\alpha_s(T) = \alpha_g(T) = \alpha_u(T) = \alpha$, when applied to $T = \mathbb{H}_{2;2;2}$ itself, but this is not particularly interesting: the point of these different methods is that the asymptotic rank of matrix multiplication tensors is not well-understood, but the methods allow us to prove bounds on α by studying other tensors.

4.6 The Known Approaches to Matrix Multiplication

We conclude this chapter by giving an overview of the two known approaches to designing MM algorithms: the Laser Method and the Group-theoretic Method. Each of these methods applies to particular tensors \mathbb{T} to yield bounds in $\alpha_s(T)$; in particular, we will see that the Solar Method is a generalization of both of these methods. We only give high-level overviews here, mostly just for context, and we recommend the original papers (cited below) for more details.

4.6.1 The Laser Method

Strassen [Str86] called his approach for reducing MM tensors to other tensors the Laser Method. In this method applied to a tensor T over $X; Y; Z$, we start by partitioning the variable sets: $X = X_1 \parallel \dots \parallel X_{k_X}$, $Y = Y_1 \parallel \dots \parallel Y_{k_Y}$, $Z = Z_1 \parallel \dots \parallel Z_{k_Z}$. For $i \in [k_X]; j \in [k_Y]; k \in [k_Z]$, let T_{ijk} be the sub-tensor of T obtained by zeroing out all variables $x \notin X_i$, $y \notin Y_j$, and $z \notin Z_k$; we call T_{ijk} a block of T . We thus obtain a partitioning

$$T = \sum_{i \in [k_X]; j \in [k_Y]; k \in [k_Z]} T_{ijk}.$$

Strassen originally considered tensors where the constituent tensors T_{ijk} are each MM tensors; we focus here on this setting, although later work showed how to remove this requirement.

The next step is to consider a large tensor power $T^{\otimes N}$ of T . We can write

$$T^{\otimes N} = \sum_{I \in \{1,2\}^{[k_X]^N}; J \in \{1,2\}^{[k_Y]^N}; K \in \{1,2\}^{[k_Z]^N}} T_{I,J,K}^{\otimes N} \quad (4.4)$$

Each of the terms on the right-hand side of (4.4) is a MM tensor, and the goal is to ultimately apply Schönhage's Asymptotic Sum Inequality (Theorem 4.1 above) to yield a bound on $\|T^{\otimes N}\|$. However, the sum on the right-hand side of (4.4) is not a disjoint sum, or in other words, the different MM tensors share variables with each other, so we cannot apply the Asymptotic Sum Inequality directly.

In the Laser Method, we seek to zero out some variables to fix this problem. We aim to choose some $I \in \{1,2\}^{[k_X]^N}$, and then for each $i \in I$, to zero out all the variables in $I \setminus \{i\}$ (and similarly for y and z variables). One must find such a zeroing out which leaves exactly a direct sum of matrix multiplication tensors. Finally, applying the Asymptotic Sum Inequality yields a bound on $\|T^{\otimes N}\|$.

We now turn to the most successful implementation of the Laser Method: the Coppersmith-Winograd approach. The Coppersmith-Winograd (CW) family of tensors $(CW_q)_{q \geq 2}$ is defined as:

$$CW_q = x_0 y_0 z_{q+1} + x_{q+1} y_0 z_0 + x_0 y_{q+1} z_0 + \sum_{i=1}^q (x_i y_0 z_i + x_0 y_i z_i + x_i y_i z_0)$$

CW_q is a tensor over $q+2$ x -variables, $q+2$ y -variables, and $q+2$ z -variables, and it is known that $\mathbb{R}(CW_q) = q+2$ (The upper bound $\mathbb{R}(CW_q) \leq q+2$ was a key contribution of [CW90], and the lower bound follows from a 'attenuating' argument).

Coppersmith and Winograd [CW90] followed the laser method. The terms of CW_q have a natural partitioning $CW_q = T_{002} + T_{020} + T_{200} + T_{011} + T_{101} + T_{110}$, where $T_{002} = x_0 y_0 z_{q+1}$; $T_{200} = x_{q+1} y_0 z_0$; $T_{020} = x_0 y_{q+1} z_0$; $T_{101} = \sum_{i=1}^q x_i y_0 z_i$; $T_{011} = \sum_{i=1}^q x_0 y_i z_i$; $T_{110} = \sum_{i=1}^q x_i y_i z_0$. This partitioning has three key properties. First, each of these parts is a MM tensor. Second, this partitioning arises from a block-partitioning of the variables where we partition $X = X_0 \parallel X_1 \parallel X_2$ where $X_0 = [x_0]$; $X_1 = [x_1; x_2; \dots; x_q]$; and $X_2 = [x_{q+1}]$; and similarly for Y and Z . Third, each sub-tensor T_{ijk} is non-zero if and only if $i + j + k = 2$.

The Coppersmith-Winograd implementation of the laser method uses these properties together with sets excluding 3-term arithmetic progressions (in conjunction with the third property above) to decide which blocks of variables to zero out in $CW_q^{\otimes N}$. The sets excluding 3-term arithmetic progressions can be used to guarantee that the result is a direct sum of many large matrix multiplication tensors, thus obtaining a bound on $\|CW_q^{\otimes N}\|$. Coppersmith and Winograd get a different bound on $\|CW_q^{\otimes N}\|$ for each q , and optimize it by picking $q = 6$. They then achieve a slightly better bound on $\|CW_q^{\otimes N}\|$ by analyzing the square $CW_q^{\otimes 2}$ in a similar way. In Theorem 5.5 in the next Chapter, we present all the details of Coppersmith and Winograd's application of the Laser Method.

The later improvements on the Coppersmith-Winograd bounds by

Stothers [DS13], Vassilevska Williams [Wil12] and Le Gall [LG14] instead used the laser method with the CW tools starting from CW_q^4 ; CW_q^8 and CW_q^{16} and CW_q^{32} , respectively. Each new analysis used different, but related, blockings and partitionings, and each ultimately optimizes the resulting bound on $\text{rank}(CW_q)$ by picking $q = 5$, and hence using CW_5 as the starting tensor. In other words, each ultimately gives an upper bound on $\text{rank}(CW_5)$. The constituent tensors of powers of CW_q are sometimes not MM tensors. To deal with this, one (recursively) performs a Coppersmith-Winograd analysis on the constituent tensors, showing that they, themselves, can zero out into large MM tensors at high enough tensor powers. As the power of CW_q which is considered grows, the number of recursive analyses needed becomes very large.

The Coppersmith-Winograd approach doesn't exploit very much about the constituent tensors T_{ijk} . In particular, the analysis remains unchanged if one replaces each T_{ijk} with another tensor T'_{ijk} over the same sets of variables $X_i; Y_j; Z_k$, as long as T'_{ijk} has the same value (i.e. can degenerate to equally large MM tensors in high tensor powers), and the modified tensor T' has the same asymptotic rank as T . In this case, the bound on the approach would give is exactly the same! For instance, when T_{ijk} is a matrix multiplication tensor $h; a; b; c$, one can replace it with another matrix multiplication tensor $h; a'; b'; c'$ as long as the new tensor uses the same variables and $a'b'c' = abc$ and as long as the asymptotic rank of the overall tensor has not increased. For instance, if we take $T_{110} = \sum_{i=1}^q \sum_{j=1}^q x_i y_j z_0$ in CW_q and replace it with $\sum_{i=1}^q \sum_{j=1}^q x_i y_{q+1-i} z_0$, then we get the rotated CW_q tensor studied in [AW18a]. This tensor still has asymptotic rank $q+2$, and thus gives the same upper bound on $\text{rank}(CW_q)$ using the CW approach.

We can thus define a family of generalized CW tensors:

Definition 4.3. The family \underline{CW}_q of tensors includes, for every permutation $\sigma \in S_q$, the tensor

$$CW_q = (x_0 y_0 z_{q+1} + x_0 y_{q+1} z_0 + x_{q+1} y_0 z_0) + \sum_{i=1}^q (x_i y_{\sigma(i)} z_0 + x_i y_0 z_i + x_0 y_i z_i):$$

The family above contains all tensors obtained from CW_q by replacing $\sum_{i=1}^q (x_i y_i z_0 + x_i y_0 z_i + x_0 y_i z_i)$ with $\sum_{i=1}^q (x_{\sigma(i)} y_{\sigma(i)} z_0 + x_{\sigma(i)} y_0 z_{\sigma(i)} + x_0 y_{\sigma(i)} z_{\sigma(i)})$ for any choice of $\sigma \in S_q$, by a simple renaming of variables.

The constituent tensor T_{110} of CW_q is $\sum_{i=1}^q x_i y_{\sigma(i)} z_0$, which is still a $h; q; 1$ tensor. Thus, for any such tensor from the family \underline{CW}_q , if its asymptotic rank is $q+2$, then the Coppersmith-Winograd approach would give exactly the same bound on $\text{rank}(CW_q)$, as with CW_q . Unfortunately, the asymptotic rank lower bounding technique applies equally well to any tensor in \underline{CW}_q , showing that they all have asymptotic rank at least $q+2$, so our upper bounds on $\text{rank}(CW_q)$ cannot be improved just in this way.

4.6.2 The Group-theoretic Method

Cohn and Umans [CU03] pioneered a new Group-theoretic Method for matrix multiplication, which works with a finite group G rather than directly with a tensor.

Roughly, they define properties of a group such that, if G has these properties, then it is possible to zero out an underlying tensor corresponding to G into a MM tensor.

Definition 4.4. For any finite group G , the group tensor of G , denoted T_G , is a tensor over $X_G; Y_G; Z_G$ where $X_G := \{x_g \mid g \in G\}$, $Y_G := \{y_g \mid g \in G\}$, and $Z_G := \{z_g \mid g \in G\}$, given by

$$T_G := \sum_{g,h \in G} x_g y_h z_{gh}.$$

(The group tensor of G is often called the structural tensor of the group algebra $\mathbb{C}[G]$, written as $T_{\mathbb{C}[G]}$, in the literature. We write T_G here for ease of notation.)

The Group-theoretic Method first bounds the asymptotic rank of T_G using representation theory, as follows. Let d_u be the dimension of the u th irreducible representation of G (i.e. the d_u s are the character degrees of G). Then T_G can be seen to restrict from $\sum_{u=1}^M d_u; d_u; d_u$. In particular, we get that⁶

$$\mathbb{R}(T_G) \leq \sum_{u=1}^M d_u = \sum_{u=1}^M d_u^3.$$

If we could find any degeneration (e.g. a zeroing out) of T_G into $h; q; q$, it would imply that

$$q \leq \sum_{u=1}^M d_u^3,$$

which gives an upper bound on h .

Cohn and Umans defined two properties of subsets S which yield a zeroing out of T_G into a MM tensor, called the 'triple product property', and the 'simultaneous triple product property' (which zeroes out into a disjoint sum of MM tensors). Hence, bounds on h can follow from showing that a group G has large subsets with these properties. Typically one works with a family of groups (as in A_n or S_n for all $n \in \mathbb{N}$), and then one can pick the one that optimizes the bound on h , or even taken $\lim_{n \rightarrow \infty} h$, e.g. when the groups correspond to tensor powers of some tensor. We refer the reader to [Lan17, Section 3.5] for more exposition on the Group-theoretic Method and its interpretation as finding a zeroing out of group tensors.

⁶It is more straightforward to see that this holds with inequalities (\leq instead of $=$) but in fact equality holds because the corresponding restriction of T_G is invertible.

Chapter 5

Limits on the Universal Method

5.1 Asymptotic Slice Rank

In this chapter, we will prove new limitation results against the Universal Method. Our limitations will critically make use of a variant on the rank of a tensor, called slice rank. We begin in this section by introducing slice rank and its key properties.

We say a tensor T over $X; Y; Z$ has x -rank one if it is of the form

$$T = \sum_{x \in X} \sum_{y \in Y} \sum_{z \in Z} \alpha_{xyz} x \otimes y \otimes z$$

for some choices of the α_{xyz} coefficients over F . More generally, the x -rank of T , denoted $S_x(T)$, is the minimum number of tensors of x -rank one whose sum is T . We can similarly define the y -rank, S_y , and the z -rank, S_z . Then, the slice rank of T , denoted $S(T)$, is the minimum k such that there are tensors T_X, T_Y and T_Z with $T = T_X + T_Y + T_Z$ and $S_x(T_X) + S_y(T_Y) + S_z(T_Z) = k$.

Unlike tensor rank, the slice-rank is not submultiplicative in general, i.e. there are tensors A and B such that $S(A \otimes B) > S(A) S(B)$. For instance, it is not hard to see that $S(CW_5) = 3$, but since it is known [Wil12, LG14] that $S(CW_5) \leq 2 \cdot 3^3$, it follows (e.g. from Theorem 5.1 below) that $S(CW_q^n) \leq 7^{n-2} \cdot 3^{3n} \cdot 5 \cdot 15^{n-o(n)}$. We are thus interested in the asymptotic slice rank $\mathfrak{S}(T)$, of tensors T , which is defined similarly to the asymptotic rank as

$$\mathfrak{S}(T) := \limsup_{n \rightarrow \infty} [S(T^{\otimes n})]^{1/n}$$

We note a few simple properties of slice rank which will be helpful in our proofs:

Lemma 5.1. For tensors A and B :

- (1) $S(A) \leq S_x(A) \leq R(A)$,
- (2) $S_x(A \otimes B) \leq S_x(A) S_x(B)$,
- (3) $S(A + B) \leq S(A) + S(B)$, and $S_x(A + B) \leq S_x(A) + S_x(B)$,

(4) $S(A \otimes B) = S(A) \max\{S_x(B); S_y(B); S_z(B)\}$, and

(5) If A is a tensor over $X; Y; Z$, then $S_x(A) = \sum_j X_j$ and hence $S(A) = \min\{\sum_j X_j; \sum_j Y_j; \sum_j Z_j\}$.

Proof. (1) and (2) are straightforward. (3) follows since the sum of the slice rank (resp. x-rank) expressions for A and for B gives a slice rank (resp. x-rank) expression for $A + B$. To prove (4), let $m = \max\{S_x(B); S_y(B); S_z(B)\}$, and note that if $A = A_x + A_y + A_z$ such that $S_x(A_x) + S_y(A_y) + S_z(A_z) = S(A)$, then

$$A \otimes B = A_x \otimes B + A_y \otimes B + A_z \otimes B;$$

and so

$$\begin{aligned} S(A \otimes B) &= S(A_x \otimes B) + S(A_y \otimes B) + S(A_z \otimes B) \\ &= S_x(A_x) S_x(B) + S_y(A_y) S_y(B) + S_z(A_z) S_z(B) \\ &= S_x(A_x) m + S_y(A_y) m + S_z(A_z) m = S(A) m; \end{aligned}$$

Finally, (5) follows since, for instance, any tensor with one only x-variable has x-rank 1. \square

5.2 Limits on the Universal Method from Asymptotic Slice Rank

Asymptotic slice rank is interesting in the context of MM algorithms and the Universal Method because of the following facts. First, degenerations cannot increase slice rank:

Proposition 5.1 ([TS16, Corollary 2]). If A and B are tensors such that $B \in A$, then $S(B) \leq S(A)$, and hence $S(B) \leq S(A)$.

Second, the independent tensor $\sum_{i,j,k} x_i y_j z_k$ has asymptotic slice rank q :

Proposition 5.2 ([Tao16, Lemma 1]; see also [BCC17a, Lemma 4.7]) For any positive integer q , we have $S(\sum_{i,j,k} x_i y_j z_k) = S(\sum_{i,j,k} x_i y_j z_k) = q$.

Third, MM tensors have (monomial) degenerations to large independent tensors:

Proposition 5.3 ([Str86, Theorem 4]) For any positive integers $a; b; c$ there is a $q = \frac{3}{4}abc = \max\{a; b; c\}$ such that $\sum_{i,j,k} x_i y_j z_k \in_{\text{md}} \sum_{i,j,k} x_i y_j z_k$.

Proof. Assume first that $a = 2m + 1$, $b = 2n + 1$, and $c = 2p + 1$ are all odd, and assume without loss of generality that $a \geq b \geq c$. We write

$$\sum_{i=1}^a x_i \sum_{j=1}^b y_j \sum_{k=1}^c z_k = \sum_{i=1}^m x_i \sum_{j=1}^n y_j \sum_{k=1}^p z_k + \dots$$

We define our monomial degeneration (using the notation of Section 4.4) via the map $m : X \times Y \times Z \rightarrow \mathbb{C}$ defined as follows:

$$\begin{aligned} m(x_{ij}) &= i^2 + 2ij + 3j^2, \\ m(y_{jk}) &= j^2 + 2jk + 3k^2, \text{ and} \\ m(z_{ki}) &= k^2 + 2ki + 3i^2. \end{aligned}$$

For any term $x_{ij} y_{jk} z_{ki}$ in $\mathbb{C}[x, y, z]$, we thus have $m(x_{ij}) + m(y_{jk}) + m(z_{ki}) = (i+j+k)^2 + 9p^2$. This exponent is always at least $9p^2$, and the term $x_{ij} y_{jk} z_{ki}$ is included in the resulting tensor D of the monomial degeneration if and only if $i + j + k = 0$. We can see that if $i + j + k = 0$, then any two of i, j, k determines the third, meaning any one of the variables x_{ij}, y_{jk}, z_{ki} determines the other two, and so D is indeed an independent tensor. Finally, there is a triple of (i, j, k) , $|i| \leq n; |j| \leq m; |k| \leq p$ with $i + j + k = 0$ for each pair (i, j) , $|i| \leq n; |j| \leq m$ with $|i| + |j| \leq p$. Since $p \geq n, m$, we can see there are at least $4ab$ such pairs, as desired. The cases where a, b, c are not all odd are similar. \square

Combined, these three facts show that MM tensors have large asymptotic slice rank. In fact, it is as large as possible given its number of variables:

Corollary 5.1. For any positive integers a, b, c we have $\mathcal{S}(a; b; c) = abc = \max\{a, b, c\}$.

Proof. Assume without loss of generality that $c \geq a, b$. For any positive integer n , we have that $(a; b; c)^n = (a^n; b^n; c^n) \in D_{h:75} a^n b^n$, meaning $\mathcal{S}(a; b; c)^n \geq 0.75 a^n b^n$ and hence $\mathcal{S}(a; b; c) \geq (0.75)^{1/n} ab$. Since this holds for all $n \in \mathbb{N}$, we see that $\mathcal{S}(a; b; c) \geq ab$. Meanwhile, $(a; b; c)$ has ab different x -variables, so it must have $\mathcal{S}_x(a; b; c) \geq ab$ and more generally, $\mathcal{S}(a; b; c)^n \leq \mathcal{S}_x(a; b; c)^n = (ab)^n$, which means $\mathcal{S}(a; b; c) \leq ab$. \square

To summarize: we know that degenerations cannot increase asymptotic slice rank, and that matrix multiplication tensors have a high asymptotic slice rank. Hence, if T is a tensor such that $\gamma_u(T)$ is 'small', meaning a power of T has a degeneration to a large matrix multiplication tensor, then T itself must have 'large' asymptotic slice rank. To be more precise:

Theorem 5.1. For any tensor T ,

$$\gamma_u(T) \geq 2 \frac{\log(\mathcal{R}(T))}{\log(\mathcal{S}(T))}.$$

Proof. By definition of $\gamma_u(T)$, for every $\epsilon > 0$, there are $n, q \in \mathbb{N}$ with $T^n \in D_{h:q} h; q; q$ and $q \leq \mathcal{R}(T)^{n(\gamma_u(T) + \epsilon)}$. By Proposition 5.1 and Corollary 5.1, it follows that $\mathcal{S}(T)^{q^2} \leq \mathcal{R}(T)^{2n(\gamma_u(T) + \epsilon)}$. Rearranging gives that $\gamma_u(T) + \epsilon \geq \frac{2 \log(\mathcal{R}(T))}{\log(\mathcal{S}(T))}$, and since this holds for all $\epsilon > 0$, the desired result follows. \square

Corollary 5.2. For any tensor T , if $\text{rank}_u(T) = 2$, then $\mathfrak{S}(T) = \mathfrak{R}(T)$. Moreover, for every constant $s < 1$, every tensor T with $\mathfrak{S}(T) \leq \mathfrak{R}(T)^s$ must have $\text{rank}_u(T) \geq 2/s > 2$.

By Theorem 5.1, in order to prove lower bounds on $\text{rank}_u(T)$, it suffices to prove upper bounds on $\mathfrak{S}(T)$! In Section 5.3, we will give a number of new tools for doing so.

5.3 Combinatorial Tools for Asymptotic Slice Rank Upper Bounds

We now give three general tools for proving upper bounds on $\mathfrak{S}(T)$ for many tensors T . Each of our tools applies to a large class of tensors, but we will see in particular that all three of them apply to the Coppersmith-Winograd tensor CW_q .

The general idea for the three tools is to find partitions $T = A + B$ of our tensors, such that at least one of $\mathfrak{S}(A)$ and $\mathfrak{S}(B)$ is low, and use this to show that $\mathfrak{r}(T)$ is itself low. If \mathfrak{S} were subadditive, i.e. if it were the case that $\mathfrak{S}(T) \leq \mathfrak{S}(A) + \mathfrak{S}(B)$ when $T = A + B$, then this would be relatively straightforward. Unfortunately, \mathfrak{S} is not subadditive in general, and even in many natural situations:

Example 5.1. Let q be any positive integer, and define the tensors $T_1 := \sum_{i=0}^q x_0 y_i z_i$, $T_2 := \sum_{i=1}^{q+1} x_i y_0 z_i$, and $T_3 := \sum_{i=1}^{q+1} x_i y_i z_{q+1}$. We can see that T_1 has only one x -variable, T_2 has only one y -variable, and T_3 has only one z -variable, and so $\mathfrak{S}(T_1) = \mathfrak{S}(T_2) = \mathfrak{S}(T_3) = 1$. However, $T_1 + T_2 + T_3 = CW_q$, so the three tensors give a partition of the Coppersmith-Winograd tensor! Hence, for instance, using the fact that $\text{rank}_u(CW_5) \geq 2.373$ we see that $\mathfrak{S}(CW_5) \leq \mathfrak{R}(CW_5)^{2/\text{rank}_u(CW_5)} \approx 7^{2/2.373} \approx 5.15$.

Throughout this section, we will nonetheless describe a number of general situations where, if T can be written as $T = A + B$, then bounds on $\mathfrak{S}(A)$ and $\mathfrak{S}(B)$ are sufficient to give bounds on $\mathfrak{r}(T)$.

5.3.1 Bounds from Variable-Deficient Partitions

We know that tensors T without many of one type of variable have small $\mathfrak{S}(T)$. For instance, if T is over $X; Y; Z$, and $|X|$ is 'small', then $\mathfrak{S}(T) \leq |X|$ is also small. We begin by showing that if T can be written as a sum of a few tensors, each of which does not have many of one type of variable, then we can still prove an upper bound on $\mathfrak{S}(T)$.

If $X; Y; Z$ are minimal for T , then the measure of T , denoted $\mu(T)$, is given by $\mu(T) := |X| |Y| |Z|$. We state two simple facts about μ :

Fact 5.1. For tensors A and B ,

$$\mu(A + B) \leq \mu(A) + \mu(B), \text{ and}$$

$$\text{if } A \text{ is minimal over } X; Y; Z, \text{ then } \mathfrak{S}(A) \leq \min\{|X|; |Y|; |Z|\} \leq \mu(A)^{1/3}.$$

Theorem 5.2. Suppose T is a tensor, and T_1, \dots, T_k are tensors with $T = T_1 + \dots + T_k$. Then, $\mathcal{S}(T) \leq \sum_{i=1}^k \mathcal{S}(T_i)^{1=3}$.

Proof. Note that

$$T^n = \sum_{(P_1, \dots, P_n) \in \mathcal{P}_{T_1, \dots, T_k}^n} \prod_{i=1}^n P_i$$

It follows that

$$\begin{aligned} \mathcal{S}(T^n) &= \sum_{(P_1, \dots, P_n) \in \mathcal{P}_{T_1, \dots, T_k}^n} \mathcal{S}\left(\prod_{i=1}^n P_i\right) \\ &= \sum_{(P_1, \dots, P_n) \in \mathcal{P}_{T_1, \dots, T_k}^n} \left(\mathcal{S}(P_1) \dots \mathcal{S}(P_n)\right)^{1=3} \\ &= \left(\mathcal{S}(T_1)^{1=3} + \dots + \mathcal{S}(T_k)^{1=3}\right)^n; \end{aligned}$$

which implies as desired that $\mathcal{S}(T) \leq \mathcal{S}(T_1)^{1=3} + \dots + \mathcal{S}(T_k)^{1=3}$. \square

5.3.2 Bounds from Block Partitions

This tool will be the most important in upper bounding the asymptotic slice rank of many tensors of interest. We show that a partitioning method similar to the Laser Method applied to a tensor T can be used to prove upper bounds on $\mathcal{S}(T)$. We begin by introducing some notation related to partitioning tensors into blocks, similar to the notation used in Subsection 4.6.1 when describing the Laser Method.

Partition Notation Throughout this subsection, we will be partitioning the terms of tensors into blocks defined by partitions of the three variable sets. Here we introduce some notation for some properties of such partitions; these definitions all depend on the particular partition of the variables being used, which will be clear from context.

Suppose T is a tensor minimal over $X; Y; Z$, and let $X = X_1 \sqcup \dots \sqcup X_{k_X}$, $Y = Y_1 \sqcup \dots \sqcup Y_{k_Y}$, $Z = Z_1 \sqcup \dots \sqcup Z_{k_Z}$ be partitions of the three variable sets. For $(i; j; k) \in [k_X] \times [k_Y] \times [k_Z]$, let T_{ijk} be T restricted to $X_i; Y_j; Z_k$ (i.e. T with $X \setminus X_i$, $Y \setminus Y_j$, and $Z \setminus Z_k$ zeroed out); T_{ijk} is called a block of T . Let $L = \{T_{ijk} \mid (i; j; k) \in [k_X] \times [k_Y] \times [k_Z]; T_{ijk} \neq 0\}$ be the set of non-zero blocks. For $i \in [k_X]$ let $L_{X_i} = \{T_{ijk} \mid (i; j; k) \in [k_X] \times [k_Y] \times [k_Z]\}$ be the set of blocks involving X_i , and define similarly L_{Y_j} and L_{Z_k} .

We will be particularly interested in probability distributions $p : L \rightarrow [0; 1]$. Let $\mathcal{P}(L)$ be the set of such distributions. For such $p \in \mathcal{P}(L)$, and for $i \in [k_X]$, let $p(X_i) := \sum_{T_{ijk} \in L_{X_i}} p(T_{ijk})$, and similarly $p(Y_j)$ and $p(Z_k)$. Then, define $p_X \in \mathbb{R}$ by

$$p_X := \sum_{i \in [k_X]} \frac{p(X_i)}{p(X_i)^{1=3}};$$

and p_Y and p_Z similarly. We can equivalently write $p_X = 2^{H(p(X))}$; where $H(p(X)) = \sum_{i \in [k_X]} p(X_i) \log p(X_i)$ is the entropy of the marginal distribution of p on the parts of X . This expression, which arises naturally in the Laser Method, will play an important role in our upper bounds and lower bounds on \mathfrak{S} .

The Main Tool We now present the main tool for bounding \mathfrak{S} in terms of the quantities we just defined.

Theorem 5.3. For any tensor T and partition of its variable sets,

$$\mathfrak{S}(T) = \limsup_{p \in \mathcal{P}(L)} \min_{p_X; p_Y; p_Z} g$$

Proof. For any positive integer n , we can write

$$T^n = \sum_{(P_1, \dots, P_n) \in \mathcal{L}^n} X_{P_1} \otimes \dots \otimes X_{P_n}$$

For a given $(P_1, \dots, P_n) \in \mathcal{L}^n$, let $\text{dist}(P_1, \dots, P_n)$ be the probability distribution on L which results from picking a uniformly random $j \in [n]$ and outputting P_j . For a probability distribution $p : L \rightarrow [0, 1]$, define $L_{n,p} := \{(P_1, \dots, P_n) \in \mathcal{L}^n \mid \text{dist}(P_1, \dots, P_n) = p\}$. Note that the number of p for which $L_{n,p}$ is nonempty is only $\text{poly}(n)$, since they are the distributions which assign an integer multiple of $1/n$ to each element of L . Let D be the set of these probability distributions.

We can now rearrange:

$$T^n = \sum_{p \in D} \sum_{(P_1, \dots, P_n) \in L_{n,p}} X_{P_1} \otimes \dots \otimes X_{P_n}$$

Hence,

$$\mathfrak{S}(T^n) \leq \sum_{p \in D} \sum_{(P_1, \dots, P_n) \in L_{n,p}} \mathfrak{S}(X_{P_1} \otimes \dots \otimes X_{P_n})$$

$$\leq \text{poly}(n) \max_{p \in D} \sum_{(P_1, \dots, P_n) \in L_{n,p}} \mathfrak{S}(X_{P_1} \otimes \dots \otimes X_{P_n})$$

For any probability distribution $p : L \rightarrow [0, 1]$, let us count the number of x -variables used in $\sum_{(P_1, \dots, P_n) \in L_{n,p}} X_{P_1} \otimes \dots \otimes X_{P_n}$. These are the tuples of the form $(x_1, \dots, x_n) \in X^n$ where, for each $i \in [k_X]$, there are exactly $n \cdot p(X_i)$ choices of j for

which $x_j \geq X_i$. The number of these is¹

$$\sum_{i \in [k_X]} \sum_{j \in [k_Y]} \sum_{k \in [k_Z]} p(X_i) p(Y_j) p(Z_k)$$

This is upper bounded by $p_X^{n+o(n)}$. It follows that $S_X \leq \sum_{(P_1, \dots, P_n) \in \mathcal{P}_{L,n,p}} p_X^{n+o(n)}$. We can similarly argue about S_Y and S_Z . Hence,

$$\begin{aligned} S(T^n) &\leq \text{poly}(n) \max_{p \in \mathcal{P}_{L,n,p}} \sum_{(P_1, \dots, P_n) \in \mathcal{P}_{L,n,p}} p_X^{n+o(n)} \\ &\leq \text{poly}(n) \max_{p \in \mathcal{P}_{L,n,p}} \min_{p_X, p_Y, p_Z} g^{n+o(n)} \\ &\leq \text{poly}(n) \limsup_{p \in \mathcal{P}(L)} \min_{p_X, p_Y, p_Z} g^{n+o(n)}. \end{aligned}$$

Hence, $S(T^n) \leq \limsup_p \min_{p_X, p_Y, p_Z} g^{n+o(n)}$, and the desired result follows. \square

We make one remark about the quantity p_X in Theorem 5.3 before continuing. Suppose T is over $X; Y; Z$ with $|X| = |Y| = |Z| = q$. For any probability distribution p we always have $p_X, p_Y, p_Z \leq q$, and moreover we only have $p_X = q$ when $p(X_i) = q$ for each i . It follows that if no probability distribution p is ϵ -close (say, in ℓ_1 distance) to having $p(X_i) = q$ for all i , $p(Y_j) = q$ for all j , and $p(Z_k) = q$ for all k , simultaneously, then we get $S(T) \leq q^{1-f(\epsilon)}$ for some increasing function f with $f(\epsilon) > 0$ for all $\epsilon > 0$. This gives a simple test for whether Theorem 5.3 can give a nontrivial upper bound on $S(T)$ for a tensor T .

Symmetric Block Partitions We make a remark about applying Theorem 5.3 to variable-symmetric tensors. This remark has implicitly been used in past work on applying the Laser Method, such as [CW90], but we prove it here for completeness. Recall the notation in Section 4.5 about variable-symmetric tensors.

If T is a variable-symmetric tensor minimal over $X; Y; Z$, then partitions $X = X_1 \sqcup \dots \sqcup X_{k_X}$, $Y = Y_1 \sqcup \dots \sqcup Y_{k_Y}$, $Z = Z_1 \sqcup \dots \sqcup Z_{k_Z}$ of the variable sets are called T -symmetric if (using the partition notation above) $k_X = k_Y = k_Z$, $|X_i| = |Y_i| = |Z_i|$ for all $i \in [k_X]$, and the block $T_{jki} = \text{rot}(T_{ijk})$ for all $(i; j; k) \in [k_X]^3$. For the L resulting from such a T -symmetric partition, a probability distribution $p \in \mathcal{P}(L)$ is called T -symmetric if it satisfies $p(T_{ijk}) = p(T_{jki})$ for all $(i; j; k) \in [k_X]^3$, and we write $\mathcal{P}^{\text{sym}}(L) \subseteq \mathcal{P}(L)$ for the set of such T -symmetric distributions. Notice in particular that any $p \in \mathcal{P}^{\text{sym}}(L)$ satisfies $p_X = p_Y = p_Z$.

Proposition 5.4. Suppose T is a variable-symmetric tensor over $X; Y; Z$, and $X = X_1 \sqcup \dots \sqcup X_{k_X}$, $Y = Y_1 \sqcup \dots \sqcup Y_{k_Y}$, $Z = Z_1 \sqcup \dots \sqcup Z_{k_Z}$ are T -symmetric partitions.

¹Recall from Proposition 2.5 that, for fixed p_i s, we have $\sum_{p_1, p_2, \dots, p_n} \prod_{i=1}^n p_i^{n+o(n)}$. Throughout this dissertation we use the convention that $p_i^{p_i} = 1$ when $p_i = 0$.

Then,

$$\mathfrak{S}(T) = \limsup_{p \in P(L)} p_X :$$

Proof. We know from Theorem 5.3 that $\mathfrak{S}(T) = \limsup_{p \in P(L)} \min\{p_X; p_Y; p_Z\} g$. We will show that for any $p \in P(L)$, there is $p^0 \in P^{sym}(L)$ such that $\min\{p_X; p_Y; p_Z\} g \leq \min\{p_X^0; p_Y^0; p_Z^0\} g$, which means that in fact, $\mathfrak{S}(T) = \limsup_{p \in P^{sym}(L)} \min\{p_X; p_Y; p_Z\} g$. Finally, the desired result will follow since, for any $p^0 \in P^{sym}(L)$, we have $p_X^0 = p_Y^0 = p_Z^0$.

Consider any $p \in P(L)$, and define the distribution $p^0 \in P^{sym}(L)$ by $p^0(T_{ijk}) := (p(T_{ijk}) + p(T_{jki}) + p(T_{kij})) / 3$ for each $T_{ijk} \in L$. In order to show that $\min\{p_X; p_Y; p_Z\} g \leq p_X^0$, we will show that $(p_X p_Y p_Z)^{1/3} \leq p_X^0$:

$$\begin{aligned} (p_X p_Y p_Z)^{1/3} &= \prod_{i \in [k_X]} \frac{\sum_j X_{ij} p(X_{ij})}{p(X_i)} \prod_{i \in [k_Y]} \frac{\sum_j Y_{ij} p(Y_{ij})}{p(Y_i)} \prod_{i \in [k_Z]} \frac{\sum_j Z_{ij} p(Z_{ij})}{p(Z_i)} \\ &= \frac{\prod_{i \in [k_X]} \sum_j X_{ij} p^0(X_{ij})}{(p(X_i)^{p(X_i)} p(Y_i)^{p(Y_i)} p(Z_i)^{p(Z_i)})^{1/3}} \\ &= \frac{\prod_{i \in [k_X]} \sum_j X_{ij} p^0(X_{ij})}{p^0(X_i)^{p^0(X_i)}} \\ &= p_X^0; \end{aligned}$$

where the second-to-last step follows from the fact that for any real numbers $a, b, c \in [0, 1]$, setting $d = (a + b + c) / 3$, we have $a^a b^b c^c \leq d^{3d}$. \square

Proposition 5.4 will help simplify calculations when we apply Theorem 5.3 to variable-symmetric tensors later in this Chapter.

5.3.3 Bounds from Parts with Low X-Rank

Finally we give our third tool. This tool will be the least important in proving bounds on \mathfrak{S} below, but in general it can help to extend upper bounds on \mathfrak{S} from tensors B for which upper bounds on $\mathfrak{S}(B)$ are known to tensors T which are slight modifications of B .

For a tensor T , let $m(T) := \max\{S_x(T); S_y(T); S_z(T)\}$. Recall from Lemma 5.1 that for any two tensors $A; B$ we have $\mathfrak{S}(A \otimes B) \leq \mathfrak{S}(A) + m(B)$.

In general, for two tensors A and B , even if $\mathfrak{S}(A)$ and $\mathfrak{S}(B)$ are 'small', it might still be the case that $\mathfrak{S}(A + B)$ is 'large', much larger than $\mathfrak{S}(A) + \mathfrak{S}(B)$. Here we show that if, not only is $\mathfrak{S}(A)$ small, but even $S_x(A)$ is small, then we can get a decent bound on $\mathfrak{S}(A + B)$.

Theorem 5.4. Suppose $T; A; B$ are tensors such that $A + B = T$. Then,

$$\mathfrak{S}(T) \leq \frac{m(A)}{(1 - p) S_x(A)} + \frac{1}{p^p};$$

where $p \in [0, 1]$ is given by

$$p := \frac{\log \frac{S_x(B)}{\mathfrak{S}(B)}}{\log \frac{m(A)}{S_x(A)} + \log \frac{S_x(B)}{\mathfrak{S}(B)}}:$$

Proof Sketch. We begin by, for any integers $n \geq k \geq 0$, giving bounds on $S(A^k B^{(n-k)})$. First, since S_x is submultiplicative, we have

$$S(A^k B^{(n-k)}) \leq S_x(A^k B^{(n-k)}) \leq S_x(A)^k S_x(B)^{n-k}:$$

Second, from the definition of m , we have

$$S(A^k B^{(n-k)}) \geq m(A^k) S(B^{(n-k)}) \geq m(A)^k \mathfrak{S}(B)^{n-k}:$$

It follows that for any positive integer n we have

$$S(T^{\otimes n}) \stackrel{X^n}{\underset{k=0}{\overset{n}{\max}}} \binom{n}{k} S(A^k B^{(n-k)}) \stackrel{X^n}{\underset{k=0}{\overset{n}{\min}}} \binom{n}{k} \min\{S_x(A)^k S_x(B)^{n-k}; m(A)^k \mathfrak{S}(B)^{n-k}\}:$$

We can see that the quantity $\binom{n}{k} \min\{S_x(A)^k S_x(B)^{n-k}; m(A)^k \mathfrak{S}(B)^{n-k}\}$ is maximized at $k = pn$, and the result follows. \square

5.4 Slice Rank Lower Bounds via the Laser Method

In the previous section, we gave three general tools for proving upper bounds on $\mathfrak{S}(T)$. Before applying them to particular tensors of interest, we begin in this section by giving a general tool for proving lower bounds on $\mathfrak{S}(T)$. Our main tool for proving lower bounds will be the Laser Method, the same technique we described in Subsection 4.6.1 which Strassen, Coppersmith, and Winograd developed for proving upper bounds on $\mathfrak{S}(T)$. The Laser Method only applies to tensors T with certain block structure, but we will show that when it applies to T , then not only does it give a lower bound on $\mathfrak{S}(T)$, but that the resulting bound matches the upper bound on $\mathfrak{S}(T)$ from one of our upper bounding tools, Theorem 5.3.

Consider any tensor T which is minimal over $X; Y; Z$, and let $X = X_1 \sqcup \dots \sqcup X_{k_X}$, $Y = Y_1 \sqcup \dots \sqcup Y_{k_Y}$, $Z = Z_1 \sqcup \dots \sqcup Z_{k_Z}$ be partitions of the three variable sets. Define T_{ijk} , L , and p_X for a probability distribution p on L , as in the top of Subsection 5.3.2. Recall in particular that T_{ijk} is T restricted to the variable sets X_i , Y_j , and Z_k .

Definition 5.1. We say that T , along with partitions of $X; Y; Z$, is a laser-ready tensor partition if the following three conditions are satisfied:

- (1) For every $(i; j; k) \in [k_X] \times [k_Y] \times [k_Z]$, either $T_{ijk} = 0$, or else T_{ijk} has a degener-

ation to a tensor $h; a; b; c$ with $ab = \sum_j X_{ij}$, $bc = \sum_j Y_{jk}$, and $ca = \sum_j Z_{kj}$ (i.e. a matrix multiplication tensor which is as big as possible given $\sum_j X_{ij}$, $\sum_j Y_{jk}$, and $\sum_j Z_{kj}$).

- (2) There is an integer ℓ such that $T_{ijk} \neq 0$ only if $i + j + k = \ell$.
- (3) T is variable-symmetric, and the partitions are \bar{T} -symmetric.

These conditions match those discussed in Subsection 4.6.1, and are exactly those required for the original Laser Method used by Coppersmith and Winograd [CW90] applies to T . We note that condition (3) is a simplifying assumption rather than a real condition on T : similar to the discussion in Section 4.5, for any tensor T and partitions satisfying conditions (1) and (2), the tensor $\text{sym}(T)$ along with the corresponding partitions on its variables (which arise from taking the products of the partitions of the variables of T), satisfies all three conditions, gives at least as good a bound on $n!$ using the Laser Method as T and the original partitions, and more generally has $n!(\text{sym}(T)) \geq n!(T)$.

The precise way in which the Laser Method applies to a laser-ready tensor partition is as follows.

Theorem 5.5 ([CW90, DS13, Wil12]) Suppose T , along with the partitions of $X; Y; Z$, is a laser-ready tensor partition. Then, for any distribution $p \in \mathcal{P}^{\text{sym}}(L)$, and any positive integer n , the tensor T^n has a degeneration into

$$\sum_{i \in [k_X]} \sum_{j \in [k_Y]} \sum_{k \in [k_Z]} p(X_i)^{p(X_i)} A^{h(a; a; a)}; \quad (5.1)$$

where

$$a = \sum_{T_{ijk} \in L} \sum_j X_{ij}^{p(T_{ijk})} A^{n-2} : \quad (5.2)$$

Proof. This proof is relatively detailed and technical; a reader willing to take the Theorem statement for granted may wish to skip reading it and instead read the overview in Subsection 4.6.1.

We begin by assuming that $p(T_{ijk})$ is an integer multiple of $1/n$ for all $i; j; k \in [k_X]$. If this is not the case, it can be achieved by slightly modifying p to a new probability distribution which changes $p(T_{ijk})$ by at most $1/n$ for each $i; j; k$. This in particular changes $p(X_i)$ by at most k_X/n for all $i \in [k_X]$, and so the changes in the quantities (5.1) and (5.2) are subsumed by the $\omega(n)$'s in the exponents.

We now proceed to describe the degeneration. We will zero out variables in three phases, leaving a tensor which easily degenerates to the desired tensor after the third phase. We will use the partition notation from Subsection 5.3.2 and Theorem 5.3.

Phase One. We say a variable $x = (x_1; \dots; x_n) \in X^n$ is p -satisfying if, for all $i \in [k_X]$, the number of $j \in [n]$ such that $x_j \in X_i$ is $n \cdot p(X_i)$, and similarly for a

$y \in Y^n$ or $z \in Z^n$. In phase one, we zero out all $x \in X^n$, $y \in Y^n$, and $z \in Z^n$ which are not p -satisfying.

Call an $I \subseteq \{X_1, \dots, X_{k_x}\}^n$ (or similar for y and z variables) a block of variables. Notice that for each block of variables I , either all its elements are p -satisfying, or none are; call I p -satisfying if all its elements are. The number of p -satisfying blocks of variables is

$$C := \prod_{i=1}^n p(X_{i_1}) \dots p(X_{i_{k_x}}) = \sum_{I \subseteq [k_x]} \prod_{i \in I} p(X_i)^{p(X_i)} A^{1 - p(X_i)};$$

which is the left-hand quantity in (5.1). For $I = \{X_{i_1}, \dots, X_{i_n}\} \subseteq \{X_1, \dots, X_{k_x}\}^n$, $J = \{Y_{j_1}, \dots, Y_{j_n}\} \subseteq \{Y_1, \dots, Y_{k_y}\}^n$, and $K = \{Z_{k_1}, \dots, Z_{k_n}\} \subseteq \{Z_1, \dots, Z_{k_z}\}^n$, if $I; J; K$ are p -satisfying, and are such that $T_{I; J; K} \in L$ for all $\ell \in [n]$, then we say $(I; J; K)$ is a surviving triple. In particular, if $(I; J; K)$ is a surviving triple, then $T_{IJK} := \prod_{\ell=1}^n T_{I; J; K}$ remains after phase one of zeroing outs. From condition (1) in the definition of a laser-ready tensor partition, we have that $T_{IJK} \in D$; thus, if it were the case that, for each p -satisfying I , there were exactly one surviving triple involving I , and similarly for each p -satisfying J and K , then the result of phase one would be a desired tensor!

This is unfortunately not yet the case, but in the remaining steps we will zero out more blocks of variables so that this will become the case. Let A be the number of surviving triples, and B be the number of surviving triples involving a fixed p -satisfying block I . Note by symmetry that B is independent of which I we pick, and whether I is a block of x , y , or z variables, and moreover that $B = A/C$. Although we could count A and B exactly using multinomial coefficients, it turns out we will only need the simple bound $B \leq 2^{O(n)}$.

Phase Two. Set $M = 4B + 1$. In this phase, we will use a result of Salem and Spencer [SS42]: There is a subset $H \subseteq [M]$ of size $|H| \geq M^{1 - o(1)}$ which does not contain any nontrivial three-term arithmetic progressions modulo M ; in other words, if $a; b; c \in H$ such that $a + b = 2c \pmod{M}$, then $a = b = c$. Let $M^0 := |H|$.

Recall from condition (2) in the definition of a laser-ready tensor partition that there is a $\delta \in \mathbb{N}$ such that $T_{ijk} \in L$ only if $i + j + k = \delta$. In particular, if $(I; J; K)$ is a surviving triple, then (using the notation above) $i + j + \ell + k = \delta$ for all $\ell \in [n]$.

Pick independently and uniformly random $w_0; w_1; \dots; w_n \in \mathbb{Z}_M$, and using them define three hash functions (which map variable blocks to integers modulo M) $h_X : \{X_1, \dots, X_{k_x}\}^n \rightarrow \mathbb{Z}_M$, $h_Y : \{Y_1, \dots, Y_{k_y}\}^n \rightarrow \mathbb{Z}_M$, $h_Z : \{Z_1, \dots, Z_{k_z}\}^n \rightarrow \mathbb{Z}_M$ by:

$$h_X(X_{i_1}, \dots, X_{i_n}) := 2 \sum_{i=1}^{X^n} w \cdot i \pmod{M};$$

$$h_Y(Y_{j_1}, \dots, Y_{j_n}) := 2w_0 + 2 \sum_{j=1}^{X^n} w \cdot j \pmod{M};$$

$$h_Z(Z_{k_1}, \dots, Z_{k_n}) := w_0 + \sum_{k=1}^{X^n} w \cdot k \pmod{M};$$

Notice that in any surviving triple $(I; J; K)$:

$$h_X(I) + h_Y(J) = 2h_Z(K) \pmod{M}, \text{ and}$$

not only is $h_X(I)$ a uniformly random value in $[M]$, but it remains a uniformly random value even conditioned on the value of $h_Y(J)$ (and similarly, and one of $h_X(I)$, $h_Y(J)$, and $h_Z(K)$ is uniformly random conditioned on any other one).

In phase two, we zero out all the x -variables in any block I such that $h_X(I) \not\equiv H$, all the y -variables in any block J such that $h_Y(J) \not\equiv H$, and all the z -variables in any block K such that $h_Z(K) \not\equiv H$. It follows from the definition of H that any surviving triple $(I; J; K)$ which also survives phase two satisfies $h_X(I) = h_Y(J) = h_Z(K)$. For each $\alpha \in H$, let S_α be the set of surviving triples $(I; J; K)$ with $h_X(I) = h_Y(J) = h_Z(K) = \alpha$.

Phase Three. Now, in phase three, for every pair of distinct surviving triples (which survived phases one and two) $(I; J; K)$ and $(I'; J'; K')$ which share the block I , we will zero out all variables in the block I (and then similarly for y and z variables). Hence, for every surviving triple $(I; J; K)$ which survives phase three, there are no other surviving triples which survive phase three which share any of J , or K . We will next show that there is a choice of the randomness above so that the number of surviving triples which survive phase three is $C = n^{o(1)}$. This (along with the discussion earlier in phase one) will complete the proof.

We begin by computing some simple expected values. First, for a fixed $\alpha \in H$, we compute the expected size of S_α . The probability that a given surviving triple (from phase one) $(I; J; K)$ also survived phase two and is in S_α is M^{-2} . Indeed, each of $h_X(I)$ and $h_Y(J)$ equals α independently with probability M^{-1} , and then given those two events, it follows that $h_Z(K) = \alpha$ as well. Hence, by linearity of expectation, the expected size of S_α is $A = M^2$.

Second, for fixed $\alpha \in H$, we compute an upper bound on the expected number of unordered pairs of distinct surviving triples (from phase one) $(I; J; K)$ and $(I'; J'; K')$ which share the block I , and which are both in S_α . There are A choices for a surviving triple $(I; J; K)$, then $A - 1$ choices for the surviving triple $(I'; J'; K')$. Both triples will be in S_α if and only if $h_X(I) = h_Y(J) = h_Y(J') = \alpha$. Since each of those three hash values is independent, this happens with probability M^{-3} . In all, an upper bound

on the expected value is $\frac{1}{2}AB=M^3$ (the $\frac{1}{2}$ is because we are counting the same pair twice). It follows that the expected number of surviving triples which are zeroed out in phase three is at most $\frac{1}{2}AB=M^3$ (since each surviving triple which is zeroed out is in at least one pair which share a block, and we need to count each of the three types of blocks).

From the calculations above, we know that for a fixed λ , the expected number of surviving triples in S which survive past phase three is at least $\frac{1}{4}A=M^2$. Hence, summing over all $\lambda \in \mathbb{H}$, the expected number of surviving triples which survive past phase three is at least $\frac{1}{4}M^0A=M^2$. Hence, there is a choice of the randomness which achieves at least this many, as desired! \square

Applying the Laser Method, we get the following key new result about the asymptotic slice rank of laser-ready tensor partitions.

Theorem 5.6. Suppose T , along with the partitions of $X; Y; Z$, is a laser-ready tensor partition. Then,

$$\mathfrak{S}(T) = \limsup_{p \in \mathbb{P}^{\text{sym}}(L)} p_X :$$

Proof. The upper bound, $\mathfrak{S}(T) \leq \limsup_{p \in \mathbb{P}^{\text{sym}}(L)} p_X$, is given by Proposition 5.4.

For the lower bound, we know from Theorem 5.5 that for all $p \in \mathbb{P}^{\text{sym}}(L)$, and all positive integers n , the tensor $T^{\otimes n}$ has a degeneration into

$$\bigotimes_{i \in [k_X]} p(X_i)^{\otimes n} A^{\otimes n} \quad \text{with } a_i = a_i;$$

where

$$a_i = \bigotimes_{T_{ijk} \in \mathcal{L}} j X_{ij} p(T_{ijk}) A^{\otimes n} :$$

By Proposition 5.3, this means $T^{\otimes n}$ has a degeneration to an independent tensor of size

$$\bigotimes_{i \in [k_X]} p(X_i)^{\otimes n} A^{\otimes n} \quad a^2 = p_X^{\otimes n} :$$

Applying Propositions 5.1 and 5.2 implies that $\mathfrak{S}(T) \leq p_X$ for all $p \in \mathbb{P}^{\text{sym}}(L)$, as desired. \square

Corollary 5.3. If T is a tensor with a laser-ready tensor partition, and applying the Laser Method to T with this partition yields an upper bound $u(T) \leq c$ for some $c > 2$, then $u(T) > 2$.

Proof. When the Laser Method for a given $p \in \mathbb{P}^{\text{sym}}(L)$ shows, as in Theorem 5.5,

that T^n has a degeneration into

$$\sum_{i=1}^n p(X_i) p(X_i) A_i \quad \text{with } a_i, a_i, a_i;$$

the resulting upper bound on $u(T)$ is that

$$\sum_{i=1}^n p(X_i) p(X_i) A_i \quad a^{u(T)} \mathbb{R}(T)^n:$$

In particular, since the left-hand side equals p_X when $u(T) = 2$, this yields $u(T) = 2$ if and only if, for every $\epsilon > 0$, there is a $p \in P^{\text{sym}}(L)$ such that $p_X \leq \mathbb{R}(T)^{1-\epsilon}$. In particular, if it does not yield $u(T) = 2$, then there is a $\delta > 0$ such that all $p \in P^{\text{sym}}(L)$ have $p_X \leq \mathbb{R}(T)^{1-\delta}$. It follows from Theorem 5.6 that $\mathbb{S}(T) \leq \mathbb{R}(T)^{1-\delta}$. Combined with Theorem 5.1, this means that $u(T) = 2 - (1 - \delta) > 2$. \square

5.4.1 Slice Rank Versus Asymptotic Subrank

For a tensor T , let $Q^0(T)$ denote the largest integer q such that there is a degeneration $T \succ_D h q$. The asymptotic subrank of T is defined as $Q(T) := \limsup_{n \rightarrow \infty} Q^0(T^n)^{1/n}$. Asymptotic Subrank is an important notion in Strassen's theory of the Asymptotic Spectrum of Tensors [Str86, Str91]. It can be thought of as 'dual' to asymptotic rank: while $\mathbb{R}(T)$ measures the 'cost' of T to convert from an independent tensor $Q(T)$ is a measure of the 'value' of T in converting back to an independent tensor.

Propositions 5.1 and 5.2 above imply that $Q(T) = \mathbb{S}(T)$ for all tensors T . Similarly, it is not hard to see that Theorem 5.1, our general bound on $u(T)$, holds with \mathbb{S} replaced by Q . One could thus conceivably hope to prove stronger lower bounds than those which we will prove in the next section by bounding Q instead of \mathbb{S} .

However, one interesting corollary of Theorem 5.6 above is that $u(T) = \mathbb{S}(T)$ for every laser-ready tensor T . In particular, every tensor T we study in the next section will be laser-ready, so such an improvement on our lower bounds on $u(T)$ using $Q(T)$ is impossible for these tensors.

More generally, there are currently no known tensors for which the best known upper bound on $Q(T)$ is smaller than the best known upper bound on $\mathbb{S}(T)$ (including the new bounds of [CVZ18, CVZ19]). Hence, novel tools for upper bounding Q would be required for such an approach to proving better lower bounds on u .

Corollary 5.4. Every tensor T with a laser-ready tensor partition has $\mathbb{S}(T) = Q(T)$.

Proof. All tensors satisfy $\mathbb{S}(T) \leq Q(T)$. In Theorem 5.6, the upper bound on $\mathbb{S}(T)$ showed that T^n has a degeneration to an independent tensor of size $\mathbb{S}(T)^n \epsilon^{o(n)}$, which implies that $Q(T) \leq \mathbb{S}(T)$. \square

5.5 Computing the Slice Ranks for Tensors of Interest

In this section, we give slice rank upper bounds for a number of tensors of interest. It follows from Theorem 5.6 above that all of the bounds we prove in this Section are tight.

5.5.1 Generalized Coppersmith-Winograd Tensors

We begin with the generalized CW tensors defined in Subsection 4.6.1 above, which for a positive integer q and a permutation $\pi \in S_q$ are given by

$$CW_{q;\pi} := x_0 y_0 z_{q+1} + x_0 y_{q+1} z_0 + x_{q+1} y_0 z_0 + \sum_{i=1}^{q-1} (x_i y_{\pi(i)} z_0 + x_i y_0 z_i + x_0 y_i z_i):$$

The usual Coppersmith-Winograd tensor CW_q results by picking π to be the identity permutation. We can see that Theorems 5.2, 5.3, and 5.4 can all apply to $CW_{q;\pi}$ to prove nontrivial upper bounds on $\mathfrak{S}(T)$.

That said, we will now use Theorem 5.3 to prove a tight bound on $\mathfrak{S}(CW_{q;\pi})$. Our bound will imply that $\mathfrak{S}_u(CW_{q;\pi}) \leq 2^{1.6805q}$ for all $q \geq 2$ and all $\pi \in S_q$. Because of Theorem 5.6, no better lower bound on $\mathfrak{S}_u(CW_{q;\pi})$ is possible by arguing about $\mathfrak{S}(CW_{q;\pi})$ or even $\mathfrak{Q}(CW_{q;\pi})$.

We begin by partitioning the variable sets of $CW_{q;\pi}$, using the notation of Theorem 5.3. Let $X_0 = \{x_0\}$, $X_1 = \{x_1, \dots, x_q\}$, and $X_2 = \{x_{q+1}\}$, so that $X_0 \sqcup X_1 \sqcup X_2$ is a partition of the x -variables of $CW_{q;\pi}$.² Similarly, let $Y_0 = \{y_0\}$, $Y_1 = \{y_1, \dots, y_q\}$, $Y_2 = \{y_{q+1}\}$, $Z_0 = \{z_0\}$, $Z_1 = \{z_1, \dots, z_q\}$, and $Z_2 = \{z_{q+1}\}$. We can see this is a $CW_{q;\pi}$ -symmetric partition with $L = \{T_{002}, T_{020}, T_{200}, T_{011}, T_{101}, T_{110}\}$.

Consider any probability distribution $p \in P^{\text{sym}}(L)$. By symmetry, we know that $p(T_{002}) = p(T_{020}) = p(T_{200}) = v$ and $p(T_{011}) = p(T_{101}) = p(T_{110}) = 1 - 3v$ for some value $v \in [0, 1/3]$. Applying Theorem 5.3, and in particular Proposition 5.4, combined with Theorem 5.6, yields:

$$\mathfrak{S}(CW_{q;\pi}) \leq \sup_{v \in [0, 1/3]} \frac{q^{2(1-3v)}}{v^v (2-3-2v)^{2-3-2v} (1-3+v)^{1-3+v}}:$$

The values for the first few q can be computed using optimization software as follows:

²The sets of partitions were 1-indexed before, but we 0-index here for notational consistency with past work.

q	$\mathfrak{S}(CW_{q;})$
1	2:7551
2	3:57165
3	4:34413
4	5:07744
5	5:77629
6	6:44493
7	7:08706
8	7:70581

Finally, using the lower bound $\mathfrak{R}(CW_{q;}) \geq q + 2$ (in fact, it is known that $\mathfrak{R}(CW_{q;}) = q + 2$), and the upper bound on $\mathfrak{S}(CW_{q;})$ we just proved, we can apply Theorem 5.1 to give lower bound $\mathfrak{!}_u(CW_{q;}) \geq 2 \log(\mathfrak{R}(CW_{q;})) = \log(\mathfrak{S}(CW_{q;})) \geq 2 \log(q + 2) = \log(\mathfrak{S}(CW_{q;}))$ as follows:

q	Lower Bound on $\mathfrak{!}_u(CW_{q;})$
1	2:16805
2	2:17794
3	2:19146
4	2:20550
5	2:21912
6	2:23200
7	2:24404
8	2:25525

It seems clear numerically that the resulting lower bound on $\mathfrak{!}_u(CW_{q;})$ is increasing with q and is always at least 2:16805; below we give a simple proof of this, concluding our main result about $CW_{q;}$.

Theorem 5.7. $\mathfrak{!}_u(CW_{q;}) \geq 2:16805$ for all $q \in \mathbb{N}$ and $q \geq 2$.

Proof. Define the function $f : [0; 1] \rightarrow \mathbb{R}$ by

$$f(v) := \frac{1}{v^{2-3} (2v)^{2-3} 2^v (1-3+v)^{1-3+v}}.$$

We already showed that

$$\mathfrak{!}_u(CW_{q;}) \geq \min_{v \in [0; 1]} 2 \frac{\log(q+2)}{\log(q^{2-3} 2^v f(v))}.$$

Moreover, we saw above that $\mathfrak{!}_u(CW_{q;}) \geq 2:16805$ for all $q \leq 8$.

Let v_q denote the argmin for the optimization problem. In particular, for $q = 8$, the argmin is $v_8 = 0:017732422$. From the $q^{2-3} 2^v$ term in the optimization problem, we see that $v_{q+1} < v_q$ for all q , and in particular, $v_q < v_8$ for all $q > 8$. It follows that $f(v_q) < f(v_8) = 2:07389$ for all $q > 8$. Thus, for all $q > 8$ we have:

$$\mathfrak{!}_u(CW_{q;}) \geq \min_{v \in [0; 1]} 2 \frac{\log(q+2)}{\log(q^{2-3} 2^v f(v_8))} = 2 \frac{\log(q+2)}{\log(q^{2-3} f(v_8))}.$$

This expression equals 2:18562 at $q = 9$, and is easily seen to be increasing with q for $q > 9$, which implies as desired that $\log_2(\mathbb{R}(CW_q)) \geq 2:16805$ for all $q \geq 9$ and hence all q . \square

5.5.2 Generalized Simple Coppersmith-Winograd Tensors

Similar to CW_q , we can define for a positive integer q and a permutation $\sigma : [q] \rightarrow [q]$ the simple Coppersmith-Winograd tensor $cw_{q;\sigma}$ given by:

$$cw_{q;\sigma} := \sum_{i=1}^q (x_i y_{\sigma(i)} z_0 + x_i y_0 z_i + x_0 y_i z_i):$$

These tensors, when σ is the identity permutation, are well-studied in the literature on MM algorithms. For instance, Coppersmith and Winograd [CW90] showed that if $\mathbb{R}(cw_{2;\text{id}}) = 2$ then $\log_2 \mathbb{R}(CW_q) \geq 2$.

We will again give a tight bound on $\mathbb{S}(cw_{q;\sigma})$ using Theorem 5.3 combined Theorem 5.6. To apply Theorem 5.3, and in particular Proposition 5.4, we again pick a partition of the variables. Let $X_0 = \{x_0\}$, $X_1 = \{x_1, \dots, x_q\}$, $Y_0 = \{y_0\}$, $Y_1 = \{y_1, \dots, y_q\}$, $Z_0 = \{z_0\}$, and $Z_1 = \{z_1, \dots, z_q\}$. This is a $cw_{q;\sigma}$ -symmetric partition with $L = \{T_{011}, T_{101}, T_{110}\}$. There is a unique $2^{\text{sym}}(L)$, which assigns probability $1/3$ to each part. It follows that

$$\mathbb{S}(cw_{q;\sigma}) = (1/3)^{1/3} (2/3)^{2/3} q^{2/3} = \frac{3}{2^{2/3}} q^{2/3}.$$

Again, we will see in the next section that this bound is tight. Using the lower bound $\mathbb{R}(cw_{q;\sigma}) \geq q+1$ from 'attening', we get the lower bound

$$\log_2(\mathbb{R}(cw_{q;\sigma})) \geq 2 \frac{\log(q+1)}{\log \frac{3}{2^{2/3}} q^{2/3}}.$$

The first few values are as follows; note that we cannot get a bound better than 2 when $q = 2$ because of Coppersmith and Winograd's remark: $\mathbb{R}(cw_2) = 2$ then $\log_2(\mathbb{R}(cw_2)) = 2$, but the best known bound is only $\mathbb{R}(cw_2) \geq 3$.

q	Lower Bound on $\log_2(\mathbb{R}(cw_{q;\sigma}))$
1	2:17795
2	2
3	2:02538
4	2:06244
5	2:09627
6	2:12549
7	2:15064

5.5.3 Cyclic Group Tensors

We next look at the group tensor T_q of the cyclic group C_q for $q \geq 2$ \mathbb{N} :

$$T_q = \sum_{i=0}^{q-1} \sum_{j=0}^{q-1} x_i y_j z_{i+j \pmod q}$$

T_q is one of the first tensors which was studied in the Group-theoretic Method [CU03]. Define also the 'lower triangular' version of T_q , called T_q^{lower} , as:

$$T_q^{\text{lower}} = \sum_{i=0}^{q-1} \sum_{j=0}^{q-1-i} x_i y_j z_{i+j}$$

It is known that $\mathcal{R}(T_q) = \mathcal{R}(T_q^{\text{lower}}) = q$, and a Coppersmith-Winograd-like analysis is possible to yield the best known bound $\mathfrak{u}(T_q) \geq \mathfrak{u}(T_q^{\text{lower}}) \geq 2.373$ (see Section 5.6 below for a proof).

While Theorem 5.3 does not give any nontrivial upper bounds on $\mathfrak{S}(T_q)$, it does give nontrivial upper bounds on $\mathfrak{S}(T_q^{\text{lower}})$, by using the 'trivial' partition of the variables into parts of size 1. Using computer optimization software, we can compute $\mathfrak{S}(T_q^{\text{lower}})$, using Theorem 5.3 where each partition contains exactly one variable, combined with Theorem 5.6, for the first few values of q :

q	$\mathfrak{S}(T_q^{\text{lower}})$
2	1.88988
3	2.75510
4	3.61071
5	4.46157

We thus get the following lower bounds on $\mathfrak{u}(T_q^{\text{lower}}) \geq 2 \log(q) = \log(\mathfrak{S}(T_q^{\text{lower}}))$:

q	Lower Bound on $\mathfrak{u}(T_q^{\text{lower}})$
2	2.17795
3	2.16805
4	2.15949
5	2.15237

These numbers match the lower bounds obtained by [AW18a, BCC7a] in their study of T_q ; our Theorem 5.3 can be viewed as an alternate tool to achieve those lower bounds. The bound approaches 2 as $q \rightarrow \infty$, as it is known that $\log(\mathfrak{S}(T_q)) = \log(q) = 1 + o(1)$ as $q \rightarrow \infty$ [BCC+17a].

There is a simple monomial degeneration $T_q \xrightarrow{D_{\text{md}}} T_q^{\text{lower}}$, and it is shown in [CVZ18, Theorem 4.16] that there is a restriction $T_q^{\text{lower}} \xrightarrow{D_{\text{md}}} T_q$ over the field F_q , which implies that our bounds above also hold for T_q over F_q .

5.5.4 Lower Triangular Tensors

More generally, we can give a strong characterization of lower triangular tensors for which Theorem 5.1 can prove $\mathfrak{S}(T) > 2$.

Definition 5.2. For $X = \langle x_0, \dots, x_{q-1} \rangle$, $Y = \langle y_0, \dots, y_{q-1} \rangle$ and $Z = \langle z_0, \dots, z_{q-1} \rangle$, a tensor T over $X; Y; Z$ is lower triangular if

For every $i, j \in \{0, \dots, q-1\}$, there is at most one $k \in \{0, \dots, q-1\}$ such that $x_i y_j z_k \in T$, and

For every $i, j \in \{0, \dots, q-1\}$ with $i + j < q-1$, $x_i y_j z_k \notin T$ for any $k \in \{1, \dots, q\}$.

Terms $x_i y_j z_k$ with $i + j = q-1$ are called diagonal terms

Theorem 5.8. For $X = \langle x_0, \dots, x_{q-1} \rangle$, $Y = \langle y_0, \dots, y_{q-1} \rangle$ and $Z = \langle z_0, \dots, z_{q-1} \rangle$, a lower triangular tensor T over $X; Y; Z$ has $\mathfrak{S}(T) = q$ if and only if it has q diagonal terms, no two of which share x -variables.

Proof. First, consider any lower diagonal tensor whose diagonal terms do not share z -variables. Evidently $\mathfrak{S}(T) \geq |X| = q$. There is a simple monomial degeneration from T to the tensor consisting only of its diagonal terms, given by $m(x_i) = m(y_i) = z_i$ and $m(z_i) = 1$ for all i . Since no two of the diagonal terms share z -variables, this is a monomial degeneration from T to an independent tensor of size q , which implies that $\mathfrak{S}(T) = q$.

Second, consider any lower diagonal tensor with $\mathfrak{S}(T) = q$. Let $f : \{0, \dots, q-1\} \rightarrow \{0, \dots, q-1\}$ be the map defining which z -variable appears in each term, i.e. such that $x_i y_j z_{f(i,j)}$ is the only term containing $x_i y_j$ for each i, j (we assume that such a term exists for each i, j ; if T is missing any such terms, then the proof is even simpler). By Theorem 5.3 (with each part in the partitions of the variables having size 1), we know that for every $\epsilon > 0$, there is a probability distribution $p : X \times Y \times Z \rightarrow [0, 1]$ whose support is on the terms of T , such that for any fixed i , $p(x_i) := \sum_{x_i y_j z_k} p(x_i y_j z_k) = 1 - \epsilon$, and similarly for $p(y_j)$ and $p(z_k)$. Summing this lower bound for all x -variables other than x_i also shows that $p(x_i) \geq 1 - \epsilon(q-1)$ for each i , and similarly for $p(y_j)$ and $p(z_k)$.

We now prove that for each $j \in \{0, \dots, q-1\}$, we have $p(x_{q-1-j} y_j z_{f(q-1-j, j)}) \geq 1 - \epsilon(q-1)$, where we are thinking of q as a constant, so the \mathcal{O} hides factors of q . We prove this by strong induction on j . For the base case, when $j = 0$, notice that the term $x_{q-1} y_0 z_{f(q-1, 0)}$ is the only term containing x_{q-1} , and $\text{supp}(x_{q-1} y_0 z_{f(q-1, 0)}) = p(x_{q-1}) \geq 1 - \epsilon$, as desired.

For the inductive step, note that for each $j^0 < j$, we have by assumption that $p(x_{q-1-j^0} y_{j^0} z_{f(q-1-j^0, j^0)}) \geq 1 - \epsilon(q-1)$. Therefore, for each such j^0 ,

³We write $x_i y_j z_k \in T$ to denote that the coefficient of $x_i y_j z_k$ in T is nonzero.

$$\begin{aligned}
& p(x_{q-1} y_j z_{(i,j)}) \\
& \sum_{i=0}^{q-1} p(x_i y_j z_{(i,j)}) \\
& = p(y_j) \sum_{i=0}^{q-1} p(x_{q-1-i} y_j z_{(q-1-i,j)}) = (1-q^{-1}) \sum_{i=0}^{q-1} p(x_{q-1-i} y_j z_{(q-1-i,j)}) = O(\epsilon):
\end{aligned}$$

It follows as desired that

$$\begin{aligned}
p(x_{q-1} y_j z_{(q-1,j)}) &= p(x_{q-1} y_j) \sum_{i=0}^{q-1} p(x_{q-1-i} y_j z_{(i,j)}) \\
&= p(x_{q-1} y_j) O(\epsilon) \\
&= O(\epsilon):
\end{aligned}$$

Now, assume to the contrary that there is k such that $k \notin f(q-1, j; j)$ for any j . Thus,

$$p(z_k) = \sum_{j=0}^{q-1} p(x_{q-1} y_j z_{(q-1,j)}) = \sum_{j=0}^{q-1} (1-q^{-1}) p(x_{q-1} y_j z_{(q-1,j)}) = O(\epsilon):$$

Picking a sufficiently small $\epsilon > 0$ contradicts Theorem 5.3. □

5.6 Upper Lower Bounds for Group Tensors

In the previous section, we mainly focused on applying our slice rank tools to various tensors which have been used in conjunction with the Laser Method to prove bounds on rank_s . In this section, we conclude the chapter by showing some implications for both upper and lower bounds on $\text{rank}_u(T_G)$ when T_G is a group tensor of a finite group G .

5.6.1 Tri-Colored Sum-Free Sets

A number of recent works (eg. [BCC⁺17a, BCC⁺17b, AW18a]) have explored connections between lower bounds on matrix multiplication algorithms, and a notion from extremal combinatorics called a 'tri-colored sum-free set'. We will make use of them here in our study of $\text{rank}_u(T_G)$ as well.

Definition 5.3. For a group G , a tri-colored sum-free set in G is a set $S \subseteq G^3$ of triples of elements of G such that:

for all $(a; b; c) \in S$, we have $ab = c$, and

for all $(a_1; b_1; c_1); (a_2; b_2; c_2); (a_3; b_3; c_3) \in S$ which are not all the same triple, we have $a_1 b_2 \neq c_3$.

In the literature, tri-colored sum-free sets are sometimes also called multiplicative matchings

In a recent breakthrough, Ellenberg and Gijswijt [EG17] used techniques introduced by Croot, Lev, and Pach [CLP17] to show that there is a constant $c < 3$ such that tri-colored sum-free sets in F_3^n have size at most $O(c^n)$. Since then, there has been an explosion of work in the area, and this result has been extended by Sawin [Saw18] to hold for all nontrivial groups G , even nonabelian groups:

Theorem 5.9 ([Saw18, Theorem 1]) Let G be any nontrivial finite group. There is a constant $c_G < 1$ such that for any positive integer n , any tri-colored sum-free set in G^n has size at most $(c_G |G|)^n$.

There are a number of families of groups \mathcal{S} where even stronger upper bounds than this are known; we refer the reader to the introduction of [BCC17b] for an exposition of these bounds.

We begin with the main connection between group tensors and tri-colored sum-free sets; this was essentially remarked in [BCC17a], but we reprove it here for completeness:

Lemma 5.2. For any finite group G and $q \geq 2$, if $T_G \leq_{\text{rank}} \text{hqj}$, then G has a tri-colored sum-free set of size $\Omega(q)$.

Proof. Let D be the subset of the terms of T_G so that $T_G \leq_{\text{rank}} D = \text{hqj}$. Let $S := \{(a; b; c) \in G^3 \mid x_a y_b z_c \in D\}$. We will show that S is a tri-colored sum-free set in G . First, recall that every $(a; b; c) \in T_G$ has $ab = c$, and $D \subseteq T_G$, and so every $(a; b; c) \in S$ has $ab = c$ as well. Second, assume to the contrary that there are $(a_1; b_1; c_1); (a_2; b_2; c_2); (a_3; b_3; c_3) \in S$, not all the same triple, such that $a_1 b_2 = c_3$. This means that none of $x_{a_1}; y_{b_2}$, or z_{c_3} were zeroed out to get from T_G to D . But, $x_{a_1} y_{b_2} z_{c_3} \in T_G$, and so we must have $x_{a_1} y_{b_2} z_{c_3} \in D$. Since D is independent, this means that $(a_1; b_1; c_1); (a_2; b_2; c_2);$ and $(a_3; b_3; c_3)$ must all be the same triple, contradicting how we picked them. \square

In fact, a more technical proof can strengthen Lemma 5.2 to work for degenerations instead of just zeroing outs:

Lemma 5.3 ([BCC⁺17a]). For any finite group G and $q \geq 2$, if $T_G \leq_{\text{rank}} \text{hqj}$, then for $n \geq 2$, G^n has a tri-colored sum-free set of size $\Omega(q)^{n - o(n)}$.

We can use this to give lower bounds $\text{oh}_u(T_G)$ for any finite group G :

Corollary 5.5. For any tensor T and any nontrivial finite group G such that $T \leq_{\text{rank}} T_G$, we have $\mathfrak{S}(T) < |G|$.

Proof. Since $T \leq_{\text{rank}} T_G$, we have $\mathfrak{S}(T) \leq \mathfrak{S}(T_G)$. Letting $c_G < 1$ be the constant from Theorem 5.9 for G , we know that for any positive integer n , any tri-colored sum-free set in G^n has size at most $(c_G |G|)^n$. Hence, by Lemma 5.3, we have $\mathfrak{S}(T_G^n) \leq (c_G |G|)^{n - o(n)}$. It follows that $\mathfrak{S}(T_G) \leq |G|$, as desired. \square

Theorem 5.10. For any finite group G , we have $\text{rk}(T_G) > 2$.

Proof. Simply pick $T = T_G$ in Corollary 5.5, which is known to have $\text{rk}(T_G) \geq |G|$ (see Subsection 4.6.2), combined with Theorem 5.1. \square

This shows that no fixed group tensor T_G can be used to show $\text{rk} = 2$ using the Universal Method. That said, it does not rule out showing $\text{rk} = 2$ by using a sequence $G_1; G_2; \dots$ of groups such that $\lim_{i \rightarrow \infty} \text{rk}(T_{G_i}) = 2$. Prior work has already made a similar remark for showing $\text{rk} = 2$ by finding large 'simultaneous triple product property' constructions in G via the Group-theoretic Method, and some natural sequences of groups have already been ruled out [BCC17b].

5.6.2 Asymptotic Slice Rank and Tri-Colored Sum-Free Set Constructions for All Finite Groups

One of the key components of our lower bounding framework is Proposition 5.3, in which we showed that matrix multiplication tensors have degenerations to large independent tensors. In this subsection, we will instead use Proposition 5.3 in a different way: to show that some other tensors of interest also have degenerations to nontrivially-large independent tensors. In particular, we will show this for the group tensor T_G of any finite group G , which will imply a nontrivially-large tri-colored sum-free set in G^n for sufficiently large n . We start with the main additional idea needed for this application:

Theorem 5.11. For every finite group G of order $|G| = q$, there is a monomial degeneration T_G into a tensor T which is a generalized Coppersmith-Winograd tensor with parameter $q - 2$.

Proof. Let $1 \in G$ be the identity, and let $g \in G$ be any other element. We will give three maps $\text{maps} : X_G \rightarrow \mathbb{Z}$, $\text{maps} : Y_G \rightarrow \mathbb{Z}$, and $\text{maps} : Z_G \rightarrow \mathbb{Z}$ such that for any $a; b \in G$ we have $\text{maps}(x_a) + \text{maps}(y_b) + \text{maps}(z_{ab}) = 0$, and then define our monomial degeneration by the map $m : X_G \times Y_G \times Z_G \rightarrow \mathbb{C}$ given by $m(x_a) = (x_a)^{+d}$, $m(y_b) = (y_b)^{+d}$, and $m(z_c) = (z_c)^{+d}$ for any $a; b; c \in G$ and a sufficiently large constant $d \in \mathbb{N}$, so that $x_a y_b z_{ab}$ will remain in the result of the monomial degeneration if and only if $\text{maps}(x_a) + \text{maps}(y_b) + \text{maps}(z_{ab}) = 0$. The maps are given as follows:

$$\text{maps}(x_1) = \text{maps}(y_1) = \text{maps}(z_1) = 0,$$

$$\text{maps}(x_g) = \text{maps}(y_g) = \text{maps}(z_g) = 2, \text{ and}$$

$$\text{maps}(x_h) = \text{maps}(y_h) = \text{maps}(z_h) = 1 \text{ for all } h \in G \setminus \{1; g\}.$$

Let T be the monomial degeneration of T_G defined by $\text{maps}; \text{maps}; \text{maps}$. Define the permutation $\sigma : G \setminus \{1; g\} \rightarrow G \setminus \{1; g\}$ which sends $h \in G$ to $\sigma(h) := h^{-1}g$. We can see that:

$$x_1 y_1 z_1 \in T \text{ since } \text{maps}(x_1) = \text{maps}(y_1) = \text{maps}(z_1) = 0.$$

$x_1 y_h z_h \geq 2$ for all $h \in G \setminus \{g\}$ (including $h = g$), since $(x_1) = 0$ while $(y_h) = (z_h) = 1$.

$x_h y_1 z_h \geq 2$ for all $h \in G \setminus \{g\}$ similarly.

$x_h y_{(h)} z_g \geq 2$ for all $h \in G \setminus \{g\}$, since $(x_h) = (y_{(h)}) = 1$, while $(z_g) = 2$.

Meanwhile,

$x_{h_1} y_{h_2} z_{h_3} \geq 2$ for any $h_1, h_2, h_3 \in G \setminus \{g\}$ with $h_1 h_2 = h_3$, since $(h_1) = (h_2) = 1$ and $(h_3) = 1$, so the three sum to 1.

$x_h y_{h^{-1}} z_1 \geq 2$ for any $h \in G \setminus \{g\}$ since $(x_h) = (y_{h^{-1}}) = 1$ while $(z_1) = 0$, so the three sum to 2.

$x_g y_{h_1} z_{h_2} \geq 2$ for any $h_1, h_2 \in G \setminus \{g\}$ with $gh_1 = h_2$, since $(x_g) = 2$, $(y_{h_1}) = 1$, and $(z_{h_2}) = 1$, so the three sum to 2.

$x_{h_1} y_g z_{h_2} \geq 2$ for any $h_1, h_2 \in G \setminus \{g\}$ with $gh_1 = h_2$ similarly.

$x_g y_{g^{-1}} z_1 \geq 2$ since $(x_g) = 2$, $(y_{g^{-1}}) = 1$, and $(z_1) = 0$, so the three sum to 3.

$x_{g^{-1}} y_g z_1 \geq 2$ similarly.

$x_g y_g z_{g^2} \geq 2$ since $(x_g) = (y_g) = 2$, and definitely $(z_{g^2}) = 2$, so the three sum to at least 2.

This covers all the entries of T_G , showing that we have defined a valid monomial degeneration to

$$T = x_1 y_1 z_1 + x_1 y_g z_g + x_g y_1 z_g + \sum_{h \in G \setminus \{g\}} (x_1 y_h z_h + x_h y_1 z_h + x_h y_{(h)} z_g):$$

This is indeed a generalized Coppersmith-Winograd tensor with parameter $q = 2$, as desired. \square

An immediate consequence of this monomial degeneration is that applying the Solar, Galactic or Universal method on T_G for any finite group G with $R(T_G) = |G|$ yields the same upper bounds on ω as the best known analysis of $CW_{|G|}$. Picking an appropriate group G where group operations are known to be efficient in practice could help lead to a more practical matrix multiplication algorithm.

Next, we will use the fact that matrix multiplication tensors, and hence Coppersmith-Winograd tensors, have large asymptotic slice rank, to show that for any finite group G , T_G also has a relatively large slice rank, and hence that \mathbb{C}^n has relatively large tri-colored sum-free sets for large enough n .

Theorem 5.12. Define $f : \mathbb{N} \rightarrow \mathbb{R}$ by $f(q) = \log_q \frac{4(q+2)^3}{27}$. For every positive integer q , and every tensor CW_q which is a generalized Coppersmith-Winograd tensor of parameter q , we have $\mathcal{S}(CW_q) \geq (q+2)^{2=f(q)}$.

Remark 5.1. For $q \geq 3$, we have $f(q) < 3$, and so $\mathfrak{S}(CW_{q;}) \leq (q+2)^{2-3}$.

Remark 5.2. In the proof of Theorem 5.12, we use a simpler lower bound $\text{bg}(CW_q)$ than is known, for ease of reading; it is, of course, possible to use the better known upper bounds $\text{ord}_s(CW_q)$ from [CW90, Wil12, LG14] in the proof and improve the result.

Proof of Theorem 5.12. Define $f : \mathbb{N} \rightarrow \mathbb{R}$ by $f(q) = \log_q \frac{4(q+2)^3}{27}$. In [CW90, Section 6], Coppersmith and Winograd show that $\text{ord}_s(CW_{q;}) \geq f(q)$. Hence, by Theorem 5.1, we get

$$\mathfrak{S}(CW_{q;}) \leq \mathfrak{R}(CW_{q;})^{2-f(q)} \leq (q+2)^{2-f(q)}.$$

□

Theorem 5.13. For every (not necessarily abelian) finite group G , there is a constant $c_{Gj} > 2-3$, depending only on $|G|$, such that $\mathfrak{S}(T_G) \leq |G|^{c_{Gj}}$. In particular, for $n \geq N$, G^n has a tri-colored sum-free set of size at least $|G|^{c_{Gj}n - o(n)}$.

Proof. The only finite groups G of order $|G| < 5$ are C_1, C_2, C_3, C_4 , and C_2^2 . For each of these groups, the result is shown, eg. by [KSS18]. For $|G| \geq 5$, we know from Theorem 5.11 that T_G has a monomial degeneration to a generalized Coppersmith-Winograd tensor of parameter $|G| \geq 2$, and so the result follows by Theorem 5.12. □

Part II

Probabilistic Polynomials and Hamming Nearest Neighbors

Chapter 6

Background and Overview

The polynomial method is a powerful tool in circuit complexity. The idea of the method is to transform all circuits of some class into simple polynomials which represent the circuit in some way. If the polynomial is always sufficiently simple (e.g. has low degree), and one can prove that a certain Boolean function cannot be represented so simply, one concludes that the circuit class is unable to compute

Recently, these tools have found surprising uses in algorithm design. If a subproblem of an algorithmic problem can be modeled by a simple circuit, and that circuit can be transformed into a simple polynomial (or simple distribution on polynomials), then fast algebraic algorithms can be applied to evaluate or manipulate the polynomial quickly. This approach has led to advances on problems such as All-Pairs Shortest Paths [Wil14a], Orthogonal Vectors [WY14, AWY15] and Constraint Satisfaction [Wil14d].

In these applications, the key step is to randomly convert simple circuits into probabilistic polynomials. If f is a Boolean function on n variables, and R is a commutative ring, a probabilistic polynomial over R for f with error $1-s$ and degree d is a distribution D of degree d polynomials over R such that for all $x \in \{0, 1\}^n$, $\Pr_{p \in D} [p(x) = f(x)] \geq 1 - s$. Razborov [Raz87] and Smolensky [Smo87] introduced the notion of a probabilistic polynomial. They showed that AND, OR, and XOR gates of unbounded fan-in have simple constant degree probabilistic polynomials, and hence that any low-depth AC^0 circuit consisting of these gates can be transformed into a low degree probabilistic polynomial. All the prior work on polynomial method algorithms uses this transformation.

In this Part, we develop new probabilistic polynomial constructions in order to solve a variety of algorithmic problems. We focus especially on polynomial representations of threshold functions. The threshold function TH_s determines whether at least a fraction s of its input bits are 1s. Threshold functions are among the simplest Boolean functions that do not have constant degree probabilistic polynomials: Razborov and Smolensky showed that the MAJORITY function (a special case of a threshold function) requires degree $\Omega(n \log s)$. Nonetheless, as we will see throughout this Part, there are many important problems which can be reduced to evaluating circuits involving threshold gates on many inputs, and so further study of polynomial representations of threshold functions is warranted. In-

deed, threshold functions have been extensively studied in theoretical computer science for many years, and there are numerous applications of linear and polynomial threshold functions to complexity and learning theory (a sample includes [BRS91, BS92, ABFR94, Bei95, KS01, OS10, She14]).

6.1 Our Results

6.1.1 Polynomial Constructions

We begin in Chapter 7 by giving a number of new polynomial constructions. We consider three different notions of polynomials representing TH . Each achieves different trade-offs between polynomial degree, the randomness required, and how accurately the polynomial represents TH . One key type of circuit which will recur in most of our applications is an OR of many thresholds; each of the polynomials we construct can be used to represent such a circuit by summing up t/s copies of the polynomial, one for each threshold gate (where s is the error parameter of the construction). Each construction leads to improved algorithms in our applications.

Probabilistic Polynomials. We begin with probabilistic polynomials. Razborov and Smolensky showed over 30 years ago that TH on n inputs with error $1/s$ requires degree $\Omega(\sqrt{n \log s})$. We show that their lower bound is tight by giving a matching construction. Our probabilistic polynomial construction is efficiently samplable using only $\text{polylog}(ns)$ random bits, which will allow us to use it to design deterministic algorithms in some cases.

Theorem 6.1. For any $0 < \epsilon < 1$, there is a probabilistic polynomial for the function TH of degree $O(\sqrt{n \log s})$ on n bits with error $1/s$ over any commutative ring R that can be efficiently sampled using only $O(\log n \log(ns))$ random bits.

Polynomial Threshold Functions. Second, we consider deterministic Polynomial Threshold Functions (PTFs). A PTF for a Boolean function f is a polynomial (not a distribution on polynomials) $p : \{0, 1\}^n \rightarrow R$ such that $p(x)$ is smaller than a fixed value when $f(x) = 0$, and $p(x)$ is larger than the value when $f(x) = 1$. In our applications, we seek PTFs with good threshold behavior, such that $|p(x) - 1| < \epsilon$ when $f(x) = 0$, and $p(x)$ is very large otherwise. We can achieve almost the same degree as for a probabilistic polynomial, and even better degree when we focus on ϵ -approximate thresholds rather than exact thresholds:

Theorem 6.2. We can construct a polynomial $P_{s;t,\epsilon} : R \rightarrow R$ of degree $O(\sqrt{(1/\epsilon) \log s})$, such that

$$\text{if } x \in \{0, 1\}^n \text{ and } f(x) = 0, \text{ then } |P_{s;t,\epsilon}(x) - 1| < \epsilon;$$

$$\text{if } x \in \{0, 1\}^n \text{ and } f(x) = 1, \text{ then } P_{s;t,\epsilon}(x) > 1 + \epsilon;$$

$$\text{if } x \in \{0, 1\}^n \text{ and } f(x) = 1, \text{ then } P_{s;t,\epsilon}(x) \leq s.$$

For the exact setting with $\epsilon = 1/t$, we can alternatively bound the degree by $O(\sqrt{t \log(st)})$.

By summing multiple copies of the polynomial from Theorem 6.2, we immediately obtain a PTF with the same degree for the OR of ϵ threshold functions (needed in our applications). This theorem follows directly from known extremal properties of Chebyshev polynomials, as well as the lesser known discrete Chebyshev polynomials. Chebyshev polynomials are well-known to yield good approximate polynomials for computing certain Boolean functions over the reals [NS94, Pat92, KS01, She13, Val15] (see Section 6.2 below for more background).




Figure 6-1: A plot of the sixth Chebyshev polynomial, $T_6(x) = 32x^6 - 48x^4 + 18x^2 - 1$. Chebyshev polynomials have the property that, on the interval $[-1, 1]$, they always output a value from $[-1, 1]$. Among all polynomials of a fixed degree with this property, the Chebyshev polynomial takes on the largest possible value at all inputs outside $[-1, 1]$. This makes it useful for applications like Theorem 6.2, where we want a large separation between inputs on either side of a threshold.

Probabilistic PTFs. Third, we introduce a new (natural) notion of a probabilistic PTF for a Boolean function f . This is a distribution on PTFs, where for each input x , a PTF drawn from the distribution is highly likely to agree with f on x . Combining the techniques from probabilistic polynomials for TH and the deterministic PTFs in a simple way, we construct a probabilistic PTF with good threshold behavior whose degree is lower than both the deterministic PTF and the degree bounds attainable by probabilistic polynomials (surprisingly breaking the square-root barrier of the Razborov-Smolensky lower bound):

Theorem 6.3. We can construct a distribution $\mathcal{L}_{n,s;t,\epsilon}$ on polynomials $L_{n,s;t,\epsilon} : \{0, 1\}^n \rightarrow \mathbb{R}$ of degree $O((1/\epsilon)^{1/3} \log s)$, such that for every $x \in \{0, 1\}^n$, when we draw a random $L_{n,s;t,\epsilon} \in \mathcal{L}_{n,s;t,\epsilon}$:

if $\prod_{i=1}^n x_i \leq t$, then $|L_{n;s;t}| \leq 1$ with probability at least $1 - \epsilon$;
 if $\prod_{i=1}^n x_i \geq 2(t; t + \epsilon^n)$, then $|L_{n;s;t}| > 1$ with probability at least $1 - \epsilon$;
 if $\prod_{i=1}^n x_i \geq t + \epsilon^n$, then $|L_{n;s;t}| \leq s$ with probability at least $1 - \epsilon$.

For the exact setting with $\epsilon = 1/n$, we can alternatively bound the degree by $O(n^{1.3} \log^{2.3}(ns))$.

The PTFs of Theorem 6.3 can be sampled using only $O(\log(n) \log(ns))$ random bits as well; their lower degree will allow us to design faster randomized algorithms for a variety of problems. For emphasis, we will sometimes refer to PTFs as deterministic PTFs to distinguish them from probabilistic PTFs.

6.1.2 Algorithmic Applications

Next, by combining these polynomials with the aforementioned polynomial method in algorithm design (and in particular, making use of fast rectangular matrix multiplication algorithms to quickly evaluate polynomials on many inputs), we design new faster algorithms for many different problems in Chapter 8.

Batch Hamming Nearest Neighbor Search

Recall the Hamming nearest neighbor problem (HNN): given a set D of n database points in the d -dimensional hypercube $\{0, 1\}^d$, we wish to preprocess D to support queries of the form $q \in \{0, 1\}^d$, where a query answer is a point $u \in D$ that differs from q in a minimum number of coordinates. Minsky and Papert [MP69, Chapter 12.7] called this the Best Match problem, and it has been widely studied since. Like many situations where one wants to find points that are most similar to query points, HNN is fundamental to modern computing, especially in search and error correction [Ind04]. However, known exact solutions to the problem require a data structure of $2^{(d)}$ size (storing all possible queries) or query time $\Omega(n \cdot \text{poly}(\log n))$ (trying nearly all the points in the database). This is one of many examples of the curse of dimensionality phenomenon in search, with corresponding data structure lower bounds. For instance, Barkol and Rabani [BR02] show a size-query tradeoff for HNN in d dimensions in the cell-probe model: if one uses b cells of size ℓ to store the database and probes at most t cells in a query, then either $s = 2^{(\ell \cdot d \cdot t)}$ or $b = n^{(1) \cdot t}$.

During the late 90's, a new direction opened in the search for better nearest neighbor algorithms. The driving intuition was that it may be easier to find and generally good enough to have approximate solutions: points with distance within $(1 + \epsilon)$ of the optimum. Utilizing novel hashing and dimensionality reduction techniques, this beautiful line of work has had enormous impact [Kle97, IM98, KOR00, Pan06, AI06, Val15, AINR14, AR15]. Still, when turning to approximations, the exponential-in- d dependence generally turns into an exponential-in- ϵ^{-1} dependence, leading to a curse of approximation [Pat08], with lower bounds matching this intuition [CCGL99,

[CR04, AIP06]. For example, Andoni, Indyk, and Patrascu [AIP06] prove that any data structure for $(1 + \epsilon)$ -approximate HNN using $O(1)$ probes requires $\Omega(n^{1-\epsilon^2})$ space.

In our first application, we design new algorithms for the natural d -line version of HNN. We design faster algorithms for both the exact and the approximate version, in the Hamming metric as well as in other metrics like ℓ_1 and the Jaccard distance.

Offline Hamming Nearest Neighbor Search. We first revisit exact nearest neighbors in the Hamming metric. We study the natural d -line problem of answering n Hamming nearest neighbor queries at once, on a database of size n . We call this the Batch Hamming Nearest Neighbor problem (BHNN). Here the aforementioned data structure lower bounds no longer apply there is no information bottleneck. Nevertheless, known algorithms for BHNN still run in either about $n^2 d^{(1)}$ time (try all pairs) [GL01, MKZ09] or about $n 2^{(d)}$ time (build a table of all possible query answers). Using our probabilistic PTFs, we improve over both these bounds for $\log n \leq d \leq \log^3 n = \log^5 \log n$.

Theorem 6.4. Given n red and n blue points in $\{0, 1\}^d$ for $d = c \log n$ and $\log^3 n = \log^5 \log n$, we can find an (exact) Hamming nearest/farthest blue neighbor for every red point in randomized time $n^{2 - \Omega(\frac{c}{\log^3 c})}$.

Using the same ideas, we are also able to derandomize our algorithm, to achieve deterministic time $n^{2 - \Omega(c \log^2 c)}$. When $d = c \log n$ for constant c , these algorithms both have truly subquadratic running times. We then apply simple reductions to achieve similar running times for finding closest pairs in ℓ_1 for vectors with small integer entries, and pairs with maximum inner product or Jaccard coefficient, as well as Bichromatic Min Inner Product: given an integer k and a collection of red and blue Boolean vectors, determine if there is a red and blue vector with inner product at most k .

It is important to keep in mind that sufficiently fast d -line Hamming closest pair algorithms would yield a breakthrough in satisfiability algorithms, so there is a potential limit. Indeed, we show:

Theorem 6.5. Suppose there is $\epsilon > 0$ such that for all constant c , Bichromatic Hamming Closest Pair can be solved in $2^{o(d)} n^{2 - \epsilon}$ time on a set of n points in $\{0, 1\}^{c \log n}$. Then the Strong Exponential Time Hypothesis (SETH) is false.

The proof is actually a reduction from the (harder-looking) Orthogonal Vectors problem, where it is well-known that $n^{2 - \epsilon}$ time would refute SETH [Wil05]. Our algorithm for Theorem 6.4 shows that for all c , there is a $\epsilon > 0$ such that Offline Hamming Nearest Neighbor search in dimension $d = c \log n$ takes $O(n^{2 - \epsilon})$ time. Theorem 6.5 says that showing that there is a universal $\epsilon > 0$ that works for all c would disprove the Strong Exponential Time Hypothesis.

Offline Approximate Nearest Neighbor Search. The problem of finding high-dimensional approximate nearest neighbors has received even more attention than the

exact variant. Locality-sensitive hashing yields data structures that can $(1 + \epsilon)$ -factor approximate nearest neighbors to any query point in $\Theta(dn^{1+\epsilon})$ (randomized) time after preprocessing in $\Theta(dn + n^{2+\epsilon})$ time and space, for not only Hamming space but also ℓ_1 and ℓ_2 space [HIM12, AI06]. Thus, a batch of n queries can be answered in $\Theta(dn^{2+\epsilon})$ randomized time. Exciting recent work on locality-sensitive hashing [AINR14, AR15] has improved the constant factor in the $(1 + \epsilon)$ bound, but not the growth rate in n . In 2012, Gregory Valiant [Val15] reported a surprising algorithm running in $\Theta(dn + n^{2+\epsilon})$ randomized time for the offline version of the problem in ℓ_2 . We obtain a still faster algorithm for the offline problem, with ϵ improved to about $\epsilon = 1/3$:

Theorem 6.6. Given n red and n blue points in $[U]^d$ and $\epsilon = \frac{\log^6 \log n}{\log^3 n}$, we can find a $(1 + \epsilon)$ -approximate ℓ_1 or ℓ_2 nearest/farthest blue neighbor for each red point in $(dn + n^{2+\epsilon}) \text{ poly}(\log(nU))$ randomized time.

Valiant's algorithm, like the previous polynomial method algorithms, relied on fast matrix multiplication. It also used Chebyshev polynomials but in a seemingly more complicated way. Our new probabilistic PTF construction is inspired by our attempt to unify Valiant's approach with the probabilistic method, which leads to an improvement of Valiant's algorithm. (We also almost succeed in derandomizing Valiant's $n^{2+\epsilon}$ result in the Hamming case, except for an initial dimension reduction step; see Remark 8.3 in Section 8.3.)

Numerous applications to high-dimensional computational geometry follow; for example, we can approximate the diameter or Euclidean minimum spanning tree of a given set of n points in roughly the same running time.

The Light Bulb Problem. The last problem related to nearest neighbor search we study is the Light Bulb Problem, introduced by Leslie Valiant in 1988 [Val88]: Given as input a set S of n vectors from $\{-1, 1\}^d$, which are all independently and uniformly random except for two planted vectors (the correlated pair) which have inner product at least ϵd for some $0 < \epsilon < 1$, the goal is to find the correlated pair. This is a basic formulation of the problem of finding correlated variables in data analysis, and the best known algorithms for more general problems like finding correlations on the Euclidean sphere [Cha02] and learning sparse parities with noise [Val15, Appendix A] come from reductions to the Light Bulb problem.

The dimension d of the vectors is called the sample complexity of the problem, since it corresponds to the number of data points which must be gathered about the variables in order to determine which are correlated. When d is too small (for instance, $d < \log(n)$), then the problem is information-theoretically impossible. By standard concentration inequalities, there is a constant $c > 1$ such that, whenever $d \geq c \log n$, the correlated pair is the closest pair of vectors with high probability. We would like to design algorithms for this $d = O(\log n)$ regime.

It is not hard to see that the Light Bulb Problem is a special case of the $(1 + \epsilon)$ approximate Hamming nearest neighbor problem which we solved in Theorem 6.6 above. However, whereas before we were concentrating on the case when

" is very small, the Light Bulb problem can be seen as the case where ϵ is instead a large constant. In other words, the result in Theorem 6.6 was optimizing for a different parameter than is necessary for the Light Bulb problem. Using techniques like Locality-Sensitive Hashing [IM98, PRR95, Dub10], one can solve the Light Bulb Problem in time $n^{2 - O(\epsilon)}$. For constant $\epsilon > 0$, this gives a truly subquadratic running time, but the running time becomes quadratic as $\epsilon \rightarrow 0$.

In a breakthrough result, G. Valiant [Val15] gave an algorithm solving the Light Bulb Problem in time $O(n^{(5 - \epsilon)/(4 - \epsilon)} + nd) < O(n^{1.615} + nd)$, where $\epsilon < 2/373$ is the exponent of matrix multiplication, for any constant $\epsilon > 0$, no matter how small. Thereafter, Karppa et al. [KKK16] gave an improved algorithm with a running time of $O(n^{2/3 + \epsilon} + nd) < O(n^{1.582} + nd)$. Both of these algorithms work when the sample complexity matches, up to a constant, the information-theoretically necessary $d = (\log n)$.

Here, we give a new randomized algorithm with a simple analysis which matches the best known running time $O(n^{2/3 + \epsilon})$ and sample complexity $d = (\log n)$. Previous algorithms for the problem made use of sophisticated random sampling techniques, but we show that these are unnecessary when approaching the problem using the polynomial method instead.

By leveraging our simpler analysis, we also give new faster deterministic algorithms for the problem. However, as the inputs to the Light Bulb Problem come from a random distribution, we need to be careful about what a deterministic algorithm means. We give algorithms in two different settings:

an algorithm running in the same time $O(n^{2/3 + \epsilon}) < O(n^{1.582})$ for sample complexity $d = (\log n)$ which is correct on almost all instances (i.e. the probability of drawing an instance where the algorithm fails is $\leq \text{poly}(n)^{-c}$), and

an algorithm running in time $O(n^{4/5 + \epsilon}) < O(n^{1.899})$ for sample complexity $d = (\log n)$ which must correctly solve every instance, given the promise that the pairs of vectors other than the correlated pair are not much more correlated than one would expect random vectors to be.

See Subsection 8.12 for more details. In both of these settings, the previous best known running time [KKK16] was at best $O(n^{1.996})$.

Satisfiability Algorithms

Next, we apply our polynomials to design faster satisfiability algorithms in a number of different settings which involve threshold functions and counting.

MAX-SAT. We begin with MAX-SAT, the problem of finding an assignment that satisfies the maximum number of clauses in a given CNF formula with n variables. In the sparse case when the number of clauses is m , a series of papers have given faster exact algorithms, for example, achieving $2^{n - O(c \log c)}$ time by Dantsin and Wolpert [DW06a], $2^{n - O((c \log c)^2)}$ time by Sakai et al. [SST15a], and $2^{n - O(c^2)}$ time by Chen and Santhanam [CS15]. Using the polynomial method and our new probabilistic PTF construction, we obtain the following improved result:

Theorem 6.7. Given a CNF formula with n variables and $c n^4 = \log^{10} n$ clauses, we can find an assignment that satisfies the maximum number of clauses in randomized $2^n \cdot n = O(c^{1=3} \log^{7=3} c)$ time.

For general dense instances, the problem becomes tougher. Williams [Wil05] gave an $O(2^{0.792n})$ -time algorithm for MAX-2-SAT, but an $O(2^{(1-\epsilon)n})$ -time algorithm for MAX-3-SAT (for a universal $\epsilon > 0$) has remained open; currently the best reported time bound [SST15b] is $2^n \cdot (n = \log n)^{1=3}$, which can be slightly improved to $2^n \cdot (n = \log n)$ with more care. We make new progress on not only MAX-3-SAT but also MAX-4-SAT:

Theorem 6.8. Given a weighted 4-CNF formula F with n variables with positive integer weights bounded by $p(n)$, we can find an assignment that maximizes the total weight of clauses satisfied in F , in randomized $2^n \cdot n = O(\log^2 n \log^2 \log n)$ time. In the sparse case when the clauses have total weight n , the time bound improves to $2^n \cdot n = O(\log^2 c \log^2 \log c)$.

LTF-LTF Circuit SAT Algorithms and Lower Bounds. Using our small sample space for probabilistic polynomials for threshold functions (Theorem 6.1), we construct a new circuit satisfiability algorithm for circuits with linear threshold functions (LTFs) which improves over several prior results. Let $AC^0[d; m]$ LTF LTF[$S_1; S_2; S_3$] be the class of circuits with a layer of S_3 LTFs at the bottom layer (nearest the inputs), a layer of S_2 LTFs above the bottom layer, and a size S_1 $AC^0[m]$ circuit of depth d above the two LTF layers.¹

Theorem 6.9. For every integer $d > 0$, $m > 1$, and $\epsilon > 0$, there is an $n > 0$ and an algorithm for satisfiability of $AC^0[d; m]$ LTF LTF[$2^n; 2^n; n^2$] circuits that runs in deterministic $2^n \cdot n^\epsilon$ time.

Williams [Wil14b] gave a comparable SAT algorithm for ACC^0 LTF circuits of 2^n size, where $\epsilon > 0$ is sufficiently small.² Theorem 6.9 strictly generalizes the previous algorithm, allowing another layer of 2^n linear threshold functions below the existing LTF layer. Theorem 6.9 also trivially implies deterministic SAT algorithms for LTF LTF circuits of up to $n^{2-\epsilon(1)}$ gates, improving over the recent SAT algorithms of Chen, Santhanam, and Srinivasan [CSS16] which only work for $1+\epsilon$ -wire circuits for $\epsilon > 1$, and the SAT algorithms of Impagliazzo, Paturi, and Schneider [IPS13].

Applying the known connection between circuit satisfiability algorithms and circuit lower bounds for E^{NP} problems [Wil13, Wil14c, JMV15], the following is immediate:

Corollary 6.1. For every $d > 0$, $m > 1$, and $\epsilon \in (0; 1)$, there is an $n > 0$ such that the class E^{NP} does not have non-uniform circuits in $AC^0[d; m]$ LTF LTF[$2^n; 2^n; n^2$].

¹Recall that for an integer $m \geq 2$, $AC^0[m]$ refers to constant-depth unbounded fan-in circuits over the basis $\{AND; OR; MOD_m\}$, where MOD_m outputs 1 if the sum of its input bits is divisible by m .

²Recall ACC^0 is the infinite union of $AC^0[m]$ for all integers $m \geq 2$.

In particular, for every $\epsilon > 0$, E^{NP} does not have ACC^0 LTF LTF circuits where the ACC^0 LTF subcircuit has $2^{n^{o(1)}}$ size and the bottom LTF layer has $n^{2-\epsilon}$ gates.

Most notably, Corollary 6.1 proves lower bounds with $n^{2-\epsilon}$ LTFs on the bottom layer and subexponentially many LTFs on the second layer. This improves upon recent LTF LTF gate lower bounds of Kane and Williams [KW16], at the cost of raising the complexity of the hard function from TC_3^0 to E^{NP} . Suguru Tamaki [Tam16] has recently reported similar results for depth-two circuits with both symmetric and threshold gates.

A Powerful Randomized SAT Algorithm. Finally, combining the probabilistic PTF (Theorem 6.3) and probabilistic polynomial (Theorem 6.1) for threshold functions, we give a randomized SAT algorithm for a rather powerful class of circuits. The class $MAJ-ACC^0-LTF-ACC^0-LTF$ denotes the class of circuits with a majority gate at the top, along with two layers of linear threshold gates, and arbitrary $O(1)$ -depth ACC^0 circuitry between these three layers. This circuit class is arguably much more powerful than TC_3^0 ($MAJ-MAJ-MAJ$), based on known low-depth circuit constructions for arithmetic functions (e.g. [CSV84, MT98, MT99]).

Theorem 6.10. For all $\epsilon > 0$ and integers $d \geq 1$, there is a $\delta > 0$ and a randomized satisfiability algorithm for $MAJ-ACC^0-LTF-ACC^0-LTF$ circuits of depth d running in $2^{n^{d-\delta}}$ time, on circuits with the following properties:

- the top MAJ gate, along with every LTF on the middle layer, has $O(n^{6-5\epsilon})$ fan-in, and
- there are $O(2^n)$ many AND/OR gates (anywhere) and LTF gates at the bottom layer.

Theorem 6.10 applies the probabilistic PTF of degree about n^{d-3} (Theorem 6.3) to the top MAJ gate, probabilistic polynomials over \mathbb{Z} of degree about $n^{1-2\epsilon}$ (Theorem 6.1) to the middle LTFs, and weight reduction to the bottom LTFs; the rest can be represented with $\text{poly}(n)$ probabilistic degree.

It would not be surprising if the above circuit class contained strong pseudorandom function candidates; that is, it seems likely that the Natural Proofs barrier applies to this circuit class. Hence from the circuit lower bounds perspective, the problem of derandomizing the SAT algorithm of Theorem 6.10 is extremely interesting.

6.2 Other Related Work

Chebyshev Polynomials in Theoretical Computer Science. New applications of Chebyshev polynomials to algorithm design are a key component of the algorithms in this Part. This is certainly not a new phenomenon in itself; here we briefly survey some prior related usages of Chebyshev polynomials. First, Nisan and Szegedy [NS94] used Chebyshev polynomials to compute the OR function on Boolean variables with an approximating polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$, such that for all $x \in \{0, 1\}^n$ we have $|OR(x) - p(x)| \leq \epsilon$, yet $\deg(p) = O(\frac{1}{\epsilon})$. They also proved the degree bound is

tight up to constants in the big-O; Paturi [Pat92] generalized the upper and lower bound to all symmetric functions.

This work has led to several advances in learning theory. Building on the polynomials of Nisan and Szegedy, Klivans and Servedio [KS01] showed how to compute an OR of t ANDs of w variables with a PTF of degree $O(\sqrt{tw \log t})$, similar to our degree bound for computing an OR of t MAJORITIES of w variables of Theorem 7.5 (however, note our bound in the exact setting is a bit better, due to our use of discrete Chebyshev polynomials). They also show how to compute an OR of t ANDs on n variables with a deterministic PTF of $O(n^{1/3} \log t)$ degree, similar to our cube-root-degree probabilistic PTF for the OR of MAJORITY of Theorem 7.6 in the exact setting. However, it looks difficult to generalize Klivans-Servedio's $O(n^{1/3} \log t)$ degree bound to compute an OR of MAJORITY: part of their construction uses a reduction to decision lists which works for conjunctions but not for MAJORITY functions. Klivans, O'Donnell and Servedio [KOS04] show how to compute an AND of MAJORITY on n variables with a PTF of degree $O(\sqrt{nw \log k})$. By a simple transformation via De Morgan's law, there is a polynomial for OR of MAJORITY with the same degree. Their degree is only slightly worse than ours in terms of k (because we use discrete Chebyshev polynomials).

In streaming algorithms, Harvey, Nelson, and Onak [HNO08] use Chebyshev polynomials to design efficient algorithms for computing various notions of entropy in a stream. As a consequence of a query upper bound in quantum computing, Ambainis et al. [ACR⁺10] show how to approximate any Boolean formula of size n with a polynomial of degree $n^{1/2 + o(1)}$, improving on earlier bounds of O'Donnell and Servedio [OS10] that use Chebyshev polynomials. Sachdeva and Vishnoi [SV13] give applications of Chebyshev polynomials to graph algorithms and matrix algebra. Linial and Nisan [LN90] use Chebyshev polynomials to approximate inclusion-exclusion formulas, and Sherstov [She08] extends this to arbitrary symmetric functions.

Further Applications of our Polynomials for Threshold Functions. Since the publication of a preliminary version of our polynomial constructions [AW15, ACW16], other researchers have applied them to even more problems in matching and computational geometry. Moeller et al. [MPS16] use them to give a subquadratic time algorithm for finding stable matchings in the case when the preference lists are given by a succinct representation rather than a quadratic-size list of lists. Chan [Cha18] applies Chebyshev polynomials to a variety of problems in low-dimensional computational geometry such as approximate nearest neighbor search and constructing ϵ -kernels.

Hardness of Approximate Nearest Neighbor Search. In our Theorem 6.5, we proved a lower bound on the running time of algorithms for exact batch nearest neighbor search assuming the Strong Exponential Time Hypothesis (SETH). In follow-up work, Rubinfeld [Rub18] showed a similar hardness result for approximate batch nearest neighbor search: assuming SETH, for every $\epsilon > 0$, there is an $n > 0$ such that the $(1 + \epsilon)$ -approximate batch Hamming nearest neighbor problem on

input points requires (n^2) time. Assuming SETH, this gives a limit to how much one can improve the running time of our algorithm in Theorem 6.6.

6.3 Bibliographic Details

This Part of the dissertation is based on the results in three previously published papers:

‘Probabilistic Polynomials and Hamming Nearest Neighbors’ with Ryan Williams [AW15], which appeared in FOCS 2015,

‘Polynomial Representations of Threshold Functions and Algorithmic Applications’ with Timothy M. Chan and Ryan Williams [ACW16], which appeared in FOCS 2016, and

‘An Illuminating Algorithm for the Light Bulb Problem’ [Alm19a], which appeared in SOSA 2019.

Subsection 7.2.1, Section 7.4, Subsection 7.5.1, and Section 8.2 present results from [AW15]. Section 8.4 presents results from [Alm19a]. Subsection 7.5.2, Section 7.6, and the remainder of Chapter 8 present results from [ACW16]. The earlier Sections of Chapter 7 give an introduction to the theory of probabilistic polynomials; we cite the original sources of the results therein when appropriate.

Chapter 7

Probabilistic Polynomials

7.1 Multilinear Polynomials Computing Boolean Functions

The main topic of this chapter is polynomial representations of Boolean functions. We focus in particular on multilinear polynomials.

Definition 7.1. A multilinear polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ over a commutative ring R is an exact polynomial for the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ if we have $p(x) = f(x)$ for all $x \in \{0, 1\}^n$.

Example 7.1. The polynomial $p(x_1; x_2; \dots; x_n) = x_1 x_2 \dots x_n$ exactly computes $\text{AND}(x_1; x_2; \dots; x_n)$, since p outputs 1 when all its inputs are 1, and it outputs 0 when any of its inputs is 0. Similarly, since we can write $\text{OR}(x_1; x_2; \dots; x_n) = 1 - \text{AND}(1 - x_1; 1 - x_2; \dots; 1 - x_n)$, it follows that $\text{OR}(x_1; x_2; \dots; x_n)$ is exactly computed by the polynomial $1 - p(1 - x_1; 1 - x_2; \dots; 1 - x_n) = 1 - (1 - x_1)(1 - x_2) \dots (1 - x_n)$.

The exact polynomial for a Boolean function f can be seen as a change of basis of f . To be more precise, we need some definitions.

For a subset $T \subseteq [n]$, let $I(T) \in \{0, 1\}^n$ denote the indicator vector for T , which has $I(T)_i = 1$ when $i \in T$ and $I(T)_i = 0$ when $i \notin T$.

For an n -input Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let $V(f) \in \{0, 1\}^{2^n}$, whose entries are indexed by subsets of $[n]$, be the truth table vector of f , which has $V(f)_T = f(I(T))$ for all $T \subseteq [n]$.

For $n \in \mathbb{N}$, define the matrix $M_{\text{SUB};n} \in \{0, 1\}^{2^n \times 2^n}$, whose rows and columns are indexed by subsets of $[n]$, and whose entry $M_{\text{SUB};n}[T; S]$ for $T, S \subseteq [n]$ is given by

$$M_{\text{SUB};n}[T; S] = \begin{cases} 1 & \text{if } S \subseteq T; \\ 0 & \text{otherwise.} \end{cases}$$

Since $M_{\text{SUB};n}$ is an upper-triangular matrix with all 1s on the diagonal, it has determinant 1, so it has full rank and a unique inverse over any commutative ring.

Finally, any multilinear polynomial $p : R^n \rightarrow R$ over a ring R can be written as

$$p(x_1, \dots, x_n) = \sum_{S \subseteq [n]} \sum_{i \in S} c_S x_i \quad (7.1)$$

Let $(c) \in R^{2^n}$, whose entries are indexed by subsets $S \subseteq [n]$, denote the coefficient vector of p , given by $(c)_S = c_S$, the coefficient from (7.1).

Evidently, a Boolean function f is in bijection with its truth table vector $V(f)$, and a multilinear polynomial p is in bijection with its coefficient vector (c) . When p is an exact polynomial for f , then these two vectors are a change of basis of one another:

Proposition 7.1. The multilinear polynomial $p : R^n \rightarrow R$ over a commutative ring R is an exact polynomial for the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ if and only if $V(f) = M_{\text{SUB};n} (c)$.

Proof. For any subset $T \subseteq [n]$, we have

$$p(\mathbb{1}(T)) = \sum_{S \subseteq T} (c)_S = [M_{\text{SUB};n} (c)]_T$$

We therefore have $p(\mathbb{1}(T)) = f(\mathbb{1}(T))$ if and only if $[M_{\text{SUB};n} (c)]_T = V(f)_T$. This must hold for all $T \subseteq [n]$ for p to exactly compute f . \square

Corollary 7.1. For any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and any commutative ring R , there is a unique multilinear polynomial $p : R^n \rightarrow R$ such that $p(x) = f(x)$ for all $x \in \{0, 1\}^n$.

So far, we have written any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ in the 'AND basis', i.e. the basis of monomials $\prod_{i \in S} x_i$ for $S \subseteq [n]$ which compute the AND on a subset of the inputs. One can similarly see that over any commutative ring R , every Boolean function f also has a unique representation over some other choices of basis, including:

The 'NOR basis' of functions $\prod_{i \in S} (1 - x_i)$ for $S \subseteq [n]$.

The 'XOR basis' of functions $\prod_{i \in S} x_i = \frac{1}{2} \prod_{i \in S} (1 + 2x_i)$ for $S \subseteq [n]$ (whenever the characteristic of the ring R is not positive and even).

7.2 Typically Correct Polynomials

One of the main goals in the polynomial method is to design low-degree polynomial representations of Boolean functions. As we saw in Example 7.1 above, even very simple Boolean functions like AND and OR on n inputs require degree n to compute exactly, the maximum possible degree of a multilinear polynomial on n inputs. In order to achieve a lower degree, we need to weaken the requirements on our polynomials. One natural way to do so is to require the polynomial be correct only on most inputs:

Definition 7.2. For any $\epsilon \in [0, 1]$, a multilinear polynomial $p : \mathbb{R}^n \rightarrow \mathbb{R}$ over a commutative ring R is a $(1 - \epsilon)$ -correct polynomial for the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ if we have $p(x) = f(x)$ for at least a $(1 - \epsilon)$ fraction of all $x \in \{0, 1\}^n$. The $(1 - \epsilon)$ -correct degree of f over R is the minimum degree of such a polynomial p .

Many functions have much lower $(1 - \epsilon)$ -correct degree than exact degree. For instance, although AND on n inputs has exact degree n , its $(1 - \epsilon)$ -correct degree is 0 for any $\epsilon > 0$, since the polynomial $p(x) = 0$ computes it correctly on all but one point from $\{0, 1\}^n$. Such simple constructions like this are, unfortunately, not particularly useful in applications.

That said, there are interesting Boolean functions whose $(1 - \epsilon)$ -correct degree is less trivial. Consider, for instance, the majority function MAJ on n inputs. Although MAJ has exact degree n , we will show in the remainder of this section that it has ϵ -typically correct degree $O(n \log(1/\epsilon))$ for all $\epsilon > 0$.

7.2.1 Interpolating Polynomials for Symmetric Functions

We begin in this subsection by proving the upper bound, that MAJ on n variables has $(1 - \epsilon)$ -correct degree $O(n \log(1/\epsilon))$. The key new polynomial construction we will need is an interpolating polynomial for correctly computing symmetric Boolean functions on inputs of certain Hamming weights. Such a polynomial can be derived from prior work (at least over fields [Sri13]), but for completeness, we prove its existence here.

Lemma 7.1. For any integers n, r, k with $n = k + r$ and any integers c_0, \dots, c_r , there is a multivariate polynomial $p : \{0, 1\}^n \rightarrow \mathbb{Z}$ of degree $\leq r$ with integer coefficients such that $p(x) = c_i$ for all $x \in \{0, 1\}^n$ with Hamming weight $|x| = k + i$.

Notice that it is not immediately obvious from univariate polynomial interpolation that the polynomial p exists as described, since the univariate polynomial $q : \mathbb{R} \rightarrow \mathbb{R}$ such that $q(k + i) = c_i$ typically has rational (non-integer) coefficients. Lemma 7.1 is more general than a result claimed without proof by Srinivasan ([Sri13], Lemma 14). It also generalizes of a theorem of Bhatnagar et al. ([BGL06], Theorem 2.8).

Proof. Our polynomial p will have the form

$$p(x_1, \dots, x_n) = \sum_{i=0}^{r-1} a_i \sum_{\substack{S \subseteq [n] \\ |S|=i}} \prod_{j \in S} x_j$$

for some constants $a_0, \dots, a_{r-1} \in \mathbb{Z}$. Hence, we will get that for any $x \in \{0, 1\}^n$:

$$p(x) = \sum_{i=0}^{r-1} \binom{|x|}{i} a_i$$

Define the matrix:

$$M = \begin{pmatrix} 0 & k+1 & k+1 & \dots & k+1 & 1 \\ p_1(x_1) & p_2(x_1) & \dots & p_r(x_1) & c_1 \\ p_1(x_2) & p_2(x_2) & \dots & p_r(x_2) & c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p_1(x_r) & p_2(x_r) & \dots & p_r(x_r) & c_r \end{pmatrix}$$

The conditions of the stated lemma are that

$$M = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 & 1 \\ a_0 & a_1 & \dots & a_{r-1} & c_1 \\ a_1 & a_2 & \dots & a_r & c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{r-1} & a_r & \dots & a_{r+1} & c_r \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_r \end{pmatrix} A$$

By Lemma 7.2 (proved below), M always has determinant 1. Because M is a matrix with integer entries and determinant 1, its inverse M^{-1} is also an integer matrix. Multiplying through by M^{-1} above gives integer expressions for the c_i , as desired. \square

Lemma 7.2. For any univariate polynomials $p_1; p_2; \dots; p_r$ such that p_i has degree $i - 1$, and any pairwise distinct $x_1; x_2; \dots; x_r \in \mathbb{Z}$, the matrix

$$M = \begin{pmatrix} 0 & p_1(x_1) & p_2(x_1) & \dots & p_r(x_1) & 1 \\ p_1(x_2) & p_2(x_2) & \dots & p_r(x_2) & c_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ p_1(x_r) & p_2(x_r) & \dots & p_r(x_r) & c_r \end{pmatrix}$$

has determinant

$$\det(M) = \prod_{i=1}^r c_i \prod_{1 \leq i < j \leq r} (x_j - x_i);$$

where c_i is the coefficient of x^{i-1} in p_i .

Proof. For i from 1 up to $r - 1$, we can add multiples of column i of M to the subsequent columns in order to make the coefficient of x^{i-1} in all the other columns 0. The resulting matrix is

$$M^0 = \begin{pmatrix} 0 & c_1 & c_2 x_1 & \dots & c_r x_1^{r-1} & 1 \\ c_1 & c_2 x_2 & \dots & c_r x_2^{r-1} & c_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ c_1 & c_2 x_r & \dots & c_r x_r^{r-1} & c_r \end{pmatrix}$$

This is a Vandermonde matrix which has the desired determinant. \square

Since MAJ is a symmetric function, we can use Lemma 7.1 to get a $(1 - \epsilon)$ -correct degree upper bound for it:

Corollary 7.2. Over any commutative ring R , MAJ on n inputs has $(1 - \epsilon)$ -correct degree $O\left(\frac{n}{\epsilon \log(1/\epsilon)}\right)$ for all $\epsilon \in (0, 1)$.

Proof. Applying Lemma 7.1, we know there is a polynomial $p : \{0, 1\}^n \rightarrow \mathbb{Z}$ over \mathbb{Z} of degree $O\left(\frac{n}{\epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$ with integer coefficients such that $p(x) = \text{MAJ}(x)$ for all $x \in \{0, 1\}^n$ such that $\sum_{i=1}^n x_i = 2j$ for $j \in \left[\frac{n}{2} - \frac{n}{2\epsilon}, \frac{n}{2} + \frac{n}{2\epsilon}\right]$. We can view p as a polynomial over \mathbb{R} (possibly by taking p mod the characteristic of \mathbb{R} when \mathbb{R} has positive characteristic) which is correct for all such x as well.

It remains to show that the fraction of $x \in \{0, 1\}^n$ with $\sum_{i=1}^n x_i = 2j$ for $j \in \left[\frac{n}{2} - \frac{n}{2\epsilon}, \frac{n}{2} + \frac{n}{2\epsilon}\right]$ is at least $1 - \epsilon$. This is equivalent to showing that, if we draw $x \in \{0, 1\}^n$ uniformly at random, then the probability that $\sum_{i=1}^n x_i = 2j$ for $j \in \left[\frac{n}{2} - \frac{n}{2\epsilon}, \frac{n}{2} + \frac{n}{2\epsilon}\right]$ is at most ϵ . By Hoeffding's inequality (Lemma 2.1 from the Preliminaries) we see that

$$\Pr \left[\sum_{i=1}^n x_i \leq \frac{n}{2} - \frac{n}{2\epsilon} \right] \leq \exp \left(- \frac{2 \left(\frac{n}{2\epsilon} \right)^2}{n} \right) = \exp \left(- \frac{n}{2\epsilon^2} \right) < \frac{\epsilon}{2};$$

By symmetry, we also have $\Pr \left[\sum_{i=1}^n x_i \geq \frac{n}{2} + \frac{n}{2\epsilon} \right] < \frac{\epsilon}{2}$, and so in total, we have that the probability of $\sum_{i=1}^n x_i = 2j$ for $j \in \left[\frac{n}{2} - \frac{n}{2\epsilon}, \frac{n}{2} + \frac{n}{2\epsilon}\right]$ is at most ϵ , as desired. \square

7.2.2 The Razborov-Smolensky Lower Bound

We now prove a matching lower bound, that MAJ on n variables requires $(1 - \epsilon)$ -correct degree $\Omega\left(\frac{n}{\epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$. This is the classic result of Razborov [Raz87] and Smolensky [Smo87]; in this subsection, we present the proof technique of Razborov. We begin with a Lemma showing that low-degree $(1 - \epsilon)$ -correct polynomials for MAJ lead to relatively low-degree $(1 - \epsilon)$ -correct polynomials for any Boolean function.

Lemma 7.3. For any commutative ring R and set $S \subseteq \{0, 1\}^n$, suppose there is a polynomial $p : \mathbb{R}^n \rightarrow R$ of degree $\deg(p) = d$, such that $p(x) = \text{MAJ}(x)$ for all $x \in S$. Then, for any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ there is a polynomial $q : \mathbb{R}^n \rightarrow R$ of degree at most $\deg(q) \leq n/2 + d$ such that $q(x) = f(x)$ for all $x \in S$.

Proof. Let $t : \mathbb{R}^n \rightarrow R$ be the exact multilinear polynomial for f over R , meaning $t(x) = f(x)$ for all $x \in \{0, 1\}^n$. We can write t out in two different ways, first over the 'AND basis' of monomials:

$$t(x) = \sum_{T \subseteq [n]} a_T \prod_{i \in T} x_i;$$

and second over the 'NOR basis':

$$t(x) = \sum_{T \subseteq [n]} b_T \prod_{i \in T} (1 - x_i);$$

where the $a_T, b_T \in R$ are the appropriate coefficients.

Notice that if $x \in \{0, 1\}^n$ is such that $\text{MAJ}(x) = 0$, then for any $T \subseteq [n]$ with

$jTj > n=2$ we have $\prod_{i \in T} x_i = 0$. Hence,

$$\text{If } \text{MAJ}(x) = 0; \text{ then } f(x) = \prod_{T \subseteq [n], |T|=2} \prod_{i \in T} x_i$$

Similarly,

$$\text{If } \text{MAJ}(x) = 1; \text{ then } f(x) = \prod_{T \subseteq [n], |T|=2} \prod_{i \in T} (1 - x_i)$$

Combining, we see that for all $x \in \{0, 1\}^n$,

$$f(x) = \prod_{T \subseteq [n], |T|=2} \prod_{i \in T} x_i + \prod_{T \subseteq [n], |T|=2} \prod_{i \in T} (1 - x_i)$$

Substituting p for MAJ above gives the desired polynomial. \square

We can now prove our lower bound:

Theorem 7.1 ([Raz87, Smo87]) There is a constant $c > 0$ such that, for every $n \geq 2$ ($n \geq 2$), every commutative ring R (other than the trivial ring), and every $(1 - \epsilon)$ -correct polynomial p for MAJ over R , the degree of p is at least $c \cdot n \log(1/\epsilon)$.

Proof. Assume to the contrary that there is such a polynomial $p : \{0, 1\}^n \rightarrow R$ of degree $d < c \cdot n \log(1/\epsilon)$. Let $S \subseteq \{0, 1\}^n$ be the set of x such that $p(x) = \text{MAJ}(x)$; by assumption we have $|S| \geq (1 - \epsilon) \cdot 2^n$. In particular, by Lemma 7.3, for every $s \in S$, there is a multilinear polynomial $p_s : \{0, 1\}^n \rightarrow R$ of degree at most $n - 2 + d$ such that $p_s(s) = 1$ and $p_s(x) = 0$ for all $x \in S^c$.

Consider the vector space V of R -linear combinations of the p_s polynomials, i.e. $V = \sum_{s \in S} a_s p_s(x)$. The p_s polynomials are linearly independent, since any linear combination p^0 of the p_s for $s^0 \in S$ will have $p^0(s) = 0$, whereas $p_s(s) = 1$. Hence, the dimension of V is at least $|S| \geq (1 - \epsilon) \cdot 2^n$.

Meanwhile, every polynomial in V has degree at most $n - 2 + d$, and so V is a subspace of the space V^0 of multilinear polynomials over R of degree at most $n - 2 + d$. This space V^0 is spanned by the multilinear monomials of degree at most $n - 2 + d$. The number of such monomials is

$$\sum_{i=0}^{n-2+d} \binom{n}{i} = 2^n - \sum_{i=0}^{n-2-d} \binom{n}{i} \leq 2^n - \binom{n}{n-2-d} = 2^n - \binom{n}{2+d}$$

Applying Corollary 2.1 from the Preliminaries, we can further upper bound this by:

$$2^n - \binom{n}{2+d} \leq 2^n \left(1 - 2^{-(2+d)}\right) < 2^n \left(1 - 2^{-(c^2 \log(1/\epsilon))}\right) = 2^n \left(1 - \epsilon^{c^2}\right):$$

If we pick a sufficiently small $\epsilon > 0$ then this is less than $2^n (1 - \epsilon)$. Then V^0 , a vector space of dimension less than $2^n (1 - \epsilon)$, contains as a subspace, a vector space of dimension at least $2^n (1 - \epsilon)$, a contradiction. \square

7.3 Probabilistic Polynomials

In the previous section, we showed that MAJ on n bits has $(1 - \epsilon)$ -correct degree $(\frac{n}{\epsilon} \log(1/\epsilon))$ over any commutative ring R . However, for other Boolean functions like AND and OR, the $(1 - \epsilon)$ -correct polynomials were trivial. This is because of a weakness of typically correct polynomials for a Boolean function: they can concentrate their errors on the 'hard' inputs of f . In the case of OR, the polynomial can simply get the answer wrong on the one point where it should output 1.

In our algorithmic applications below, we will need a stronger polynomial notion than this, in which the polynomial has a good chance of getting any given input correct:

Definition 7.3. For any $\epsilon \in [0, 1]$, and any commutative ring R , a probabilistic polynomial with error ϵ and degree d for the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a distribution P on polynomials $p : \{0, 1\}^n \rightarrow R$ of degree at most d over R such that, for every $x \in \{0, 1\}^n$, we have

$$\Pr_{p \sim P} [p(x) = f(x)] \geq 1 - \epsilon.$$

The ϵ -probabilistic degree of f over R is the minimum degree of a probabilistic polynomial with error ϵ for f .

A probabilistic polynomial is required to take all the different values of $f(x)$ into account by enforcing that ϵ is the worst case probability, among all $x \in \{0, 1\}^n$, that $p \sim P$ incorrectly outputs a value $p(x) \neq f(x)$.

For every $\epsilon \in (0, 1)$, every commutative ring R , and every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, we can see that the ϵ -probabilistic degree of f over R is at least the $(1 - \epsilon)$ -correct degree of f over R . We next give two examples where there is a separation between the two.

7.3.1 The Probabilistic Degree of OR

We noted earlier that the $(1 - \epsilon)$ -correct degree of OR is 0. In this subsection, we present classical polynomial constructions for OR, showing that the ϵ -probabilistic degree of OR on n inputs is $(\log(1/\epsilon))$ over the finite field F_q for prime q , and $O(\log n \log(1/\epsilon))$ over \mathbb{Z} . The upper bounds extend without much work to any commutative ring R as well.

Proposition 7.2 ([Raz87, Smo87]) For any prime q and any $\epsilon \in (0, 1)$, the ϵ -probabilistic degree of OR on n inputs over F_q is at most $(q - 1) \lceil \log_q(1/\epsilon) \rceil$.

Proof. We construct a probabilistic polynomial P for OR on n inputs. To draw a polynomial from P , pick $k = \lceil \log_q(1/\epsilon) \rceil$ independent uniformly random hash functions $h_1, \dots, h_k : [n] \rightarrow \mathbb{F}_q$, then output the polynomial

$$p(x_1, \dots, x_n) = \prod_{i=1}^k \sum_{j=1}^{q-1} h_j(i) x_i^{j-1}$$

The degree of p is $(q-1)k$ as desired.

We now prove correctness. When $\text{OR}(x) = 0$, then $x_i = 0$ for all $i \in [n]$, so p always outputs 0. Otherwise, if $\text{OR}(x) = 1$, then there is at least one $i \in [n]$ such that $x_i = 1$. It follows that, for each $j \in [k]$, the sum $\sum_{i=1}^n h_j(i) x_i^{j-1}$ is a uniformly random element of \mathbb{F}_q . In particular, with probability $1/q$ we have $\sum_{i=1}^n h_j(i) x_i^{j-1} \neq 0$, and hence $(\sum_{i=1}^n h_j(i) x_i^{j-1})^{q-1} = 1$ by Fermat's little theorem. If this is the case for any $j \in [k]$ then p outputs 1. Hence, since the hash functions are independent, the probability that p does not output 1 is q^{-k} , as desired. \square

We show next that, up to constant factors, the degree upper bound achieved in Proposition 7.2 is tight.

Proposition 7.3. For any commutative ring R (other than the trivial ring) and any $\epsilon \in (0, 1/4]$, the ϵ -probabilistic degree of OR on n inputs over R is at least $\log(1/\epsilon) - 1$.

Proof. Suppose that OR on n inputs has a probabilistic polynomial P of degree d for error ϵ over R . Let m be the biggest integer less than $\log(1/\epsilon)$, so that $\log(1/\epsilon) - 1 < m < \log(1/\epsilon)$, and note that $m < n$.

We construct a probabilistic polynomial Q for OR on m inputs, by drawing a random $p \in P$ and outputting $q(x_1, \dots, x_m) = p(x_1, \dots, x_m, 0, \dots, 0)$, where we have set $x_i = 0$ in p for all $i > m$. Since Q is just a restriction of the inputs to P , we have that Q is a probabilistic polynomial for OR on m inputs with error ϵ and degree d .

There are only 2^m different possible values for the binary inputs to Q , but Q has error $\epsilon < 2^{-m}$. Hence, there must be a polynomial in the support of Q that makes no errors, and exactly computes OR on m input bits. This polynomial must have degree at least m , so it follows that $d \geq m$. In particular, we get as desired that $d \geq \log(1/\epsilon) - 1$. \square

Our probabilistic polynomial construction in Proposition 7.2 above critically relied on the finite field we were working over. We next present another construction with slightly higher degree that works over \mathbb{Z} , and hence over any commutative ring R .

Proposition 7.4 ([ABFR94, Lemma 5.1]). For any $\epsilon \in (0, 1)$, the ϵ -probabilistic degree of OR on n inputs over \mathbb{Z} is at most $O(\log(1/\epsilon) \log n)$.

Proof. We first construct a distribution P^0 on polynomials $p^0 : \mathbb{Z}^n \rightarrow \mathbb{Z}$ of degree $O(\log n)$ over \mathbb{Z} such that

$$p^0(0; 0, \dots, 0) = 0 \text{ for all } p^0 \text{ in the support of } P^0, \text{ and}$$

for all $x \in \{0, 1\}^n$ with $\text{OR}(x) = 1$, we have $\Pr_{p^0}[p^0(x) = 1] \geq 1/2$.

Once we have constructed P^0 , we can then construct our desired probabilistic polynomial for OR with error ϵ and degree $O(\log(1/\epsilon) \log n)$ by drawing $k := \lceil \log(1/\epsilon) \rceil$ independent $p_1^0, \dots, p_k^0 \in P^0$, and outputting the polynomial $p(x) = \prod_{i=1}^k (1 - p_i^0(x))$.

We now construct P^0 . To draw a polynomial p^0 from P^0 , we first pick $m := \log n + 3$ random subsets $S_0, S_1, \dots, S_{m-1} \subseteq [n]$ as follows: set $S_0 = [n]$, then for each ℓ from 1 up to $m-1$, let S_ℓ be a uniformly random subset of $S_{\ell-1}$, which includes each element independently with probability $1/2$. We then output the polynomial

$$p^0(x) = \prod_{\ell=0}^{m-1} \prod_{i \in S_\ell} x_i;$$

which has degree $m = O(\log n)$.

For a given $x \in \{0, 1\}^n$, if $\text{OR}(x) = 0$, then $x_i = 0$ for all $i \in [n]$ and we will always have $p^0(x) = 0$. Otherwise, if $\text{OR}(x) = 1$, then for each $\ell \in \{0, 1, \dots, m-1\}$, define the random variables $s_\ell(x) := \prod_{i \in S_\ell} x_i$. Since $\text{OR}(x) = 1$, we have that $s_0(x) = 1$. The $s_\ell(x)$ form a sequence of nonnegative integers with $s_\ell(x) \leq s_{\ell-1}(x)$ for all $\ell \in [m-1]$. Moreover, each $s_\ell(x)$ is distributed as the sum of $s_{\ell-1}(x)$ independent random values from $\{0, 1\}$. If there is any $\ell \in \{0, 1, \dots, m-1\}$ such that $s_\ell(x) = 1$, then p^0 will output 1. Notice that exactly one of the following must be the case:

$s_\ell(x) > 1$ for all $\ell \in \{0, 1, \dots, m-1\}$. Since $s_0(x) = n$, we can apply a union bound to see that this occurs with probability at most $n^{-2} \sum_{\ell=1}^{m-1} \binom{m-1}{\ell} = 4/n^{-2 \log(n)} = 1/4$.

$s_0(x) = 1$.

There is some $\ell \in [m-1]$ such that $s_\ell(x) > 1$ and $s_{\ell-1}(x) = 1$. Note that, given the values $s_{\ell-1}(x)$, we have that $s_\ell(x) = 0$ with probability $2^{-s_{\ell-1}(x)}$, and $s_\ell(x) = 1$ with probability $s_{\ell-1}(x) \cdot 2^{-s_{\ell-1}(x)}$. Hence, conditioned on $s_{\ell-1}(x) = 1$, we have that $s_\ell(x) = 1$ with probability $s_{\ell-1}(x) \cdot (s_{\ell-1}(x) + 1)^{-2} = 3/4$.

Since one of the latter two cases occurs with probability at least $3/4$, and in each of those cases, there is an $\ell \in \{0, 1, \dots, m-1\}$ such that $s_\ell(x) = 1$ with probability at least $2/3$, it follows that p^0 will output 1 with probability at least $\frac{3}{4} \cdot \frac{2}{3} = \frac{1}{2}$, as desired. \square

7.3.2 The Probabilistic Degree of Biased Threshold Functions

In the previous subsection, we saw that the OR function has $(1 - \epsilon)$ -correct degree 0, and ϵ -probabilistic degree $O(\log(1/\epsilon) \log n)$ over, say, F_2 . The difference between these two degrees can grow unboundedly as $\epsilon \rightarrow 0$, but it is only a constant when ϵ is a constant. In this subsection, we present a different Boolean function for which the two degrees differ by an unbounded amount even in the constant regime.

Let $TH_{3=4} : \{0, 1\}^n \rightarrow \{0, 1\}$ be the Boolean function

$$TH_{3=4}(x) = \begin{cases} 1 & \text{if } |x| \geq 3n/4; \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 7.5. For sufficiently large n , the function $TH_{3=4}$ has $(2=3)$ -correct degree $O(\sqrt{n})$.

Proof. The polynomial $p(x) = 0$ is correct on every input $x \in \{0, 1\}^n$ with $|x| \geq 3n/4$. The number of inputs it gets wrong is at most (using the bound from Proposition 2.3 from the Preliminaries): $\sum_{i=0}^{n-3n/4} \binom{n}{i} 2^{H(1=4)n - \Omega(n)} = 2^{0.82n + o(n)}$, which is less than $\frac{1}{3}2^n$ for big enough n . \square

Proposition 7.6. The function $TH_{3=4}$ on n inputs has $(1=3)$ -probabilistic degree $O(\sqrt{n})$.

Proof. When half of its inputs are restricted to 0, the function $TH_{3=4}$ on n inputs becomes the function MAJ on $n/2$ inputs. Similar to the proof of Proposition 7.3, it follows that the $(1=3)$ -probabilistic degree of $TH_{3=4}$ on n is lower bounded by the $(1=3)$ -probabilistic degree of MAJ on $n/2$ inputs. But, by Theorem 7.1 (and the fact that the $(1=)$ -correct degree is a lower bound on the probabilistic degree) this is $O(\sqrt{n})$, as desired. \square

7.4 Probabilistic Polynomials for Majority

In the previous section, we saw examples of Boolean functions whose $(1=)$ -correct degree was much lower than their probabilistic degree. We saw in Section 7.2 that MAJ on n inputs has $(1=)$ -typically correct degree $O(\sqrt{n \log(1=)})$ over any commutative ring R . This raises the question: does MAJ on n inputs have $(1=)$ -probabilistic degree $O(\sqrt{n \log(1=)})$ as well?

In this section, we show that the answer is yes by giving a new probabilistic polynomial construction for MAJ. As we will see in the next Chapter, this probabilistic polynomial construction will be crucial in a number of our applications, including to nearest neighbor search algorithms. The previous best construction, by Srinivasan [Sri13], achieved degree $O(\sqrt{n \log(1=)} \text{ polylog}(n))$; we will see that the extra $\text{polylog}(n)$ factor would have been prohibitive in most of our applications (see Remark 8.1 in the next Chapter for more details).

Theorem 7.2. There is a probabilistic polynomial over \mathbb{Z} for MAJ on n variables with error ϵ and $\text{deg}(n; \epsilon) = O(\sqrt{n \log(1=\epsilon)})$. Furthermore, a polynomial can be sampled from the probabilistic polynomial distribution in $O(\sum_{i=0}^{d(n; \epsilon)} \binom{n}{i})$ time.

Notation. For $\alpha \in [0, 1]$, define $TH_{\alpha} : \{0, 1\}^n \rightarrow \{0, 1\}$ to be the threshold function $TH_{\alpha}(x_1, \dots, x_n) := [|\{i : x_i = 1\}| \geq \alpha n]$. In particular, $TH_{1/2} = \text{MAJ}$. We also define

NEAR_{ε; g} : f 0; 1g^n ! f 0; 1g, such that NEAR_{ε; g} (x) := [jxj=n 2 [n - ε; n + ε]]. Intuitively, NEAR_{ε; g} checks whether jxj=n is near n/2, with error ε.

In the remainder of this Section, we prove Theorem 7.2. To do so, we construct a probabilistic polynomial for TH over Z[x₁; ::; x_n] which has degree O(n log(1/ε)) and on each input is correct with probability at least 1 - ε.

Intuition for the construction. First, let us suppose jxj=n is not too close to n/2; in particular jxj=n is not within ε = O(n log(1/ε)) of n/2. Then, if we construct a new smaller vector x' by sampling 1=10 of the entries of x, it is likely that jx'j=(n=10) lies on the same side of n/2 as jxj=n. This suggests a recursive strategy: we can use our polynomial construction on the sample x'. Second, if jxj=n is close to n/2, then by interpolating, we can use an exact polynomial of degree O(n log(1/ε)) (which we call A_{n; g}) that is guaranteed to give the correct answer. To decide which of the two cases we are in, we will use a probabilistic polynomial for NEAR (on a smaller number of variables), which can itself be written as the product of two probabilistic polynomials for TH. The degree incurred by recursive calls can be adjusted to have tiny overhead, with the right parameters.

In comparison, Srinivasan [Sri13] takes a number theoretic approach. For different primes p, his polynomial uses ε⁻¹ probabilistic polynomials in order to determine the Hamming weight of the input (mod p). Then, it uses an exact polynomial inspired by the Chinese Remainder Theorem to determine the true Hamming weight of the input, and whether it is at least n=2. This approach works on a more general class of functions than ours, called W-sum determined, which are determined by a weighted sum of the input coordinates. However, the number of primes being considered inherently means that this type of approach will incur extra logarithmic degree increases. In fact, we also give a better probabilistic degree for every symmetric function.

Probabilistic Polynomial Definition. Let n be an integer for which we want to compute TH. Let A_{n; g} : f 0; 1g^n ! Z be an exact polynomial with integer coefficients of degree at most 2g^n n + 1 which gives the correct answer to TH for any vector x with jxj 2 [n - g^n n; n + g^n n], and can give arbitrary answers to other vectors. Such a polynomial A_{n; g} exists by Lemma 7.1 above.

Let M_{m; ε} : f 0; 1g^m ! Z denote the probabilistic polynomial for TH with error degree as described above for all m < n. We can assume as a base case that when m is constant, we simply use the exact polynomial for TH.

Define

$$S_{m; ε; g} (x) := (1 - M_{m; ε; g} (x)) M_{m; ε; g} (x)$$

Assuming M_{m; ε} works as prescribed (with ε error), this is a probabilistic polynomial for NEAR_{ε; g} with error at most 2ε. For x 2 f 0; 1g^n, let x' 2 f 0; 1g^{n=10} be a vector of length n=10, where each entry is an independent and uniformly random entry of x. Hence, each entry of x' is a probabilistic polynomial in x of degree 1. Let a = 10 log(1/ε). Our probabilistic polynomial for TH on n variables is defined

to be:

$$M_{n; \epsilon}(x) := A_{n; \epsilon, 2a}(x) S_{n=10; \epsilon; a=\frac{p}{n}; \epsilon=4}(\mathbf{x}) + M_{n=10; \epsilon; \epsilon=4}(\mathbf{x}) (1 - S_{n=10; \epsilon; a=\frac{p}{n}; \epsilon=4}(\mathbf{x})).$$

Note that \mathbf{x} denotes the same randomly chosen vector in each of its appearances, and $S_{n=10; \epsilon; a=\frac{p}{n}; \epsilon=4}$ denotes the same draw from the random polynomial distribution in both of its appearances.

Degree of $M_{n; \epsilon}$. First we show by induction on n that $M_{n; \epsilon}$ has degree at most $\frac{p}{n} \ln(1/\epsilon)$. Assume that $M_{m; \epsilon}$ has degree at most $\frac{p}{m} \ln(1/\epsilon)$ for all $m < n$. The degree of $M_{n; \epsilon}$ is thus equal to

$$\begin{aligned} & \max \deg A_{n; \epsilon, 2a}(x) S_{n=10; \epsilon; a=\frac{p}{n}; \epsilon=4}(\mathbf{x}) + \deg M_{n=10; \epsilon; \epsilon=4}(\mathbf{x}) (1 - S_{n=10; \epsilon; a=\frac{p}{n}; \epsilon=4}(\mathbf{x})) \\ &= \deg(S_{n=10; \epsilon; a=\frac{p}{n}; \epsilon=4}(\mathbf{x})) + \max \{ \deg(A_{n; \epsilon, 2a}(x)); \deg(M_{n=10; \epsilon; \epsilon=4}(\mathbf{x})) \} \\ &= 2 \frac{p}{n} \ln(1/\epsilon) + \max \{ 4 \frac{p}{n}; 2 \frac{p}{n} \ln(1/\epsilon) \} \\ &= 2 \frac{p}{n} \ln(1/\epsilon) + \max \{ 4 \frac{p}{n}; 2 \frac{p}{n} \ln(1/\epsilon) \} \\ &= 3 \frac{p}{n} \ln(1/\epsilon) \end{aligned}$$

Time to compute $M_{n; \epsilon}$. Computing $A_{n; \epsilon, 2a}$ can be done in $\text{poly}(n)$ time as described in Lemma 7.1, as can sampling from x . Given the three recursive polynomials, we can then compute $M_{n; \epsilon}$ in three multiplications. Each recursive polynomial has degree at most $\frac{p}{n} \ln(1/\epsilon)$, and hence at most $\frac{p}{n} \ln(1/\epsilon)$ monomials. Since the time for these multiplications dominates the time for the recursive computations, the total time is $\mathcal{O}(\frac{p}{n} \ln(1/\epsilon))$ using the fast Fourier transform, as desired.

Correctness. Now we prove that $M_{n; \epsilon}$ correctly simulates TH with probability at least $1 - \epsilon$, on all possible inputs. We begin by citing a lemma explaining our choice of the parameter ϵ .

Lemma 7.4. If $x \in \{0, 1\}^n$ with $|x| = w$, and $\mathbf{x} \in \{0, 1\}^{n=10}$ is a vector each of whose entries is an independent and uniformly random entry of $\{0, 1\}$ with $|x| = (n=10) = \frac{n}{4}$, then for every $\epsilon < 1/4$,

$$\Pr |v - w| \geq \frac{p}{n} \ln(1/\epsilon);$$

where $a = \frac{p}{n} \ln(1/\epsilon)$.

¹By replacing each variable with increasing powers of a single variable, we can reduce multivariate polynomial multiplication to single variable polynomial multiplication.

Proof. Each entry of \mathbf{x} is drawn from a binomial distribution with probability w of giving a 1. Hence, applying Hoeffding's inequality, Lemma 2.1 from the Preliminaries, with $p = w$, $m = n=10$, and $k = \frac{n}{10}(w - a = \bar{n}) = \frac{nw}{10} - \frac{a \bar{n}}{10}$ yields:

$$\Pr[|v - w - a = \bar{n}| \geq \frac{nw}{10} - \frac{a \bar{n}}{10}] \leq \exp\left(-2 \frac{\left(\frac{nw}{10} - \frac{a \bar{n}}{10}\right)^2}{n}\right);$$

which simplifies to $\exp\left(-\frac{a^2}{5}\right) = \exp(-2 \ln(1 + \epsilon)) = \epsilon^2 < \frac{1}{4}$. □

We now move on to the main proof of correctness, which proceeds by induction on n . By symmetry, we may assume we have an input vector $\mathbf{x} \in \{0, 1\}^n$ with $|\mathbf{x}| = n$, and we want to show that $M_{n; a}(\mathbf{x})$ outputs 1 with probability at least $1 - \epsilon$. We assume $\epsilon < 1/4$ so that we may apply Lemma 7.4.

For notational convenience, define the intervals:

$$I_0 = [a = \bar{n}; \infty); I_1 = [0; a = \bar{n}); I_2 = [a = \bar{n}; 2a = \bar{n}); I_3 = [2a = \bar{n}; 1];$$

Note that depending on the values of ϵ and a , some of these intervals may be empty; this is not a problem for our proof.

Let $w = |\mathbf{x}|/n$. Let \mathbf{x} be the random subvector of \mathbf{x} selected in $M_{n; a}$ (recall we use the same \mathbf{x} in each of the three locations it appears in the definition of M). Let $v = |\mathbf{x}|/(n=10)$. Our proof strategy is to consider different cases depending on the value of w . For each case, we show there are at most four events such that, if all events hold then $M_{n; a}$ outputs the correct answer, and each event does not hold with probability at most $\frac{\epsilon}{4}$. By the union bound, this implies that $M_{n; a}$ gives the correct answer with probability at least $1 - \epsilon$. The cases are as follows:

1. $w \geq \frac{1}{4}$ ($|\mathbf{x}| = n$ is very close to $\frac{n}{4}$). By Lemma 7.4, we know that with probability at least $1 - \frac{\epsilon}{4}$, we have $|v - w| \leq \frac{\epsilon}{4}$. In other words, $v \in [a = \bar{n} - \frac{\epsilon}{4}, a = \bar{n} + \frac{\epsilon}{4}]$.

$v \geq a = \bar{n} - \frac{\epsilon}{4}$, then with probability at least $1 - \frac{\epsilon}{4}$, we have $S_{n=10; a = \bar{n} - \frac{\epsilon}{4}}(\mathbf{x}) = 1$, by our inductive assumption that $S_{n=10; a = \bar{n} - \frac{\epsilon}{4}}$ is a probabilistic polynomial for NEAR $_{a = \bar{n}}$ with error probability at most $\frac{\epsilon}{4}$. In this case, $M_{n; a}(\mathbf{x}) = A_{n; 2a}(\mathbf{x})$, which is 1 by definition of A .

$v \leq a = \bar{n} + \frac{\epsilon}{4}$, then with probability at least $1 - \frac{\epsilon}{4}$, we have $S_{n=10; a = \bar{n} + \frac{\epsilon}{4}}(\mathbf{x}) = 0$, in which case $M_{n; a}(\mathbf{x}) = M_{n=10; a = \bar{n} + \frac{\epsilon}{4}}(\mathbf{x})$. But, by the inductive hypothesis, this is 1 with probability at least $1 - \frac{\epsilon}{4}$, since $v > \frac{1}{4}$ in this case.

Since we are in one of these two cases with probability $1 - \frac{\epsilon}{4}$, and each gives the correct answer with probability $1 - \frac{\epsilon}{4}$, the correct answer is given in this case with probability $1 - \epsilon$.

2. $w \leq \frac{1}{4}$ ($|\mathbf{x}| = n$ is close to 0). In this case we have $v \leq \frac{1}{4}$, therefore $A_{n; 2a}(\mathbf{x}) = 1$. Hence, if $S_{n=10; a = \bar{n} - \frac{\epsilon}{4}}(\mathbf{x}) = 1$ then $M_{n; a}(\mathbf{x})$ returns the correct

answer. If $S_{n=10; ; a=^P \bar{n}; = 4}(x) = 0$, then we return $M_{n=10; ; = 4}(x)$. By Lemma 7.4, we have \dots with probability at least $1 - \frac{1}{4}$, and in this case, $M_{n=10; ; = 4}(x) = 1$ with probability $1 - \frac{1}{4}$. Hence, M returns the correct value with probability at least $1 - \frac{2}{4}$, no matter what the value of $S_{n=10; ; a=^P \bar{n}; = 4}(y)$ happens to be.

3. $w \geq 2$ ($|x| = n$ is far from \dots). By Lemma 7.4, we have \dots with probability at least $1 - \frac{1}{4}$. In this case, \dots , and so $M_{n=10; ; = 4}(x) = 1$ with probability $1 - \frac{1}{4}$. Moreover, since \dots , it follows that $S_{n=10; ; a=^P \bar{n}; = 4}(x) = 0$ with probability $1 - \frac{2}{4}$, in which case $M_{n; ; } (x) = M_{n=10; ; = 4}(x)$. Overall, $M_{n; ; } (x) = M_{n=10; ; = 4}(x) = 1$ with probability $1 - \dots$.

This completes the proof of correctness, and the proof of Theorem 7.2.

7.5 Further Probabilistic Polynomial Constructions

In this Section, we give two additional constructions of probabilistic polynomials which strengthen Theorem 7.2 from the previous section. Both of them make use of the following observation about the proof of correctness of Theorem 7.2: The only randomness used by the construction is the sampled vector v at each layer of the recursion, and moreover, the polynomial will always give a correct answer as long as $\sum_{j=1}^n |x_j| = n - j \leq a = \frac{1}{n}$ at each layer of recursion. This condition is true for the probabilistic polynomial for TH no matter what $v \in [0; 1]^n$ is.

7.5.1 Symmetric Functions

Recall that the Boolean function $f : \{0; 1\}^n \rightarrow \{0; 1\}$ is symmetric if the value of $f(x)$ depends only on $|x|$, the Hamming weight of x . We now describe how to use the probabilistic polynomial for TH to derive a probabilistic polynomial for any symmetric function with the same degree d .

Theorem 7.3. Every symmetric function $f : \{0; 1\}^n \rightarrow \{0; 1\}$ on n variables has a probabilistic polynomial of $O(\frac{n}{\log(1/\epsilon)})$ degree and error ϵ over \mathbb{Z} .

Proof. For every $0 \leq i \leq n$, let $f_i \in \{0; 1\}$ denote the value of $f(x)$ when x has Hamming weight i . Define:

$$A = \{0 \leq i \leq n \mid f_i = 1 \text{ and } f_{i-1} = 0\};$$

$$B = \{0 \leq i \leq n \mid f_i = 0 \text{ and } f_{i-1} = 1\};$$

Then, f can be written exactly as:

$$f(x) = f_0 + \sum_{i \in B} \text{TH}_{i=n}(x) - \sum_{j \in A} \text{TH}_{j=n}(x); \tag{7.2}$$

We replace each TH in (7.2) with a probabilistic polynomial of Theorem 7.2 with error ϵ . However, we make sure that in all of the different probabilistic polynomials for

TH, we make the same choice for the sampled vector at each layer of recursion. We can then apply the proof of Theorem 7.2, to see that every one of the probabilistic polynomials will give the correct answer as long as $\sum_{j=1}^n |x_j - \bar{x}| < a \sqrt{n}$ at each of the $\log_{10}(n)$ layers of recursion (this is a property only of the sampling, and independent of ϵ). Just as in the original proof, this will happen with error at most ϵ , as desired. \square

7.5.2 Derandomization

The polynomial for TH in Theorem 7.2 used $\log_{10}(n)$ random bits in order to randomly sample x from x at each layer of recursion. In this subsection, we show it can be implemented using only $\text{poly}(\log(n); \epsilon)$ random bits. The key will be to sample x with limited independence combined with a Chernoff bound for bits with limited independence (Lemma 2.2 from the Preliminaries). In particular, we use the following strengthening of Lemma 7.4 from the previous Section:

Lemma 7.5. If $x \in \{0, 1\}^n$ with $\sum_{j=1}^n x_j = w$, and $\epsilon \in (0, 1)$ is a vector each of whose entries is k -wise independently chosen entry x_j , where $k = \lceil 20 \epsilon^{-3} \log(1/\epsilon) \rceil$, with $\sum_{j=1}^n x_j = v$, then for every $\delta < 1/4$,

$$\Pr |v - w| \geq \frac{a}{n} \sqrt{\frac{\delta}{4}}$$

where $a = \frac{1}{10} \sqrt{\frac{\delta}{\ln(1/\epsilon)}}$.

Proof. Apply Lemma 2.2 with $X = \sum_{j=1}^k x_j$, $\mu = E[\sum_{j=1}^k x_j] = wn/k$, and $\epsilon = \frac{\delta}{40 \log(1/\epsilon) \sqrt{n}}$. \square

Theorem 7.4. For any $0 < \epsilon < 1$, there is a probabilistic polynomial for the threshold function TH of degree $O(\frac{1}{\epsilon} \log(1/\epsilon))$ on n bits with error ϵ that can be randomly sampled using $O(\log(n) \log(1/\epsilon))$ random bits.

Proof. We follow the construction of Theorem 7.2 exactly, with only one modification. In the original proof, x was a sample of $\log_{10}(n)$ bits of x , chosen independently at random. Here, we instead pick x to be a sample of $\log_{10}(n)$ bits chosen k -wise independently, for $k = \lceil 20 \epsilon^{-3} \log(1/\epsilon) \rceil$.

The only requirement of the randomness in the proof of Theorem 7.2 is that it satisfies Lemma 7.4, a concentration inequality for sampling x from x . Our new Lemma 7.5 is identical to Lemma 7.4, except that it replaces the old method of sampling x with new k -wise sampling; the remainder of the proof of correctness is exactly as before.

Recall that in our recursive polynomial construction, we divide n by 10 and divide by 4 each time we move from one recursive layer to the next. At the j th recursive level of our construction, for $1 \leq j < \log_{10}(n)$, we thus need to $O(\log(4^j))$ -wise independently sample $\log_{10}(n)$ entries from a vector of length $n \cdot 10^{-j}$. Summing across all of the layers, we need a total of $O(n)$ samples from a k -wise independent space,

where k is never more than $O(\log(n))$. This can be done all together using $O(n)$ samples from $\{1, 2, \dots, n\}$ which are $O(\log(n))$ -wise independent. Using standard constructions², this requires $O(\log(n) \log(n))$ random bits. \square

7.6 PTFs for ORs of Threshold Functions

The primary way we will apply our probabilistic polynomial for the threshold function TH in the next Chapter is to efficiently compute ORs of thresholds. Suppose any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ has a probabilistic polynomial P of degree d and error ϵ over any commutative ring R . Then, for any $s \geq 1$, we can construct a probabilistic polynomial for $\text{OR}_s f$, the OR of s independent copies of f , of degree sd and error at most $s\epsilon$: On input $x^{(1)}; x^{(2)}; \dots; x^{(s)} \in \{0, 1\}^n$, we compute $\sum_{i=1}^s P(x^{(i)})$ by drawing $ap \sim P$, letting $q : R^s \rightarrow R$ be the exact polynomial for OR_s of degrees, and outputting $q(p(x^{(1)}); p(x^{(2)}); \dots; p(x^{(s)}))$. Indeed, by a union bound, the polynomial $p(x^{(i)})$ for each $i \in [s]$ will give the correct answer with error at most ϵ , and so q will exactly compute the OR.

When R has characteristic 0 (say, for instance, $R = \mathbb{R}$), we can do even better: we can construct a 'probabilistic polynomial' for $\text{OR}_s f$ of degree d and error at most $s\epsilon$, by outputting $p(x^{(1)}) + p(x^{(2)}) + \dots + p(x^{(s)}) - 1$. With error at most $s\epsilon$, this polynomial will output 1 when the $\text{OR}_s f$ is false, and a nonnegative value when it is true. This is not a probabilistic polynomial for $\text{OR}_s f$ as we defined it earlier, since it only outputs a nonnegative value in the 'true' case, rather than necessarily outputting 1. However, this 'thresholding' behavior between true and false values still allows us to determine whether the $\text{OR}_s f$ was true or false.

Inspired by this remark, in this Section, we show how to construct low-degree polynomial threshold functions (PTFs) representing TH that have good threshold behavior, and consequently obtain low-degree PTFs for an OR of many threshold functions. By only aiming for such thresholding behavior, we will be able to further decrease the degree of our polynomial representations TH.

Definition 7.4. A polynomial threshold function (PTF) for the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a polynomial $p : \{0, 1\}^n \rightarrow \mathbb{R}$ such that, for every $x \in \{0, 1\}^n$, we have $p(x) \geq 0$ if $f(x) = 1$, and $p(x) < 0$ if $f(x) = 0$. The PTF degree of f is the minimum degree of a PTF for f .

Example 7.2. The function TH has PTF degree 1, via the PTF $p(x_1; \dots; x_n) = x_1 + \dots + x_n - n$. However, it is unclear how to use this to construct a PTF for $\text{OR}_s \text{TH}$; we will need a PTF with better thresholding behavior for TH for this.

²For example, one can pick a uniformly random degree k (single-variable) polynomial over F_q for some prime power $q \geq n$, and output its values on n distinct points from F_q .

7.6.1 Deterministic Construction

We begin by reviewing some basic facts about Chebyshev polynomials. The degree q Chebyshev polynomial of the first kind

$$T_q(x) := \sum_{i=0}^{\lfloor q/2 \rfloor} \binom{q}{2i} (x^2 - 1)^i x^{q-2i}.$$

Fact 7.1. For any $\epsilon \in (0, 1)$,

if $x \in [-1, 1]$, then $|T_q(x)| \leq 1$;

if $x \in (1, 1 + \epsilon)$, then $T_q(x) > 1$;

if $x = 1 + \epsilon$, then $T_q(x) = \frac{1}{2}(e^{q\epsilon} + e^{-q\epsilon})$.

Proof. The first property easily follows from the known formula $T_q(x) = \cos(q \arccos(x))$ for $x \in [-1, 1]$. The second and third properties follow from another known formula $T_q(x) = \cosh(q \operatorname{arccosh}(x))$ for $x > 1$, which for $x = 1 + \epsilon$ implies $T_q(x) = \cosh(q\epsilon) = \frac{1}{2}(e^{q\epsilon} + e^{-q\epsilon})$. \square

In certain scenarios, we obtain slightly better results using a (lesser known) family of discrete Chebyshev polynomials defined as follows [Hir03, page 59]:

$$D_{q;t}(x) := \sum_{i=0}^{\lfloor t/x \rfloor} \binom{q}{i} \left(\frac{x}{t}\right)^i x^{q-i}.$$

(See also [Sze75, pages 33–34] or Chebyshev's original paper [Che99] with an essentially equivalent definition up to rescaling.)

Fact 7.2. Let $c_{q;t} = (t+1)^{q+1} = q! \cdot \frac{1}{8(t+1)\ln(t+1)}$.

if $x \in [0, 1; \dots; t]$, then $|D_{q;t}(x)| \leq c_{q;t}$;

if $x = 1$, then $D_{q;t}(x) = e^{q^2/(8(t+1))} c_{q;t}$.

Proof. From [Hir03, page 61],

$$\begin{aligned} \sum_{k=0}^{\lfloor t/x \rfloor} D_{q;t}(k)^2 &= \sum_{k=0}^{\lfloor t/x \rfloor} \binom{2q}{k} \left(\frac{x}{t}\right)^k x^{2q-k} \\ &= \frac{2q!}{q!(q-k)!} \frac{q!}{k!} \frac{(t+1+k)(t+q)}{(2q+1)(2q-k)} \frac{(t+1-k)}{1} \\ &= \frac{(t+1)((t+1)^2 - k^2)((t+1)^2 - k^2)}{(2q+1)(q!)^2} \\ &= \frac{(t+1)^{2q+2}}{(q!)^2}. \end{aligned}$$

Thus, for every integer $x \in [0; t]$, we have $|D_{q;t}(x)| \leq (t+1)^{q+1} = q! = c_{q;t}$.

For $x \in [0, 1]$, we have $\binom{t}{i} x^i = \frac{(x)(x+1)\dots(x+i-1)}{i!} \binom{t}{i}$, and by the Chu-Vandermonde identity,

$$\begin{aligned}
 D_{q;t}(x) &= \sum_{i=0}^q \binom{t+1}{i} x^i = \sum_{i=0}^q \binom{t+1+q}{i} \frac{(x)(x+1)\dots(x+i-1)}{i!} \\
 &= \frac{(t+1)^q (1 + \frac{1}{t+1}) (1 + \frac{2}{t+1}) \dots (1 + \frac{q}{t+1})}{q!} \\
 &= \frac{C_{q;t}}{t+1} e^{\frac{1+2+\dots+q}{2(t+1)}} = e^{q(q+1)/(4(t+1))} \ln(t+1) C_{q;t} e^{q^2/(8(t+1))} C_{q;t} \quad \square
 \end{aligned}$$

Using the Chebyshev polynomials and the Discrete Chebyshev polynomials, we can design PTFs for TH with 'nice' thresholding behavior. Our construction also extends to 'approximate threshold functions', where there is a middle range of input values between the 'true' and 'false' inputs where we are allowed to output any value. Our construction involves two parameters: s , which corresponds to the number of different copies of the threshold function whose OR we want to take, and ϵ , the 'gap' in the approximate threshold.

Theorem 7.5. We can construct a polynomial $P_{s;t,\epsilon} : \mathbb{R} \rightarrow \mathbb{R}$ of degree $O(\frac{p}{1-\epsilon} \log s)$, such that

- if $x \in [0, 1 - \epsilon]$, then $|P_{s;t,\epsilon}(x) - 0| \leq \epsilon$;
- if $x \in [(1 + \epsilon)t, t]$, then $P_{s;t,\epsilon}(x) > 1 - \epsilon$;
- if $x \in [(1 + \epsilon)t, t]$, then $P_{s;t,\epsilon}(x) \leq s$.

For the exact setting with $\epsilon = 1/t$, we can alternatively bound the degree by $O(\frac{p}{t} \log(st))$.

Proof. Set $P_{s;t,\epsilon}(x) := T_q(x/t)$, where T_q is the Chebyshev polynomial, for a parameter q to be determined. The first two properties are obvious from Fact 7.1. On the other hand, if $x \in [(1 + \epsilon)t, t]$, then Fact 7.1 shows that $P_{s;t,\epsilon}(x) \leq \frac{1}{2} e^{q \frac{p}{m}}$ s , provided we set $q = \frac{p}{1-\epsilon} \ln(2s)$. This achieves $O(\frac{p}{1-\epsilon} \log s)$ degree.

When $\epsilon = 1/t$ the above yields $O(\frac{p}{t} \log s)$ degree; we can reduce the $\log s$ factor by instead defining $P_{s;t,\epsilon}(x) := \frac{D_{q;t}(t-x)}{C_{q;t}}$. Now, if $x \in [t+1, t]$, then $P_{s;t,\epsilon}(x) \leq e^{q^2/(8(t+1))} s$ by setting $q = \frac{p}{8(t+1) \ln(\max\{s; t+1\})}$. \square

Using Theorem 7.5, we can construct a low-degree PTF for computing an OR of n thresholds of n bits:

Corollary 7.3. Given $p; s; t; \epsilon$, we can construct a polynomial $P : [0, 1]^{ns} \rightarrow \mathbb{R}$ of degree at most $d := O(\frac{p}{1-\epsilon} \log s)$ and at most n^d monomials, such that

$$\text{if } \sum_{i=1}^s \prod_{j=1}^n x_{ij} > t \text{ is false, then } |P(x_{11}, \dots, x_{1n}, \dots, x_{s1}, \dots, x_{sn})| \leq s;$$

if $\sum_{i=1}^s \sum_{j=1}^n x_{ij} \geq t + \epsilon n$ is true, then $P(x_{11}, \dots, x_{1n}, \dots, x_{s1}, \dots, x_{sn}) > 2s$.

For the exact setting with $\epsilon = 1/n$, we can alternatively bound by $O(\sqrt{n \log(ns)})$.

Proof. Define $P(x_{11}, \dots, x_{1n}, \dots, x_{s1}, \dots, x_{sn}) := \prod_{i=1}^s P_{n;3s;t,\epsilon} \prod_{j=1}^n x_{ij}$; where $P_{n;3s;t,\epsilon}$ is from Theorem 7.5. It is not hard to see that the stated properties hold. (In the second case, the output is at least $2s - \epsilon n > 2s$.) \square

7.6.2 Probabilistic Construction

Finally, we give our last polynomial construction, by combining the PTF from Theorem 7.5 from the previous Subsection with our probabilistic polynomial from Theorem 7.2 above. We will construct a distribution of PTFs to randomly draw from, which will allow us to achieve noticeably lower degree than either the PTFs from the previous section, or the probabilistic polynomial from before.

Definition 7.5. For any $\epsilon \in [0, 1]$, a probabilistic PTF with error ϵ and degree d for the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a distribution P on polynomials $p : \{0, 1\}^n \rightarrow \mathbb{R}$ of degree at most d over \mathbb{R} such that, for every $x \in \{0, 1\}^n$, we have

If $f(x) = 1$, then $\Pr_{p \sim P} [p(x) \geq 0] \geq 1 - \epsilon$, and

If $f(x) = 0$, then $\Pr_{p \sim P} [p(x) < 0] \geq 1 - \epsilon$.

The ϵ -probabilistic PTF degree of f is the minimum degree of a probabilistic PTF with error ϵ for f .

Our construction will use a 'random sampling' approach similar to Theorem 7.2 from before.

Restatement of Theorem 7.2 We can construct a distribution $Q_{n;s;t}$ on polynomials $Q_{n;s;t} : \{0, 1\}^n \rightarrow \mathbb{R}$ of degree $O(\sqrt{n \log s})$, such that for every $x \in \{0, 1\}^n$, when we draw a random $Q_{n;s;t}$ $Q_{n;s;t}$:

if $\sum_{i=1}^n x_i \leq t$, then $Q_{n;s;t}(x_1, \dots, x_n) = 0$ with probability at least $1 - \epsilon$;

if $\sum_{i=1}^n x_i > t$, then $Q_{n;s;t}(x_1, \dots, x_n) = 1$ with probability at least $1 - \epsilon$.

Theorem 7.6. We can construct a distribution $L_{n;s;t,\epsilon}$ on polynomials $L_{n;s;t,\epsilon} : \{0, 1\}^n \rightarrow \mathbb{R}$ of degree $O((1/\epsilon)^3 \log s)$, such that for every $x \in \{0, 1\}^n$, when we draw a random $L_{n;s;t,\epsilon}$ $L_{n;s;t,\epsilon}$:

if $\sum_{i=1}^n x_i \leq t$, then $|L_{n;s;t,\epsilon}(x_1, \dots, x_n)| \leq \epsilon$ with probability at least $1 - \epsilon$;

if $\sum_{i=1}^n x_i \geq 2(t + \epsilon n)$, then $L_{n;s;t,\epsilon}(x_1, \dots, x_n) > 1$ with probability at least $1 - \epsilon$;

if $\sum_{i=1}^n x_i \leq t + \epsilon n$, then $L_{n;s;t,\epsilon}(x_1, \dots, x_n) \leq \epsilon$ with probability at least $1 - \epsilon$.

For the exact setting with $n = 1 = n$, we can alternatively bound the degree by $O(n^{1=3} \log^{2=3}(ns))$.

Proof. Let r and q be parameters to be set later. Draw a random sample x_1, \dots, x_n of size n . Let

$$t_R := \frac{tr}{n} \quad c_0^p \overline{r \log s} \quad \text{and} \quad t := t - 2c_0 \frac{n}{r} \quad p \overline{\log s}$$

for a sufficiently large constant c_0 . Define

$$L_{n;s;t} (x_1, \dots, x_n) := Q_{r;2s;t_R} (f_{x_i} g_{i2R}) \quad P_{s;t^0;n^0} \prod_{i=1}^n x_i \quad t ;$$

where $P_{s;t^0;n^0}$ is the polynomial from Theorem 7.5, with $t^0 := t - t = ((n = r) \overline{p \log s})$ and $n^0 := n - t^0 = ((n = r) \overline{p \log s})$, and $Q_{r;2s;t_R}$ is a polynomial drawn from $Q_{r;2s;t_R}$ from Theorem 7.2.

To verify the stated properties, consider three cases:

Case 1: $\prod_{i=1}^n x_i < t$. By a Chernoff bound, with probability at least $1 - \epsilon$, we have $\prod_{i=1}^n x_i < t - n + c_0 \overline{r \log s} - t_R$ (assuming that $r \log s$). Thus, with probability at least $1 - \epsilon$, we have $Q_{r;2s;t_R} (f_{x_i} g_{i2R}) = 0$ and so $L_{n;s;t} (x_1, \dots, x_n) = 0$.

Case 2: $\prod_{i=1}^n x_i \in [t, t]$. With probability at least $1 - \epsilon$, we have $Q_{r;2s;t_R} (f_{x_i} g_{i2R}) \in [0, 1]$ and so $L_{n;s;t} (x_1, \dots, x_n) \in [0, 1]$.

Case 3: $\prod_{i=1}^n x_i > t$. By a Chernoff bound, with probability at least $1 - \epsilon$, we have $\prod_{i=1}^n x_i > t + n + c_0 \overline{r \log s} = t_R$. Thus, with probability at least $1 - \epsilon$, we have $Q_{r;2s;t_R} (f_{x_i} g_{i2R}) = 1$ and so $L_{n;s;t} (x_1, \dots, x_n) > 1$ for $\prod_{i=1}^n x_i \in [t, t + n]$, or $L_{n;s;t} (x_1, \dots, x_n) \leq \epsilon$ for $\prod_{i=1}^n x_i > t + n$.

The degree of $L_{n;s;t}$ is

$$O \left(\frac{p}{r \log s} + \frac{q}{(1 - \frac{p}{r})^p \log s \log s} \right)$$

and we can set $r = (1 - \epsilon)^{2=3} \log s$. For the exact setting, the degree is

$$O \left(\frac{p}{r \log s} + \frac{q}{(1 - \frac{p}{r})^p \log s \log(ns)} \right)$$

and we can set $r = \frac{1}{n^{2=3} \log^{1=3}(ns)}$. □

Remark 7.1. Using the same techniques as in Theorem 7.4, we can sample a probabilistic polynomial from Theorem 7.6 with only $O(\log(n) \log(ns))$ random bits.

Finally, we can construct a probabilistic PTF for an OR of thresholds:

Corollary 7.4. Given n, s, t, ϵ , we can construct a distribution L on polynomials $L : \{0, 1\}^{ns} \rightarrow \mathbb{R}$ of degree at most $d := O((1/\epsilon)^{1/3} \log s)$ with at most s^n monomials, such that

if $\prod_{i=1}^s \prod_{j=1}^n x_{ij} = t$ is false, then $L(x_{11}, \dots, x_{1n}, \dots, x_{s1}, \dots, x_{sn}) \leq \epsilon$ with probability at least $2/3$;

if $\prod_{i=1}^s \prod_{j=1}^n x_{ij} = t + \epsilon^n$ is true, then $L(x_{11}, \dots, x_{1n}, \dots, x_{s1}, \dots, x_{sn}) > 2\epsilon$ with probability at least $2/3$.

For the exact setting with $\epsilon = 1/n$, we can alternatively bound d by $O(n^{1/3} \log^{2/3}(ns))$.

Proof. Draw $L_{n, 3s, t, \epsilon}$ from the distribution from Theorem 7.6, then define $L(x_{11}, \dots, x_{1n}, \dots, x_{s1}, \dots, x_{sn}) := \prod_{i=1}^s L_{n, 3s, t, \epsilon}(x_{i1}, \dots, x_{in})$. \square

Remark 7.2. The coefficients of the polynomials from Theorem 7.2 are $\text{poly}(n)$ -bit integers, and it can be checked that the coefficients of all our deterministic and probabilistic PTFs are rational numbers with $\text{poly}(n)$ -bit numerators and a common $\text{poly}(n)$ -bit denominator, and that the same bound for the number of monomials holds for the construction time, up to $\text{poly}(n)$ factors. That is, computations with these polynomials have low computational overhead relative to

Chapter 8

Algorithmic Applications

8.1 Sparse Polynomials and Rectangular Matrix Multiplication

In this Chapter, we give new algorithmic applications of the polynomials we constructed in the previous chapter. Because of the prevalence of threshold functions throughout algorithms and complexity, we will be able to apply them in a variety of settings: nearest neighbor search, high-dimensional computational geometry, the Light Bulb problem from data science, MAX-SAT, circuit SAT, and threshold circuit lower bounds.

A key insight in many of the applications is a way to quickly evaluate a sparse polynomial (i.e. a polynomial with few monomials) on a combinatorial rectangle of inputs by using rectangular matrix multiplication. This makes use of the following simple reduction from evaluating a polynomial to computing an inner product:

Proposition 8.1. For $d \geq 1$, and any commutative ring R , let $p : R^d \rightarrow R$ be any polynomial with t monomials. Then, there are maps $\alpha : R^d \rightarrow R^t$ such that, for any $x, y \in R^d$, we have $p(x; y) = \sum_{i=1}^t \alpha_i(x) y_i$. Moreover, if p has degree d , then α can be computed in $O(t^2 d)$ arithmetic operations over R .

Proof. Since p has t monomials, there are maps $\alpha : [d] \rightarrow [t]$, $\beta : [d] \rightarrow [t]$, $\gamma : [t] \rightarrow R$ such that

$$p(x; y) = \sum_{\ell=1}^t \alpha_\ell(x) \beta_\ell(y) \gamma_\ell$$

We define $\alpha : R^d \rightarrow R^t$, for $\ell \in [t]$, by:

$$\alpha_\ell(x) := \sum_{i=1}^d \alpha_{\ell,i} x_i^{\beta_{\ell,i}}$$

Similarly define $h : R^d \rightarrow R^t$, for $t \geq [t]$, by:

$$h(y) := \sum_{j=1}^d y_j^{b(j;t)}$$

We can see that $p(x; y) = h(x) \cdot h(y)$ as desired. \square

By combining Proposition 8.1 with rectangular matrix multiplication algorithms, we can quickly evaluate sparse polynomials on many pairs of inputs; such a technique was first used in [Wil14a], and also implicitly in [Wil14c].

Proposition 8.2. For $d \geq N$, and any commutative ring R , let $p : R^{2d} \rightarrow R$ be any polynomial of degree $\leq t$ with t monomials. Given as input two sets $A, B \subseteq R^d$, using $O((|A| + |B|) \cdot t)$ arithmetic operations over R , we can reduce the problem of evaluating p on all pairs $(x; y) \in A \times B$ to the problem of $|A| \cdot |B|$ matrix multiplication over R .

Proof. Letting $h : R^d \rightarrow R^t$ be the maps from Proposition 8.1, we compute $h(x)$ for each $x \in A$, and $h(y)$ for each $y \in B$, then multiply the resulting two matrices. \square

We will typically apply Proposition 8.2 in conjunction with Coppersmith's very efficient rectangular matrix multiplication algorithm:

Lemma 8.1 ([Cop82]). For all sufficiently large N , multiplication of an $N \times N$ matrix with an $N \times N$ matrix can be done in $O(N^2 \log^2 N)$ arithmetic operations over any field.

Such a rectangular matrix multiplication algorithm requires some care; simply applying the idea from Proposition 4.1 above to Coppersmith's rank expression would yield an algorithm which uses $O(N^{2+\epsilon})$ arithmetic operations for any $\epsilon > 0$. A proof of how one can perform only $O(N^2 \log^2 N)$ arithmetic operations can be found in the appendix of [Wil14b]. We will typically use this approach when R is either a finite field, or else R but with relatively small coefficients, so that the arithmetic operations can be performed in $\text{polylog}(n)$ time:

Lemma 8.2 ([Wil14a]). Given a polynomial $P(x_1, \dots, x_d; y_1, \dots, y_d)$ with at most $n^{0.17}$ monomials such that either:

P is over a finite field F_q with $q = \text{polylog}(n)$, or

P is over R , and all its coefficients are $\text{polylog}(n)$ -bit numbers,

then given two sets of inputs $A = \{a_1, \dots, a_n\} \subseteq \{0, 1\}^d$, $B = \{b_1, \dots, b_n\} \subseteq \{0, 1\}^d$, we can evaluate P on all pairs $(a_i; b_j) \in A \times B$ in $O(n^2 + d \cdot n^{1.17})$ time.

8.2 Exact Batch Nearest Neighbor Search

We now apply our polynomial constructions to solve the exact batch nearest neighbor problem. We begin by focusing on the related *closest pair* problem.

Let M be a metric on $\{0, 1\}^d$. We define the *Bichromatic M -Metric Closest Pair* problem to be: given an integer k and a collection of red and blue vectors in $\{0, 1\}^d$, determine if there is a pair of red and blue vectors with distance at most k under metric M . This problem arises frequently in algorithms on a metric space. In what follows, we shall assume that the metric M can be computed on two points of d dimensions in time $\text{poly}(d)$. Define the Boolean function

$$M\text{-dist}_k(x_{1;1}, \dots, x_{1;d}, \dots, x_{s;1}, \dots, x_{s;d}, y_{1;1}, \dots, y_{1;d}, \dots, y_{s;1}, \dots, y_{s;d}) \\ := \bigvee_{i,j=1,\dots,s} [M(x_{i;1}, \dots, x_{i;d}, y_{j;1}, \dots, y_{j;d}) \leq k]$$

That is, $M\text{-dist}_k$ takes two collections of s vectors as input, and outputs 1 if and only if there is a pair of vectors (one from each collection) that have distance at most k under metric M . For example, the *Hamming-dist $_k$* function tests if there is a pair of vectors with Hamming distance at most k .

We observe that sparse probabilistic polynomials for computing $M\text{-dist}_k$ imply subquadratic time algorithms for finding close bichromatic pairs in metric M .

Theorem 8.1. Suppose for all k, d , and n , there is an $s = s(d; n)$ such that $M\text{-dist}_k$ on $2sd$ variables has a probabilistic PTF with at most $n^{0.17}$ monomials, $\text{polylog}(n)$ -bit coefficients, and error at most $1/3$, where each sample can be produced in $\mathcal{O}(n^2 = s^2)$ time. Then *Bichromatic M -Metric Closest Pair* on n vectors in d dimensions can be solved in $\mathcal{O}(n^2 = s^2 + s^2 \text{poly}(d))$ randomized time.

Proof. We have an integer k and sets $R, B \subseteq \{0, 1\}^d$ such that $|R| = |B| = n$, and wish to determine if there is $u \in R$ and $v \in B$ such that $M(u; v) \leq k$. First, partition both R and B into $dn = se$ groups, with at most s vectors in each group. By assumption, for all k , there is a probabilistic PTF for $M\text{-dist}_k$ with $2sd$ variables, $n^{0.17}$ monomials, $\text{polylog}(n)$ -bit coefficients, and error at most $1/3$. Let p be a polynomial sampled from this distribution. Our idea is to efficiently evaluate p on all $\mathcal{O}(n^2 = s^2)$ pairs of groups from R and B , by feeding as input to p all s vectors x_i from a group of R and all s vectors y_i from a group of B .

Since the number of monomials is $n^{0.17}$, we can apply Lemma 8.2, evaluating p on all pairs of groups in time $\mathcal{O}(n^2 = s^2)$. For each pair of groups from R and B , this evaluation determines if the pair of groups contain a bichromatic pair of distance at most k , with probability at least $2/3$.

To obtain a high probability answer, sample $\ell = 10 \log n$ polynomials p_1, \dots, p_ℓ for $M\text{-dist}_k$ independently from the distribution, in $\mathcal{O}(n^2 = s^2)$ time (by assumption). Evaluate each p_i on all pairs of groups from R and B in $\mathcal{O}(n^2 = s^2)$ time by the above paragraph. Compute the majority value of p_1, \dots, p_ℓ on all pairs of groups, again in $\mathcal{O}(n^2 = s^2)$ time. By a Chernoff-Hoeffding bound, the majority value reported for a pair of groups is correct with probability at least $1 - n^{-3}$. Therefore with probability

at least $1 - \epsilon$, we correctly determine for all pairs of groups from R and B whether the pair contains a bichromatic pair of vectors with distance at most k .

Given a pair of groups R^0 and B^0 which are reported to contain a bichromatic pair of close vectors, we can simply brute force to find the closest pair in R^0 and B^0 in $s^2 \text{poly}(d)$ time. (In principle, we could also perform a recursive call, but this doesn't asymptotically help us in our applications.) \square

We will use our probabilistic PTF from the previous chapter:

Lemma 8.3. For sufficiently large s and d , the Hamming-dist $_k$ function on $2sd$ variables has a probabilistic PTF of degree $O(d^{1+3} \log^{2+3}(ds))$, error at most ϵ , and at most $O(s^{2+O(d^{1+3} \log^{2+3}(ds))})$ monomials, whose coefficients are $\text{polylog}(n)$ bit numbers. Moreover, we can sample from the probabilistic PTF distribution in time polynomial in the number of monomials.

Proof. Applying Corollary 7.4 ('exact setting') with $n = d$, $s = s^2$, and $t = k$, gives us a probabilistic PTF degree $O(d^{1+3} \log^{2+3}(ds))$, error at most ϵ , and at most $O(s^{2+O(d^{1+3} \log^{2+3}(ds))})$ monomials for, on input $z \in \{0, 1\}^{2ds^2}$, testing the predicate $\sum_{i,j \in [s]} \sum_{\ell=1}^d z_{i,j,\ell} \geq k$. We use this to solve Hamming-dist $_k$ function on inputs $x, y \in \{0, 1\}^{2ds}$ by substituting in $z_{i,j,\ell} = (x_i - y_j)^\ell$ for all $i, j \in [s]$ and $\ell \in [d]$. Hence, for any $i, j \in [s]$, the quantity $\sum_{\ell=1}^d z_{i,j,\ell}$ computes exactly the Hamming distance between x_i and y_j , and so the result is a probabilistic PTF for the Hamming-dist $_k$ function.

Since we substituted in a polynomial with 4 monomials for each $z_{i,j,\ell}$, this increases the number of monomials in the resulting probabilistic PTF by a factor of $4^{O(d^{1+3} \log^{2+3}(ds))}$, but this factor is subsumed by the binomial coefficient $O(d^{1+3} \log^{2+3}(ds))$. \square

Putting it all together, we obtain a faster algorithm for Bichromatic Hamming Closest Pair:

Theorem 8.2. For n vectors of dimension $d = c(n) \log n$, Bichromatic Hamming Closest Pair can be solved in $m^{2+1/O(c(n) \log^3 c(n))}$ time by a randomized algorithm that is correct with high probability.

Proof. Let $d = c \log n$ in the following, with the implicit understanding that c is a function of n . We apply the reduction of Theorem 8.1 and the probabilistic PTF for Hamming-dist $_k$ of Lemma 8.3.

The reduction of Theorem 8.1 requires that the number of monomials in our probabilistic polynomial is at most $n^{0.17}$, while the monomial bound for Hamming-dist from Theorem 8.3 is $m = O(s^{2+O(d^{1+3} \log^{2+3}(s))})$ for some universal constant a , provided that $s > d$ are sufficiently large. Therefore our primary task is to maximize the value of s such that $m = n^{0.17}$. This will minimize the total running time of $O(n^2 = s^2)$. With hindsight, let us guess $s = n^{1-(u^P c \log^3 c)}$ for a constant u , and focus on the

large binomial in the monomial estimator. Then,

$$\begin{aligned} \frac{2d}{a^{1-3} \log^{2-3}(s)} &= \frac{2c \log n}{a(c \log n)^{1-3} (\log n)^{2-3} = (u^p \bar{c} \log^{3-2} c)^{2-3}} \\ &= \frac{2c \log n}{a \log n = (u^{2-3} \log c)} : \end{aligned}$$

For notational convenience, let $\beta = a(u^{2-3} \log c)$. By Proposition 2.2 from the Preliminaries, we have

$$\frac{2c \log n}{\log n} \frac{2ce^{\log n}}{\log n} = n^{\log(\frac{2ce}{\beta})}.$$

Plugging $\beta = a(u^{2-3} \log c)$ back into the exponent, we find

$$\log \frac{2ce}{\beta} = \frac{a \log(\frac{2ce u^{2-3} \log c}{a})}{u^{2-3} \log c} : \tag{8.1}$$

The quantity (8.1) can be made arbitrarily small, by setting u sufficiently large. In that case, the number of monomials $\leq s^2 n^{\log(\frac{2ce}{\beta})}$ can be made less than $n^{0.1}$. \square

Remark 8.1. Observe that we would not have been able to prove Theorem 8.2 if the probabilistic PTF degree we applied from Lemma 8.3 had an additional $\text{poly}(\log(n))$ multiplicative factor (which would have been the case if we hadn't succeeded at removing such factors from the degree in Theorem 7.2 earlier). Indeed, propagating this extra factor through the proof of Theorem 8.2 would have resulted in an additional $\text{poly}(\log n)$ multiplicative factor in expression (8.1). It would then have been impossible to pick a constant u such that expression (8.1) is at most 0.1 .

Now we show how to solve Batch Hamming Nearest Neighbor (BHNN). In the following theorem, we assume for all pairs of vectors in our instance that the maximum metric distance is at most some value MAX . (For the Hamming distance, $\text{MAX} \leq d$.) We reduce the batch nearest neighbor query problem to the bichromatic close pair problem:

Theorem 8.3. Let E^d be some d -dimensional domain supporting a metric space M . If the Bichromatic M -Metric Closest Pair on n vectors in E^d can be solved in $T(n; d)$ time, then Batch M -Metric Nearest Neighbors on n vectors in E^d can be solved in $O(n \cdot T(\frac{n}{\text{MAX}}; d) \cdot \text{MAX})$ time.

Proof. We give an oracle reduction similar to previous work [AWY15]. Initialize a table T of size n , with the maximum metric value v in each entry. Given n database vectors D and n query vectors Q , color D red and Q blue. Break D into d groups of size at most MAX , and do the same for the set Q . For each pair $(R^0, B^0) \subseteq (D, Q)$ of groups, and for each $k = \text{MAX} - 1, \dots, 1, 0$, we initialize $D_k := D$, $Q_k := Q$, and call Bichromatic M -Metric Closest Pair on $(R^0, B^0) \subseteq (D_k, Q_k)$ with integer k . While we continue to find a pair $(x_i, y_j) \in (R^0, B^0)$ with $M(x_i, y_j) \leq k$, set $T[i] := k$

and remove y_j from Q_k and B^0 . (With a few more recursive calls, we could also find an explicit vector y_j such that $M(x_i; y_j) = k$.)

Now for each call that finds a close bichromatic pair, we remove a vector from Q_k ; we do this at most MAX times for each vector, so there can be at most $MAX \cdot n$ such calls. For each pair of groups, there are MAX oracle calls that find no bichromatic pair. Therefore the total running time is $O((n + n^2 = s^2) \cdot T(s; d) \cdot MAX)$. Setting $s = \sqrt{n}$ to balance the terms, the running time is $O(n \cdot T(\sqrt{n}; d) \cdot MAX)$. \square

The following is immediate from Theorem 8.3 and Theorem 8.2:

Theorem 8.4. For n vectors of dimension $d = c(n) \log n$, Batch Hamming Nearest Neighbors can be solved in $n^{2 - 1/O(c(n) \log^{3=2} c(n))}$ time by a randomized algorithm, with high probability.

Remark 8.2. For a deterministic algorithm, we can use the PTF from Corollary 7.3 instead of the probabilistic PTF from Corollary 7.4 when constructing our polynomial for the Hamming-distance function in Lemma 8.3. The exact same algorithm then results in a deterministic running time of $n^{2 - 1/O(c \log^2 c)}$.

8.2.1 Closest Pair in Hamming Space is Hard

The Strong Exponential Time Hypothesis (SETH) states that there is no universal $\epsilon < 1$ such that for all c , CNF-SAT with n variables and cn clauses can be solved in $O(2^{(1-\epsilon)n})$ time. We next show that, assuming SETH, there is a limit to how much one can improve our running time from Theorem 8.4.

Theorem 8.5. Suppose there is $\epsilon > 0$ such that for all constant c , Bichromatic Hamming Closest Pair can be solved in $n^{2 - \epsilon}$ time on a set of n points in $\{0, 1\}^{c \log n}$. Then SETH is false.

Proof. The proof is a reduction from the Orthogonal Vectors problem with n vectors $S \subseteq \{0, 1\}^d$, which asks whether there are $u, v \in S$ such that $u \cdot v = 0$. It is well-known that an algorithm for Orthogonal Vectors running in time $2^{o(d)} \cdot n^{2 - \epsilon}$ would refute SETH [Wil05]. Indeed, we show that Bichromatic Minimum Inner Product (finding a pair of vectors with minimum inner product, not just inner product zero) reduces to Bichromatic Hamming Closest Pair, as well as the version for maximum inner product.

First, we observe that Bichromatic Hamming Closest Pair is equivalent to Bichromatic Hamming Furthest Pair: let \bar{v} be the complement of v (the vector obtained by flipping all the bits of v). Then the Hamming distance of u and v is $H(u; v) = d - H(u; \bar{v})$. Thus by flipping all the bits in the components of the blue vectors, we can reduce from the closest pair problem to furthest pair, and vice versa.

Now we reduce Orthogonal Vectors to Bichromatic Hamming Furthest Pair. Our Orthogonal Vectors instance has red vectors S_r and blue vectors S_b , and we wish to find $u \in S_r$ and $v \in S_b$ such that $u \cdot v = 0$.

For every d^2 possible choice of $f; J = 1; \dots; d$, construct the subset $S_{r,f}$ of vectors in S_r with exactly f ones, and construct the subset $S_{b,J}$ of vectors in S_b with exactly

J ones. We will look for an orthogonal pair among $S_{r,l}$ and $S_{b,j}$ for all such l, j separately.

Recall that Hamming distance of two vectors equals the ℓ_2^2 norm distance, in $\{0, 1\}^d$. The ℓ_2^2 norm of u and v is

$$\|u - v\|_2^2 = \|u\|_2^2 + \|v\|_2^2 - 2\langle u, v \rangle$$

However, in $S_{r,l}$ all vectors have the same norm, and all vectors in $S_{b,j}$ have the same norm. Therefore, finding a red-blue pair $u \in S_{r,l}$ and $v \in S_{b,j}$ with minimum inner product is equivalent to finding a pair in $S_r \times S_b$ with smallest Hamming distance. (Similarly, maximum inner product is equivalent to Hamming closest pair.)

The reduction only requires $O(d^2)$ calls to Bichromatic Hamming Furthest Pair, with no changes to the dimension nor the number of vectors. \square

8.2.2 Metrics Beyond Hamming Distance

We now show how Theorem 8.4 can be extended to find nearest neighbors for a number of other metrics. We state our best randomized running times in this Subsection, but one could also apply Remark 8.2 to get deterministic algorithms with slightly worse running times in all of the applications below.

Recall that the ℓ_1 norm of two vectors x and y is $\sum_i |x_i - y_i|$. We can solve Batch ℓ_1 Nearest Neighbors on vectors with small integer entries by a simple reduction to Batch Hamming Nearest Neighbors, (which is probably folklore):

Theorem 8.6. For n vectors of dimension $d = c \log n$ in $\{0, 1, \dots, m\}^d$, Batch ℓ_1 Nearest Neighbors can be solved in $n^2 \cdot 1 = O(n^{c \log^3 = 2}(mc))$ time by a randomized algorithm, with high probability.

Proof. Notice that for any $x, y \in \{0, \dots, m\}^d$, the Hamming distance of their unary representations, written as m -dimensional vectors, is equal to $\|x - y\|_1$. Hence, for $x \in \{0, \dots, m\}^d$, we can transform it into a vector $x^0 \in \{0, 1\}^{md}$ by setting $(x_{m(i-1)+1}^0, x_{m(i-1)+2}^0, \dots, x_{m(i-1)+m}^0)$ equal to the unary representation of x_i , for $1 \leq i \leq d$. It is then equivalent to solve the Hamming nearest neighbors problem on these md -dimensional vectors. \square

It is also easy to extend Theorem 8.4 for vectors over $O(1)$ -sized alphabets using equidistant binary codes ([MKZ09], Section 5.1). This is useful for applications in biology, such as finding similar DNA sequences. The above algorithms also apply to computing maximum inner products:

Theorem 8.7. The Bichromatic Minimum Inner Product (and Maximum) problem with n red and blue Boolean vectors in $c \log n$ dimensions can be solved in $n^2 \cdot 1 = O(n^{c \log^3 = 2}(c))$ randomized time.

Proof. In Theorem 8.5 above, we gave a reduction from Bichromatic Minimum Inner Product to Bichromatic Hamming Furthest Pair, and showed that Bichromatic Hamming Furthest Pair is equivalent to Bichromatic Hamming

Closest Pair . The same reduction shows that Bichromatic Maximum Inner Product reduces to the closest pair version. Hence Theorem 8.4 applies, to both minimum and maximum inner products. \square

As a consequence, we can answer a batch of minimum inner product queries on a database of size n with the same time estimate, applying a reduction analogous to that of Theorem 8.3. From there, Theorem 8.7 can be extended to other important similarity measures, such as finding a pair of sets A, B with maximum Jaccard coefficient, defined as $\frac{|A \setminus B|}{|A \cup B|}$ [Bro97].

Corollary 8.1. Given n red and blue subsets of a universe of size n , we can find the pair of red and blue sets with maximum Jaccard coefficient in $O(n^2 \log^3 n)$ randomized time.

Proof. Let S be a given collection of red and blue sets over $[n]$. We construe the sets in S as vectors, in the natural way. For all possible values $d_1, d_2 = 1, \dots, d$, we will construct an instance of Bichromatic Maximum Inner Product S_{d_1, d_2}^0 , and take the best pair found, appealing to Theorem 8.7.

As in the proof of Theorem 8.5, we filter sets based on their cardinalities. In the instance S_{d_1, d_2}^0 of Bichromatic Maximum Inner Product, we only include red sets with cardinality exactly d_1 , and blue sets with cardinality exactly d_2 . For sets R, B , we have

$$\frac{|R \setminus B|}{|R \cup B|} = \frac{|R \setminus B|}{d_1 + d_2 + |R \setminus B|} \quad (8.2)$$

Suppose that we choose a red set R and blue set B that maximize $|R \setminus B|$. This choice simultaneously maximizes the numerator and minimizes the denominator of (8.2), producing the sets R and B with maximum Jaccard coefficient over the red sets with cardinality d_1 and blue sets with cardinality d_2 . Finding the maximum pair R and B over each choice of d_1, d_2 , we will find the overall R and B with maximum Jaccard coefficient. \square

8.3 Approximate Batch Nearest Neighbor Search

The same approach that we used in Theorem 8.4 for exact nearest neighbor search in Hamming space can be applied to solve approximate nearest neighbor search in Hamming space as well:

Theorem 8.8. Given n red and n blue points in $\{0, 1\}^d$ and $\epsilon = \log^6(d \log n) = \log^3 n$, we can find an approximate Hamming nearest/farthest blue neighbor with additive error at most ϵd for each red point in randomized time $n^{2 + \epsilon^{-1} \log(\frac{d}{\epsilon \log n})}$.

Proof. We mimic the proof of Theorem 8.4 up to the definition of the polynomial in Lemma 8.3. However, instead of applying the exact polynomial of Corollary 7.4, we insert the approximate polynomial construction from the same Corollary. While

the exact polynomial had degree $O(d^{1+3} \log^{2+3}(ds))$, the approximate one has degree $O((1+3)^{1+3} \log s)$. Setting

$$s := n^{1+3} := n^{(1+3) \log(\frac{d}{\log n})};$$

the number of monomials in the new polynomial is now

$$s^2 = \frac{O(d)}{O((1+3)^{1+3} \log s)} = n^2 \cdot O\left(\frac{d}{((1+3) \log n)^{1+3}}\right) = n^2 \cdot n^{O((1+3) \log \frac{d}{\log n})} = (n=s)^{O(1)};$$

for large enough. The remainder of the algorithm is the same as the proof of Theorem 8.4, and the running time is $O(n^2 = s^2) = n^2 \cdot n^{O((1+3) \log(\frac{d}{\log n}))}$. \square

Remark 8.3. For a deterministic algorithm, using Corollary 7.3 instead of Corollary 7.4, we get a deterministic running time of $n^2 \cdot n^{O((1+3) \log(\frac{d}{\log n}))}$.

The algorithm of Theorem 8.8 still has three drawbacks: (i) the exponent in the time bound depends on the dimension d , (ii) the result requires additive instead of multiplicative error, and (iii) the result is for Hamming space instead of more generally ℓ_1 or ℓ_2 . We sketch how to resolve all three issues at once, by using known dimension reduction techniques:

Theorem 8.9. Given n red and n blue points in $[U]^d$ and $\epsilon = \frac{\log^6 \log n}{\log^3 n}$, we can find a $(1 + \epsilon)$ -approximate ℓ_1 or ℓ_2 nearest/farthest blue neighbor for each red point in $(dn + n^2 \cdot n^{O((1+3) \log(1+3))}) \cdot \text{poly}(\log(nU))$ randomized time.

Proof. (The ℓ_1 case.) We first solve the decision problem for a fixed threshold value t . We use a variant of ℓ_1 locality-sensitive hashing (see [And05]) to map points from \mathbb{R}^d into low-dimensional Hamming space (providing an alternative to Kushilevitz, Ostrovsky, and Rabani's dimension reduction technique for Hamming space [KOR00]). For each red/blue point p and each $i \in \{1, \dots, k\}$, define $h_i(p) = (h_{i1}(p); \dots; h_{id}(p))$ with $h_{ij}(p) = (p_{a_{ij}} + b_j) \bmod (2t)$ where $a_{ij} \in \{1, \dots, d\}$ and $b_j \in [0; 2t)$ are independent uniformly distributed random variables. For each of the $O(n)$ hashed values of h_i , pick a random bit; let $f_i(p)$ be the random bit associated with $h_i(p)$. Finally, define $f(p) = (f_1(p); \dots; f_k(p)) \in \{0; 1\}^k$. For any fixed $p; q$

$$\Pr[h_{ij}(p) \neq h_{ij}(q)] = \frac{1}{d} \sum_{a=1}^d \min\left\{\frac{|p_a - q_a|}{2t}; 1\right\}; \text{ and so}$$

$$\Pr[f_i(p) \neq f_i(q)] = \frac{1}{2} \Pr[h_i(p) \neq h_i(q)] = \frac{1}{2} \Pr\left[\bigvee_{j=1}^k [h_{ij}(p) \neq h_{ij}(q)]\right];$$

Hence,

$$\text{If } |p - q|_1 \leq t, \text{ then } \Pr[h_{ij}(p) \neq h_{ij}(q)] \leq \frac{|p_a - q_a|}{2t} \leq \frac{1}{2d} \text{ and } \Pr[f_i(p) \neq f_i(q)] \leq \frac{1}{2} \left(1 - \left(1 - \frac{1}{2d}\right)^d\right);$$

if $\|p - q\|_1 \leq (1 + \epsilon)t$, then $\Pr[h_{ij}(p) \neq h_{ij}(q)] \leq \min\left\{\frac{\|p - q\|_1}{2dt}; \frac{1 + \epsilon}{2d}\right\}$ and $\Pr[f_i(p) \neq f_i(q)] \leq \frac{1}{2}\left(1 - \left(1 - \frac{1 + \epsilon}{2d}\right)^d\right)$.

Note that $\epsilon = \frac{1}{k}$. By a Chernoff bound, it follows (assuming $k \geq \log n$) that

if $\|p - q\|_1 \leq t$, then $\|f(p) - f(q)\|_1 \leq A_0 := \epsilon k + O(\sqrt{k \log n})$ with probability $1 - O(1/n^3)$;

if $\|p - q\|_1 \geq (1 + \epsilon)t$, then $\|f(p) - f(q)\|_1 \leq A_1 := \epsilon k + O(\sqrt{k \log n})$ with probability $1 - O(1/n^3)$.

Note that $A_1 - A_0 = O(\epsilon k)$ by setting k to be a sufficiently large constant times $(1/\epsilon)^2 \log n$. We have thus reduced the problem to an approximate problem with additive error $O(\epsilon k)$ for Hamming space of $\log n = O((1/\epsilon)^2 \log n)$ dimensions, which by Theorem 8.8 requires $n^{2 \cdot (1/\epsilon^3 = \log(1/\epsilon))}$ time. The initial cost of applying the mapping f is $O(dkn)$.

This solves the decision problem; we can then solve the original problem by calling the decision algorithm $O(\log_{1+\epsilon} U)$ times for all t 's that are powers of $1 + \epsilon$. \square

Proof. (The ℓ_2 case.) We use a version of the Johnson Lindenstrauss lemma to map from ℓ_2 to ℓ_1 (see for example [Mat08]). For each red/blue point p , define $f(p) = (f_1(p); \dots; f_k(p)) \in \mathbb{R}^k$ with $f_i(p) = \sum_{j=1}^k a_{ij} p_j$, where the a_{ij} 's are independent normally distributed random variables with mean 0 and variance 1. For each fixed $p, q \in \mathbb{R}^d$, it is known that after rescaling by a constant, $\|f(p) - f(q)\|_1$ approximates $\|p - q\|_2$ to within a $1 \pm O(\epsilon)$ factor with probability $1 - O(1/n^3)$, by setting $k = O((1/\epsilon)^2 \log n)$. It suffices to keep $O(\log U)$ -bit precision of the mapped points. The initial cost of applying the mapping f is $O(dkn)$ (which can be slightly improved by utilizing a sparse Johnson Lindenstrauss transform [AC09]). \square

Numerous applications to high-dimensional computational geometry now follow. We briefly mention just one such application, building on the work of [IM98, HIM12]:

Corollary 8.2. Given n points in $[U]^d$ and $\epsilon = \log^6 \log n = \log^3 \log n$, we can find a $(1 + \epsilon)$ -approximate ℓ_1 or ℓ_2 minimum spanning tree in $(dn + n^{2 \cdot (1/\epsilon^3 = \log(1/\epsilon))}) \text{poly}(\log(nU))$ randomized time.

Proof. Let G_r denote the graph where the vertex set is the given point set and an edge pq is present whenever p and q have distance at most r . Har-Peled, Indyk, and Motwani [HIM12] gave a reduction of the approximate minimum spanning tree problem to the following approximate connected components problem:

Given a value r , compute a partition of P into subsets with the properties that (i) two points in the same subset must be in the same component in $G_{(1+\epsilon)r}$, and (ii) two points in different subsets must be in different components in G_r .

The reduction is based on Kruskal's algorithm and increases the running time by a logarithmic factor.

To solve the approximate connected components problem, Har-Peled, Indyk, and Motwani gave a further reduction to online dynamic approximate nearest neighbor search. Since we want a reduction to offline static approximate nearest neighbor search, we proceed differently.

We first reduce the approximate connected components problem to the offline approximate nearest foreign neighbors problem:

Given a set P of n colored points with colors from $[n]$, for each point $q \in P$, find a $(1 + \epsilon)$ -approximate nearest neighbor $NF_N(q)$ among all points in P with color different from q 's color.

The reduction can be viewed as a variant of Boruvka's algorithm and is as follows: Initially assign each point a unique color and mark all colors as active. At each iteration, solve the offline approximate nearest foreign neighbors problem for points with active colors. For each q , if $NF_N(q)$ and q have distance at most $(1 + \epsilon)r$ and have different colors, merge the color class of $NF_N(q)$ and q . If a color class has not been merged to other color classes during the iteration, mark its color as inactive. When all colors are inactive, output the color classes. Otherwise, proceed to the next iteration. The correctness of the algorithm is obvious. Since each iteration decreases the number of active colors by at least a half, the number of iterations is bounded by $O(\log n)$. Thus, the reduction increases the running time by a logarithmic factor.

To finish, we reduce the offline approximate nearest foreign neighbors problem to the standard (red/blue) offline approximate nearest neighbors problem by a standard trick: For each $j = 1, \dots, \lceil \log n \rceil$, for each point $q \in P$ where the j -th bit of q 's color is 0 (resp. 1), compute an approximate nearest neighbor q' among all points $p \in P$ where the j -th bit of p 's color is 1 (resp. 0). Record the nearest among all approximate nearest neighbors found for each point q . The final reduction increases the running time by another logarithmic factor. \square

8.4 The Light Bulb Problem

In all our applications of the polynomial method for algorithm design so far, we have been focusing on optimizing the asymptotics of the exponent in the running time as a parameter of the problem (in particular, the constant factor in the dimension for the problem) grew. In this Section, we instead show techniques for optimizing the constants in the exponent of the running time. We focus our attention on the Light Bulb Problem.

Problem 8.1 (Light Bulb Problem). We are given as input a set S of n vectors from $\{-1, 1\}^d$, which are all independently and uniformly random except for two planted vectors (the correlated pair) which have inner product at least d for some $0 < \epsilon < 1$. The goal is to find the correlated pair.

Theorem 8.10. For every $\epsilon > 0$, there is a $\delta > 0$ such that the Light Bulb Problem for correlation δ can be solved in randomized time $\tilde{O}(n^{2+\epsilon})$ whenever $d = \log n$ with polynomially low error.

Proof. Our algorithm can be seen as applying the polynomial method in algorithm design for the very simple polynomial $p(x; y) = (\langle x; y \rangle)^r$.

For two constants $\epsilon, k > 0$ to be determined, we will pick $r = k^2 = 2/\epsilon$. Let $S \subseteq \{0, 1\}^d$ be the set of input vectors, and let $x^0, y^0 \in S$ denote the correlated pair which we are trying to find. For distinct $x, y \in S$ other than the correlated pair, the inner product $\langle x; y \rangle$ is a sum of uniform independent $\{0, 1\}$ values. Let $v := (k/\epsilon) \log n$. By a Chernoff bound, for large enough n , we have $|\langle x; y \rangle| \leq v$ with probability at least $1 - n^{-3}$. Hence, by a union bound over all pairs of uncorrelated vectors, we have $|\langle x; y \rangle| \leq v$ for all such x, y with probability at least $1 - 1/n$. We assume henceforth that this is the case. Meanwhile, $\langle x^0; y^0 \rangle = kv$.

Arbitrarily partition S into $m := n^{2-\epsilon}$ groups S_1, \dots, S_m of size $|S_i| = n/m = n^{1-\epsilon}$ each. We can compute the inner product between each pair of vectors which was assigned to the same group in time $\tilde{O}(m \cdot d) = \tilde{O}(n^{4-\epsilon})$, and if we find the correlated pair, we can return it and end the algorithm. Otherwise, we may assume the correlated vectors are in different groups, and we continue.

For each $x \in S$, our algorithm picks a value $a^x \in \{0, 1\}$ independently and uniformly at random. For a constant $\epsilon > 0$ to be determined, let $r = \lceil \log_k(n^{1-\epsilon}) \rceil$, and define the polynomial $p: \mathbb{R}^d \rightarrow \mathbb{R}$ by $p(z_1, \dots, z_d) = (z_1 + \dots + z_d)^r$. Our goal is, for each $(i, j) \in [m]^2$, to compute the value

$$C_{ij} := \sum_{x \in S_i} \sum_{y \in S_j} a^x a^y p(\langle x; y \rangle).$$

Solving the problem using C_{ij}

Let us first explain why we are interested in computing C_{ij} . Denote $p(x; y) := p(\langle x; y \rangle)$. Intuitively, $p(x; y)$ is computing an amplification of $\langle x; y \rangle$. C_{ij} is then summing these amplified inner products for all pairs $(x; y) \in S_i \times S_j$. We will pick our parameters so that the amplified inner product of the correlated pair is large enough to stand out from the sums of inner products of random pairs.

Let us be more precise. Recall that for uncorrelated x, y we have $|\langle x; y \rangle| \leq v$, and hence $|p(x; y)| \leq v^r$. Similarly, we have $|p(x^0; y^0)| = (kv)^r = n^{1-\epsilon} v^r$. For $x, y \in S$, define $a^{(x; y)} := a^x a^y$. Notice that, for $i \neq j$, $C_{ij} = \sum_{x \in S_i, y \in S_j} a^{(x; y)} p(\langle x; y \rangle)$, where the $a^{(x; y)}$ are pairwise independent random $\{0, 1\}$ values.

We will now analyze the random variable C_{ij} where we think of the vectors in S as fixed, and only the values a^x as random.

Consider first when the correlated pair are not in S_i and S_j . Then, C_{ij} has mean 0, and (since variance is additive for pairwise independent variables) C_{ij} has variance at most $|S_i| |S_j| \max_{x \in S_i, y \in S_j} |p(\langle x; y \rangle)|^2 = n^{2-\epsilon} v^{2r}$. For sufficiently large constant ϵ , by the Chebyshev inequality, we have that $|C_{ij}| \leq n^{1-\epsilon} v^r = 3$ with probability at least $3/4$. Let $\epsilon = n^{1-\epsilon} v^r = 3$, so $|C_{ij}| \leq 3$ with probability at least $3/4$.

Meanwhile, if $x^0 \in S_i$ and $y^0 \in S_j$, then C_{ij} is the sum of $a^{(x^0; y^0)} p(\langle x^0; y^0 \rangle)$

and a variable C^0 distributed as $C_{i;j}$ was in the previous paragraph. Hence, since $j p(hx^0, y^q) j \leq n^{1-3} v^r = 3^{-4}$, and $j C^q j \leq 2$ with probability at least 3^{-4} , we get by the triangle inequality that $j C_{i;j} j \leq 2$ with probability at least 3^{-4} .

Hence, if we repeat the process of selecting the values for each $x \in S$ independently at random $O(\log n)$ times, whichever pair $S_i; S_j$ has $j C_{i;j} j \leq 2$ most frequently will be the pair containing the correlated pair with polynomially low error, and then a brute force within this set of $O(n^{1-3})$ vectors can find the correlated pair in $O(n^{2-3})$ time. In all, by a union bound over all possible errors, this will succeed with polynomially low error.

Computing $C_{i;j}$

It remains to give the algorithm to compute $C_{i;j}$. Our overall approach will be almost identical to Proposition 8.2. However, rather than appeal directly to the statement of Proposition 8.2, we go through the details, since the running time of the reduction stated there is actually too slow for us.

We begin by rearranging the expression for $C_{i;j}$ into one which is easier to compute. Since we are only interested in the values of $C_{i;j}$ when its inputs are all in $\{0, 1\}$, we can replace ϕ with its multilinearization $\tilde{\phi}$. Let $M_1; \dots; M_t$ be an enumeration of all subsets of $[d]$ of size at most r , so $t = \sum_{i=0}^r \binom{d}{i}$. Then, there are coefficients $c_1; \dots; c_t \in \mathbb{Z}$ such that $\tilde{\phi}(x) = \sum_{s=1}^t c_s x_{M_s}$ (where, for $x \in \{0, 1\}^d$ and $M \subseteq [d]$ we define $x_M := \prod_{i \in M} x_i$). Rearranging the order of summation, we see that we are trying to compute

$$C_{i;j} = \sum_{s=1}^t \sum_{x \in S_i} \sum_{y \in S_j} a^x a^y c_s x_{M_s} y_{M_s} = \sum_{s=1}^t c_s \sum_{x \in S_i} a^x x_{M_s} \sum_{y \in S_j} a^y y_{M_s} \quad (8.3)$$

In order to compute $C_{i;j}$, we first need to compute the coefficients c_s . Notice that c_s depends only on $|M_s|$ and r . We can thus derive a simple combinatorial expression for c_s , and hence compute all of the c_s coefficients in $\text{poly}(r) = \text{polylog}(n)$ time. Alternatively, by starting with the polynomial $(z_1 + \dots + z_d)$ and then repeatedly squaring then multilinearizing, we can easily compute all the coefficients in $O(t^2 \text{polylog}(n))$ time; this slower approach is still fast enough for our purposes.

Define the matrices $A; B \in \mathbb{Z}^{m \times t}$ by $A_{i;s} = \sum_{x \in S_i} a^x x_{M_s}$ and $B_{i;s} = c_s A_{i;s}$. Notice from (8.3) that the matrix product $C := AB^T$ is exactly the matrix of the values $C_{i;j}$ we desire. A simple calculation (see Lemma 8.4 below) shows that for any $\epsilon > 0$, we can pick a sufficiently big constant $k > 0$ such that $t = O(n^{2-3+\epsilon})$. Since $m = O(n^{2-3})$, if we have the matrices $A; B$, then we can compute this matrix product by performing n^2 instances of $n^{2-3} \times n^{2-3} \times n^{2-3}$ matrix multiplication over \mathbb{Z} with $\text{polylog}(n)$ -bit entries, in $O(n^{2(2-3+\epsilon)})$ time, completing the algorithm².

¹In other words, whenever a variable appears raised to an exponent bigger than 1, we reduce that exponent mod 2 to either 0 or 1, which does not change the value of the polynomial.

²One can slightly decrease the constant $\epsilon > 0$ so that the $\text{polylog}(n)$ factors do not appear in the final running time.

Unfortunately, computing the entries of A and B naively would take $(m + t)g = (n^{5/3})$ time, which is slower than we would like. We will instead use a clever trick due to Lovett [Lov11], which was first applied in this context by Karppa et al. [KKK16]: we will compute those entries using another matrix multiplication. Let N_1, \dots, N_u be an enumeration of all subsets of $[d]$ of size at most $d/2$. For each $i \in [m]$, define the matrices $L^i, L^i \in \mathbb{Z}^{u \times g}$ (whose columns are indexed by elements $s \in S_i$) by $L^i_{s;x} = x_{N_s}$ and $L^i_{s;x} = a^x x_{N_s}$. Then, compute the product $P^i := L^i L^{i,T}$. We can see that $P^i_{s;s^0} = \sum_{x \in S_i} a^x x_{N_s} x_{N_{s^0}}$, where $N_s \oplus N_{s^0}$ is the symmetric difference of N_s and N_{s^0} . Since any set of size at most $d/2$ can be written as the symmetric difference of two sets of size at most $d/4$, each desired entry $A_{i;s}$ can be found as an entry of the computed matrix P^i . Similar to our bound on t from before (see Lemma 8.4 below), we see that for big enough constant k , we have $t = O(n^{1+3\epsilon})$. Computing the entries of the L^i matrices naively takes only $O(m(u + g)r) = O(n^4) = O(n^{4+3\epsilon})$ time, and then computing the products P^i takes $O(m \max(u, g)^2) = O(n^{(2+1)\epsilon+3\epsilon})$ time; both of these are dominated by $O(n^{2\epsilon+3\epsilon})$. This completes the algorithm! Finally, we perform the computations mentioned above in Lemma 8.4 below. \square

Lemma 8.4. For every $\epsilon > 0$, there is a $k > 0$ such that (with the same notation as in the proof of Theorem 8.10 above) we can bound $t = O(n^{2+3\epsilon})$, and $u = O(n^{1+3\epsilon})$.

Proof. Recall that $d = O(k^2 \log(n))$, and $r = \log_k(O(n^{1/3}))$. Hence, by Proposition 2.2 from the Preliminaries,

$$t = (r + 1) \binom{d}{r} (r + 1) (ed=r)^r = O(k^2 \log(k))^{\log_k(O(n^{1/3}))} = n^{2+3\epsilon} O(\log \log(k) = \log(k)).$$

For any $\epsilon > 0$ we can thus pick a sufficiently large k so that $t = O(n^{2+3\epsilon})$. We can similarly bound $\sum_{r=2}^d O(n^{1+3\epsilon})$ which implies our desired bound on u . \square

8.4.1 Deterministic Algorithms

We now present two deterministic algorithms for the Light Bulb Problem. Each is a slight variation on the algorithm from Theorem 8.10 above.

Theorem 8.11. For every $\epsilon > 0$, there is a $\delta > 0$ such that the Light Bulb Problem for correlation δ can be solved in deterministic time $O(n^{2\epsilon+3\epsilon})$ on almost all instances whenever $d = \log n$.

Recall that our goal when solving the Light Bulb Problem on almost all instances is to design a deterministic algorithm such that the probability of drawing an instance where the algorithm fails is $1 - \text{poly}(n)$.

Proof. The only randomness used by our algorithm for Theorem 8.10 was our choice of an independently and uniformly random $a^x \in \mathbb{F}_2^{1;1g}$ for each $x \in S$. Since this requires (n) random bits, and we repeat the entire algorithm $(\log n)$ times to get our desired correctness guarantee, the total number of random bits used is $(n \log n)$.

However, the only property of the x variables which we use in the proof of correctness is that they are pairwise-independent. By standard constructions, only $O(\log n)$ independent random bits are needed to generate pairwise-independent random bits. Thus, our entire algorithm actually only needs $O(\log^2 n)$ independent random bits.

Our entirely deterministic algorithm then proceeds as follows. Pick the same as in Theorem 8.10. Let $S = \{1; 1\}^d$ be the input vectors. Arbitrarily pick a subset $S^0 \subseteq S$ of $|S^0| = (\log^2 n)$ of the input vectors, and let $R = S \setminus S^0$ be the remaining vectors.

We begin by testing via brute-force whether either vector of the correlated pair is in S^0 . This can be done in $O(|S^0| |S| \cdot d) = O(n \log^2(n))$ time. If we find the correlated pair (a pair with inner product at least d), then we output it, and otherwise, we can assume that the vectors in S^0 are all uniformly random vectors from $\{1; 1\}^d$. In other words, we can use them as $|S^0| = (\log^2 n)$ independent uniformly random bits. We thus use them as the required randomness to run the algorithm from Theorem 8.10 on input vectors R . That algorithm has polynomially low error, which implies the desired correctness guarantee. \square

Theorem 8.12. There is a constant $w > 0$ such that, for every $\epsilon > 0$, there is a $\delta > 0$ such that the Promise Light Bulb Problem with parameter w for correlation can be solved in deterministic time $O(n^{4+\epsilon})$ whenever $d = \log n$.

Recall that in the Promise Light Bulb Problem with parameter w , we are promised that every pair of vectors other than the correlated pair has inner product at most $w \cdot d \log n$, and our deterministic algorithm needs to solve the problem correctly on every input with this guarantee.

Proof. The guarantee of the Promise Light Bulb Problem is that, when we pick a sufficiently large w , the uncorrelated vectors have as small inner product as we assumed they did in the first paragraph in the proof of Theorem 8.10. In other words, there is a quantity v such that $\langle x; y \rangle \leq v$ for all $x; y \in S$ other than the correlated pair, and moreover, $\langle x^0; y^0 \rangle \geq kv$ for a constant $k > 0$ with $k \rightarrow 1$ as $w \rightarrow 1$.

The algorithm is then almost identical to Theorem 8.10, except we need to remove the only use of randomness: the randomness used to pick the a^x values. To do this, we will simply pick $a^x = 1$ for all x .

In order to guarantee the correctness of our algorithm, we must now change the parameters slightly. Instead of partitioning the input into $m = n^{2/3}$ groups of size $g = n^{1/3}$, we will instead partition into $m = n^{4/5}$ groups of size $g = n^{1/5}$. Similarly, instead of picking r (the exponent in the polynomial p) to be $\log_k(O(n^{1/3}))$, we will pick $r = \log_k(3n^{2/5})$, so that $p(x^0; y^0) = (kv)^r = 3n^{2/5}v^r$.

With these choices, for any i and j such that the correlated pair are not in S_i and S_j , we have $\langle C_{i;j} \rangle \leq |S_i| |S_j| n^{2/5} = n^{2/5}v$, whereas if $x^0 \in S_i$ and $y^0 \in S_j$ then by the triangle inequality, $\langle C_{i;j} \rangle \geq p(x^0; y^0) - |S_i| |S_j| n^{2/5} \geq 2n^{2/5}v^r$. Hence, the correlated pair must be in whichever S_i and S_j with $i \neq j$ has the largest $\langle C_{i;j} \rangle$.

³For one example, to generate ℓ pairwise-independent bits, pick only ℓ bits $b_1; \dots; b_\ell$ independently and uniformly at random, and then output, for each $l \in [1; \ell]$, the product $\prod_{i \neq l} b_i$.

The algorithm to compute the $C_{i,j}$ values is identical to that of Theorem 8.10. We now get that $t = \sum_{i=0}^r \binom{d}{i} O(n^{4+5i})$ and similarly, $u = O(n^{2+5i})$, which leads to a total running time of $O(n^{4+5r})$, as desired. \square

8.5 Faster Algorithms For MAX-SAT

Next, we apply our probabilistic PTFs for threshold functions to obtain faster algorithms for MAX-SAT for sparse instances with cn clauses. We first consider MAX-k-SAT for small k before solving the general problem:

Theorem 8.13. Given a k -CNF formula F (or k -CSP instance) with n variables and $cn = n^4 = (k^4 \log^6 n)$ clauses, we can find an assignment that satisfies the maximum number of clauses (constraints) of F in randomized $2^{n^{O(k^4 c^{1/3} \log(kc))}}$ time.

Proof. We proceed as in the k -SAT algorithm of Chan and Williams [CW16]. We first solve the decision problem of testing whether there is a variable assignment satisfying more than t clauses for a fixed $t \in [cn]$. Let $s = \alpha n$ for some parameter $\alpha < 1/2$ to be set later.

For $j \in [cn]$, define the function $C_j(x_1, \dots, x_n)$ to output 1 if the j -th clause of the given formula is satisfied, and 0 otherwise. Note that each C_j can be expressed as a polynomial of degree at most k .

Say that a variable is good if it occurs in at most $2kc$ clauses. By the pigeonhole principle, at least half of the variables are good, so we can find good variables x_1, \dots, x_s . Let x_{s+1}, \dots, x_n be the remaining variables, and let $J \subseteq [cn]$ be the set of indices of all clauses C_j that contain some occurrence of a good variable; note that $|J| = O(kcs)$. Now for every variable assignment $(x_{s+1}, \dots, x_n) \in \{0, 1\}^{n-s}$, we want to compute

$$F(x_{s+1}, \dots, x_n) := \sum_{(a_1, \dots, a_s) \in \{0, 1\}^s} \sum_{j \in J} C_j(a_1, \dots, a_s, x_{s+1}, \dots, x_n) > t$$

We will achieve this by computing for every $t \in [cn]$:

$$G_t(x_{s+1}, \dots, x_n) := \sum_{(a_1, \dots, a_s) \in \{0, 1\}^s} \sum_{j \in J} C_j(a_1, \dots, a_s, x_{s+1}, \dots, x_n) > t$$

Let us define $T[x_{s+1}, \dots, x_n] := \sum_{j \in J} C_j(0, \dots, 0, x_{s+1}, \dots, x_n)$. (Observe that it is not a problem to set the good variables x_1, \dots, x_s to zero here, because we are only summing over clauses that do not contain them.) Note that T can be viewed as a polynomial in $n-s$ variables with only $\text{poly}(n)$ monomials. Therefore for all $(x_{s+1}, \dots, x_n) \in \{0, 1\}^{n-s}$, these T -values can be precomputed in $\text{poly}(n) 2^{n-s}$ time. As these T -values are measuring the contribution from the variables x_{s+1}, \dots, x_n to the number of satisfied clauses, we have

$$F(x_{s+1}, \dots, x_n) = G_{T[x_{s+1}, \dots, x_n]}(x_{s+1}, \dots, x_n):$$

Applying Corollary 7.4 (in the exact setting), we can express an \mathcal{G}_{t^0} as a sum of 2^s probabilistic PTFs of degree $k = O((kcs)^{1/3}(s + \log(kcs))^{2/3})$, where each probabilistic PTF computes an expression of the form $\prod_{j \in J} p_j(x_{s+1}; \dots; x_n)$ with error probability at most $1/(10 \cdot 2^s)$, and for all $j \in J$ we have $\deg(p_j(x_{s+1}; \dots; x_n)) \leq k$. The number of monomials in our probabilistic PTF for \mathcal{G}_{t^0} is at most

$$2^s \cdot k^{O((kcs)^{1/3}(s + \log(kcs))^{2/3})} = 2^n \cdot O\left(\frac{n}{k^{4/3}c^{1/3}}\right)^{O(k^{4/3}c^{1/3}n)} = 2^n \cdot 2^{O(k^{4/3}c^{1/3} \log \frac{1}{k})n} = 2^{O(1)n}$$

by setting ϵ to be a sufficiently small constant times $1/(k^{4/3}c^{1/3} \log(kc))$. The same bound holds for the construction time of the polynomial.

For each t^0 , we can evaluate the polynomial $f_{\mathcal{G}_{t^0}}$ at all 2^{n-s} input values by divide-and-conquer or dynamic programming using $\text{poly}(n)2^{n-s}$ arithmetic operations [Yat37, Wil14c] on $\text{poly}(n)$ -bit numbers. The total time is $2^n \cdot n^{O(k^{4/3}c^{1/3} \log(kc))}$. As before, the error probability can be lowered by taking the majority values over $O(n)$ repetitions, and the original problem can be solved by calling the decision algorithm for at most cn times. \square

Theorem 8.14. Given a CNF formula with n variables and $cn = n^{4/\log 10} n$ clauses, we can find an assignment that satisfies the maximum number of clauses in randomized $2^n \cdot n^{O(c^{1/3} \log^{7/3} c)}$ time.

Proof. We use a standard width reduction technique [SST15] originally observed by Schuler [Sch05] and studied closely by Calabro, Impagliazzo, and Paturi [CIP06]. Consider the following recursive algorithm:

If all clauses have length at most k , then call the algorithm from Theorem 8.13 and return its output.

Otherwise, pick a clause $(x_1 \vee \dots \vee x_\ell)$ with $\ell > k$. Return SAT if at least one of the two following calls return SAT:

Recursively solve the instance in which $(x_1 \vee \dots \vee x_\ell)$ is replaced by $(x_1 \vee \dots \vee x_k)$, and

recursively solve the instance in which $x_1; \dots; x_k$ are all assigned false

Sakai, Seto, and Tamaki's analysis for MAX-SAT [SST15] can be directly modified to show that the total time of this algorithm remains $2^n \cdot n^{O(k^{4/3}c^{1/3} \log(kc))}$, when the parameter k is set to be a sufficiently large constant times $\log c$. \square

For MAX- k -SAT with $k \geq 4$, we can obtain a much better dependency on the sparsity parameter c ; in fact, we obtain significant speedup even for general dense instances. The approach this time requires only our probabilistic polynomials for threshold functions. Naively, the dense case seems to require threshold functions with superlinearly many arguments, but by incorporating a few new ideas, we manage to solve MAX-4-SAT using only $O(n)$ -variate threshold functions.

Theorem 8.15. Given a weighted 4-CNF formula \mathcal{F} with n variables with positive integer weights bounded by $\text{poly}(n)$, we can find an assignment that maximizes the

total weight of clauses satisfied in F , in randomized $2^{n - n = O(\log^2 n \log^2 \log n)}$ time. In the sparse case when the clauses have total weight n , the time bound improves to $2^{n - n = O(\log^2 c \log^2 \log c)}$.

Proof. (Dense case.) Let $s = n$ for some parameter to be set later. Arbitrarily divide the n variables of F into three groups: $x = f x_1; \dots; x_{(n-s)=2} g$, $y = f y_1; \dots; y_{(n-s)=2} g$, and $z = f z_1; \dots; z_s g$. As in Theorem 8.13, it suffices to solve the decision problem of whether there exist $x; y \in \{0, 1\}^{(n-s)=2}$ and $z \in \{0, 1\}^s$ such that $f(x; y; z) > t$, for a given degree-4 polynomial f and a fixed $t \in [n^{c_0}]$ (for an appropriately large constant c_0). Since f has degree 4, observe that each term has either (a) at most one x variable, (b) at most one y variable, or (c) no z variable. We can thus write

$$f(x; y; z) = \sum_{i=1}^{(n-s)=2} f_i(x; z) y_i + \sum_{i=1}^{(n-s)=2} g_i(y; z) x_i + h(x; y)$$

where the f_i 's and g_i 's are degree-3 polynomials, and h is a degree-4 polynomial.

For every $x; y \in \{0, 1\}^{(n-s)=2}$, it suffices to compute

$$F(x; y) := \sum_{z \in \{0, 1\}^s} [f(x; y; z) > t]$$

More generally, we compute for every $t \in [n^{c_0}]$:

$$G_{t^0}(x; y) := \sum_{z \in \{0, 1\}^s} H_{z; t^0}(x; y);$$

with

$$H_{z; t^0}(x; y) := \sum_{i=1}^{(n-s)=2} f_i(x; z) y_i + \sum_{i=1}^{(n-s)=2} g_i(y; z) x_i > t^0 \quad (3)$$

Then $F(x; y) = G_{t^0}(x; y)$; we can precompute all $h(x; y)$ values in $\text{poly}(n) 2^{n-s}$ time.

The $H_{z; t^0}(x; y)$ predicate can be viewed as a weighted threshold function with $O(n)$ arguments. To further complicate matters, these weights are not fixed: they depend on x and y . We resolve the issue by extending the vectors x and y and using a binary representation trick.

For each vector $x \in \{0, 1\}^{(n-s)=2}$, define an extended vector \tilde{x} where $\tilde{x}_i = x_i$ for each $i = 1; \dots; (n-s)=2$ and $\tilde{x}_{i; j; z}$ is the j -th least significant bit in the binary representation of $f_i(x; z)$ for each $i = 1; \dots; (n-s)=2, j = 0; \dots; \ell$ and $z \in \{0, 1\}^s$, with $\ell = O(\log n)$. Note that \tilde{x} is a vector in $O(n \log n 2^s)$ dimensions. Similarly, for each vector $y \in \{0, 1\}^{(n-s)=2}$, define an extended vector \tilde{y} where $\tilde{y}_i = y_i$ for each $i = 1; \dots; (n-s)=2$ and $\tilde{y}_{i; j; z}$ is the j -th least significant bit in the binary representation of $g_i(y; z)$ for each $i = 1; \dots; (n-s)=2, j = 0; \dots; \ell$ and $z \in \{0, 1\}^s$. We can precompute all extended vectors in $\text{poly}(n) 2^s$ time.

Then

$$H_{z;t^0}(x; y) := \sum_{(t_0; \dots; t^0) \in \{0, 1\}^s} \prod_{j=0}^{s-1} \sum_{i=1}^{2^j} x_{i;j;z} y_i + \prod_{i=1}^{2^s} y_{i;j;z} x_i = t_j$$

where the outer sum is over all tuples $(t_0; \dots; t^0) \in \{0, 1\}^s$ with $\sum_{j=0}^{s-1} 2^j t_j > t^0$.

By Fact 7.6.2, for each $z \in \{0, 1\}^s$, $j = 0, \dots, s-1$, and $t_j \in \{0, 1\}$, we can construct a probabilistic polynomial (over \mathbb{R} or \mathbb{F}_2) for the predicate $\sum_{i=1}^{2^j} x_{i;j;z} y_i + \sum_{i=1}^{2^j} y_{i;j;z} x_i = t_j$ with degree $O(\sqrt{n \log s})$ with error probability at most $1/s$. By the union bound, the probability that there is an error for some $z; j; t_j$ is at most $O((1/s) 2^s \log n) = O(1/n)$, which can be made at most $1/4^s$, for example, by setting $s = n^{c_0} 2^s$ for a sufficiently large constant c_0 . Thus, the degree for each predicate $\mathcal{G}_j(\bar{s})$ (assuming $s = O(\log n)$).

For each $z \in \{0, 1\}^s$ and $t^0 \in \{0, 1\}^s$, by distributing over the product $\prod_{j=0}^{s-1}$ we can then construct a probabilistic polynomial for $H_{z;t^0}(x; y)$ with degree $O(\sqrt{ns \log n})$. For a fixed z and t^0 , such a polynomial is a function of $O(n \log n)$ free variables in x and y , and therefore has at most $O(n \log n)$ monomials. The same bound holds for the time needed to construct the probabilistic polynomial (note the number of tuples $(t_0; \dots; t^0)$ is $n^{O(\log n)}$, which is a negligible factor).

For each $t^0 \in \{0, 1\}^s$, we can thus construct a probabilistic polynomial for $G_{t^0}(x; y)$ with degree $O(\sqrt{ns \log n})$ over x and y , with the following number of monomials:

$$2^s \cdot O(\sqrt{ns \log n}) \cdot 2^n \cdot O\left(\frac{n \log n}{n \log n}\right)^{O(\sqrt{n \log n})} = 2^{0.1(n-s)^2}$$

by setting ϵ to be a sufficiently small constant times $1/(\log n \log \log n)^2$. The same bound holds for the construction time.

We can rewrite the polynomial for $G_{t^0}(x; y)$ as the dot product of two vectors (x) and (y) of $2^{0.1(n-s)^2}$ dimensions. The problem of evaluating $G_{t^0}(x; y)$ over all $x; y \in \{0, 1\}^{(n-s)^2}$ then reduces to multiplying a $2^{(n-s)^2} \times 2^{0.1(n-s)^2}$ matrix (over \mathbb{R} or \mathbb{F}_2), which can be done in $\text{poly}(n) 2^n$ time (Lemma 8.1). The total time is $2^{n - n = O(\log^2 n \log^2 \log n)}$. \square

Proof. (Sparse case.) If the clauses have total weight cn , we can refine the analysis above, as follows. Let p_i and q_i be the maximum value of $f_i(x; z)$ and $g_i(y; z)$ respectively. We know that $p_i + q_i \leq cn$. The variable $x_{i;j;z}$ is needed only when $j \leq \log(p_i)$, and the variable $y_{i;j;z}$ is needed only when $j \leq \log(q_i)$. For each $z; j; t_j$, the probabilistic polynomial for the predicate

$$\sum_{i \in I} x_{i;j;z} y_i + \sum_{i \in J} y_{i;j;z} x_i = t_j$$

has degree $O(\sqrt{n_j s})$, where n_j is the number of i 's with $p_i \geq 2^j$ or $q_i \geq 2^j$.

Observe that $n_j = O(cn/2^j)$. It follows that the degree for the $H_{z;t^0}(x; y)$

polynomial is $O(\sum_{j=0}^p \binom{p}{j} n^j s) = O(p \overline{ns} \log c + \sum_{j > \log c}^p \binom{p}{j} (cn=2^j)s) = O(p \overline{ns} \log c)$.
 The number of variables in $H_{z,t^0}(x; y)$ is at most $O(\sum_{j=0}^p \binom{p}{j} n_j) = O(n \log c + \sum_{j > \log c}^p (cn=2^j)) = O(n \log c)$.

Thus, the bound on the total number of monomials becomes

$$2^s \frac{O(n \log c)}{O(p \overline{ns} \log c)} = 2^n O\left(\frac{n \log c}{p \overline{ns} \log c}\right)^{O(p \overline{n} \log c)}$$

$$2^n \frac{2^{p \overline{n} \log c \log(1=)}}{2^{0.1(n-s)=2}}$$

by setting ϵ to be a sufficiently small constant times $1=(\log c \log \log c)^2$. □

8.6 Circuit Satisfiability Algorithms

In this Section, we give new algorithms for solving the SAT problem on some rather expressive circuit classes. First, we outline some notions used in all of these algorithms.

8.6.1 Satisfiability on a Cartesian Product

In intermediate stages of our SAT algorithms, we will study the following generalization of SAT, where the task is to find a SAT assignment in a Cartesian product of possible assignments.

Definition 8.1. Let n be even, and let $A; B \subseteq \{0, 1\}^{n=2}$ be arbitrary. The SAT problem on the set $A \times B$ is to determine if a given n -input circuit has a satisfying assignment contained in the set $A \times B$.

Recall that a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a linear threshold function (LTF) if there are $a_1, \dots, a_n; t \in \mathbb{R}$ such that for all $x \in \{0, 1\}^n$, $f(x) = 1 \iff \sum_i a_i x_i \geq t$.

Let $\text{Circuit-LTF}[Z; S]$ be the class of circuits with a layer of S LTFs at the bottom (nearest the inputs), with Z additional arbitrary gates above that layer. Let $\text{Circuit-SUM-AND}[Z; S]$ be the analogous circuit class, but with S DNFs at the bottom layer with property that each DNF always has at most one conjunct true for every variable assignment. (Thus we may think of the DNF as simply an integer sum.) We first prove that the SAT problem for Circuit-LTF can be reduced to the SAT problem for Circuit-SUM-AND , utilizing a weight reduction trick that can be traced back to Matoušek's algorithm for computing dominances in high dimensions [Mat91, Wil14b]:

Lemma 8.5. Let $A; B \subseteq \{0, 1\}^{n=2}$, with $|A| = |B| = N \leq 2^n$. Let $K \in [1; N]$ be an integer parameter. The SAT problem for $\text{Circuit-LTF}[Z; S]$ circuits on the set $A \times B$ can be reduced to the SAT problem for $\text{Circuit-SUM-AND}[Z; S]$ where each DNF has at most $O(\log K)$ terms and each AND has fan-in at most $2 \log K$, on a prescribed set $A^0 \times B^0$ with $|A^0| = |B^0| = N$ and $A^0; B^0 \subseteq \{0, 1\}^{2S \log K}$. The reduction has the

property that if the latter SAT problem can be solved in time $\tilde{O}(N^2 Z^2 = K + N^S) \text{ poly}(n)$, then the former SAT problem can be solved in time $\tilde{O}(T + N^2 Z^2 = K + N^S) \text{ poly}(n)$.

Proof. For a given circuit C of type $\text{Circuit-LTF}[Z; S]$, let the j th LTF in the bottom layer have weights $t_{j,1}; \dots; t_{j,n}; t_j$. Let the assignments in A be $a_1; \dots; a_N$, and let the assignments in B be $b_1; \dots; b_N$. Denote the k th bit of a_i and b_i as $a_i[k]$ and $b_i[k]$, respectively.

Make N^S matrices M_A and M_B , where

$$M_A(i; j) = \sum_{k=1}^{Z^2} t_{j,k} a_i[k]$$

and

$$M_B(i; j) = t_j \sum_{k=1}^{Z^2} b_{i=2+k}[k]$$

The key property of these matrices is that $M_A(i; j) = M_B(i^0; j)$ if and only if the n -variable assignment $(a_i; b_{i^0})$ makes the j th LTF output 1.

For each $j = 1; \dots; S$, let L_j be the list of all 2^N entries in the j th column of M_A and the j th column of M_B , sorted in increasing order. Partition L_j into K contiguous parts of $O(N=K)$ entries each, and think of each part of L_j as containing a set of $O(N=K)$ assignments from $A \cup B$. (So, the partition of L_j is construed as a partition of the assignments in $A \cup B$.) There are two possible cases for a satisfying assignment to the circuit C :

1. There is a satisfying assignment $(a_i; b_{i^0}) \in A \cup B$ such that for some $j = 1; \dots; S$, a_i and b_{i^0} are in the same part of L_j . By enumerating every $a_i \in A$, every $b_{i^0} \in B$, $j = 1; \dots; S$, and all $O(N=K)$ assignments b_{i^0} of B which are in the same part of L_j as a_i , then evaluating the circuit C on the assignment $(a_i; b_{i^0})$ in $Z^2 \text{ poly}(n)$ time, we can determine satisfiability for this case in $O(N^S N=K Z^2) \text{ poly}(n)$ time. If this does not uncover a SAT assignment, we move to the second case.
2. There is a satisfying assignment $(a_i; b_{i^0}) \in A \cup B$ such that for every $j = 1; \dots; S$, a_i and b_{i^0} are different parts of L_j . Then for every LTF gate $j = 1; \dots; S$ on the bottom layer of the circuit, we claim that the j -th LTF can be replaced by a sum of $O(\log K)$ ANDs on $2 \log K$ new variables. In particular, for the j -th LTF we define one new set of $\log K$ variables which encodes the index $k = 1; \dots; K$ such that a_i is in part k of L_j , and another set of $\log K$ variables which encodes the index k^0 such that b_{i^0} is in part k^0 of L_j . Then, determining $[k = k^0]$ is equivalent to determining whether $(a_i; b_{i^0})$ satisfies the j -th LTF gate. Finally, note that the predicate $[k = k^0]$ can be computed by a DNF of $O(\log K)$ conjuncts. (Take an OR over all $\ell = 0; \dots; \log K$, guessing that the ℓ -th bit is the most significant bit in which k and k^0 differ; we can verify that guess with a conjunction on $2 \log K$ variables.) On every possible input $(k; k^0) \in \{0; \dots; 1\}^{2 \log K}$, the DNF has at most one true conjunction. Thus we can construe the OR as

simply an integer sum of ANDs, as desired. Preparing these new assignments for this new SAT problem takes time $O(N^S) = \text{poly}(n)$. \square

8.6.2 Simulating LTFs with AC⁰ of MAJORITY

In our SAT algorithms, we will need a way to simulate LTFs with bounded-depth circuits with MAJORITY gates. This was also used in Williams' work on solving ACC-LTF SAT [Wil14b], as a black box. However, here we must pay careful attention to the details of the construction. In fact, we will actually have to modify the construction slightly in order for our circuit conversion to work out. Let us review the construction here, and emphasize the parts that need modification for this algorithm. Recall that MAJ denotes the majority function.

Theorem 8.16 (Follows from [MT98], Theorem 3.3) Every LTF can be computed by polynomial-size AC⁰ MAJ circuits. Furthermore, the circuits can be constructed in polynomial time given the weights of the LTF, and the fan-in of each MAJ gate can be made $n^{1+\epsilon}$, for every desired $\epsilon > 0$, and the circuit has depth $O(\log(1/\epsilon))$.

It will be crucial for our final results that the fan-in of the MAJ gates can be made arbitrarily close to linear.

Proof. We begin by revisiting the circuit construction of Maciel and Thérien [MT98], which shows that the addition of n distinct n -bit numbers can be performed with polynomial-size AC⁰ MAJ circuits. The original construction of Maciel and Thérien yields MAJ gates of fan-in $O(n^2)$, which is too large for our purposes. We can reduce the fan-in of MAJ gates to $O(n^{1+\epsilon})$ by setting the parameters differently in their construction. Let us sketch their construction in its entirety, then describe how to modify it.

Recall that SYM denotes the class of symmetric functions. First, we show that addition of n n -bit numbers can be done in AC⁰ SYM. Suppose then n -bit numbers to be added are A_1, \dots, A_n , where $A_i = A_{i;n} \dots A_{i;1}$ for $A_{j;i} \in \{0, 1\}$. Maciel and Thérien partition each A_i into m blocks of ℓ bits, where $m \cdot \ell = n$. They compute the sum S_k of the n -bit numbers in each block $k = 1, \dots, m$, i.e.

$$S_k = \sum_{i=1}^n \sum_{j=1}^{\ell} A_{i;(k-1)\ell+j} 2^{j-1};$$

and note that the desired sum is

$$Z = \sum_{k=1}^m S_k 2^{(k-1)\ell};$$

Each S_k can be represented in $\ell + \log n$ bits. Maciel and Thérien set $\ell = \log n$, so that each S_k is represented by 2ℓ bits. They then split each S_k into ℓ -bit numbers H_k and L_k such that

$$S_k = H_k 2^{\ell} + L_k;$$

Note that the high part H_k corresponds to the carry bits of S_k . They then note that if

$$y_1 := \sum_{k=1}^n H_k 2^{k-1}; \quad y_2 := \sum_{k=1}^n L_k 2^{(k-1)^2};$$

we have

- (a) $z = y_1 + y_2$, and
- (b) each bit of y_i is a function of exactly one H_k or L_k for some k . In turn, each L_k, H_k is a sum of $n^{1+\epsilon}$ $A_{i,j}$'s where each $A_{i,j}$ is multiplied by a power of two in $[0, 2^k]$. Therefore, each bit of y_i can be computed by a SYM gate of fan-in at most $n^{1+\epsilon} 2^k \leq n^2$.

We have therefore reduced the addition of n -bit numbers to adding the two $O(n)$ -bit numbers y_1 and y_2 , with a layer of SYM gates. Adding two numbers can be easily computed in AC^0 (see for example [CFL85]), so the whole circuit is of the form AC^0 SYM.

We wish to reduce the fan-in of the SYM gates to $O(n^{1+\epsilon})$ for arbitrary $\epsilon > 0$. To reduce the fan-in further, it suffices to find a construction that lets us reduce ϵ . Naturally, we can try to set $\epsilon = \epsilon \log n$ for arbitrarily small $\epsilon \in (0, 1)$. Without loss of generality, let us assume $\epsilon \log n$ is an integer. Then, each S_k is represented in $\epsilon \log n + 1 = t$ bits. Let $t = \epsilon \log n + 1$. If we then split each S_k into t ϵ -bit numbers T_k^{t-1}, \dots, T_k^0 , ranging from high-order to low-order bits, we then have

$$S_k = T_k^{t-1} 2^{(t-1)\epsilon} + \dots + T_k^1 2^\epsilon + T_k^0.$$

Defining the t numbers

$$y_i := \sum_{k=1}^n T_k^i 2^{(k+i-1)\epsilon};$$

the desired sum is $z = \sum_{i=0}^{t-1} y_i$. Just as before, each bit of y_i is a function of exactly one T_k^i for some k , which is a sum of $n^{1+\epsilon}$ $A_{i,j}$'s where each $A_{i,j}$ is multiplied by an integer in $[0, 2^k]$. Hence each bit of y_i can be computed by a SYM gate of fan-in at most $n^{1+\epsilon} 2^k = O(n^{1+\epsilon})$. So with one layer of SYM gates, we have reduced the number n -bit addition problem to the addition of t $O(n)$ -bit numbers y_0, \dots, y_{t-1} . But for $t = \epsilon \log n$, addition of t n -bit numbers can be computed by AC^0 circuits of poly(n)-size and fixed depth independent of t (see e.g. [Vol99, p.14-15]). This completes the description of our AC^0 SYM circuit.

Observe that each SYM gate can be easily represented by a OR AND MAJ circuit. In particular, the OR is over all $j \in \{0, 1, \dots, t-1\}$ such that the SYM gate outputs 1 when given j inputs are equal to 1, and the AND MAJ part computes $\sum_j x_j = j$. Again, the fan-in of each MAJ here is $O(n^{1+\epsilon})$.

We now apply the addition circuits to show how every LTF on variables can be represented by a polynomial-size AC^0 MAJ circuit. Suppose our LTF has weights w_1, \dots, w_{n+1} , computing $\sum_{j=1}^n w_j x_j \leq w_{n+1}$. By standard facts about LTFs, we may assume for all j that $|w_j| \leq 2^{bn \log_2 n}$ for some constant $b > 0$. Set $W = bn \log_2 n$.

Let D be a AC^0 MAJ circuit for adding n W -bit numbers as described above,

where each MAJ gate has fan-in $\Theta(n^{1+\epsilon})$. For all $j = 1; \dots; n$, connect to the j th W -bit input of D a circuit which, given x_j , feeds w_j to D if the input bit $x_j = 1$, and the all-zero W -bit string if $x_j = 0$. Observe this extra circuitry is only wires, no gates: we simply place a wire from x_j to all bits of the j th W -bit input where the corresponding bit of w_j equals 1.

This new circuit D^0 clearly computes the linear form $\sum_{j=1}^n w_j x_j$. The linear form can then be compared to w_{n+1} with an AC^0 circuit, since the less-than-or-equal-to comparison of two integers can be performed in AC^0 . Indeed, this function can be represented as a quadratic-size DNF (SUM AND), as was noticed in Lemma 8.5. We now have an AC^0 MAJ circuit D^0 of size $\text{poly}(W; t) \cdot n^b$ computing the LTF, where the MAJ gates have fan-in $\Theta(n^{1+\epsilon})$. \square

8.6.3 Satisfiability for ACC of LTF of LTF

Let $AC^0[d; m]$ LTF LTF[$S_1; S_2; S_3$] be the class of circuits with a layer of S_3 LTFs at the bottom layer (nearest the inputs), a layer of S_2 LTFs above the bottom layer, and a size S_1 $AC^0[m]$ circuit of depth d above the two LTF layers.

Theorem 8.17. For every integer $d > 0$, $m > 1$, and $\epsilon > 0$, there is an $n^0 > 0$ and an algorithm for satisfiability of $AC^0[d; m]$ LTF LTF[$2^{n^0}; 2^{n^0}; n^2$] circuits that runs in deterministic 2^{n^0} time.

Before giving the proof, we first sketch the ideas in this SAT algorithm for ACC^0 LTF LTF. Similar to the SAT algorithm for ACC^0 LTF circuits [Wil14b], the bottom layer of LTFs can be replaced by a layer of DNFs, via a weight reduction trick. We replace LTFs in the middle layer with AC^0 MAJ circuits (modifying a construction of Maciel and Thérien [MT98] to keep the fan-in of MAJ gates low), then replace these MAJ gates of fan-in n^2 with probabilistic F_2 -polynomials of degree $n^{1+\epsilon}$ over a small sample space, provided by Theorem 7.4. Taking a majority vote over all samples, and observing that an F_2 -polynomial is a MOD₂ AND circuit, we obtain a MAJ ACC^0 circuit, but with $2^{n^{1+\epsilon}}$ size in some of its layers. By carefully applying known depth reduction techniques, we can convert the circuit into a depth-two circuit of size $2^{n^{1+\epsilon}}$ which can then be evaluated efficiently on many inputs. (This is not obvious: applying the Beigel-Tarui depth reduction to a $2^{O(n^{1+\epsilon})}$ -size circuit would make its new size quasi-polynomial in $2^{O(n^{1+\epsilon})}$, yielding an intractable bound of $2^{n^{O(1)}}$.)

We now move on to the proof of Theorem 8.17. We use the following depth-reduction theorem of Beigel and Tarui (with important constructibility issues clarified by Allender and Gore [AG94], and recent size improvements by Chen and Papakonstantinou [CP19]):

Theorem 8.18 ([BT94, AG94]). Every SYM ACC circuit of size s can be simulated by a SYM AND circuit of $2^{(\log s)^{c_0}}$ size for some constant c_0 depending only on the depth d and MOD m gates of the ACC part. Moreover, the AND gates of the final circuit have only $(\log s)^{c_0}$ fan-in, the final circuit can be constructed from the original in $2^{O((\log s)^{c_0})}$ time, and the final symmetric function at the output can be computed in $2^{O((\log s)^{c_0})}$ time.

Proof of Theorem 8.17. Let $\epsilon > 0$ be a parameter to be set later. The plan is to start with a circuit as specified in the theorem statement, and slowly convert into a nice form that can be evaluated efficiently on many inputs.

1. Trade Variables for Circuit Size. Our first step is standard for ACCSAT algorithms [Wil14b, Wil14c]: given an $ACC^0[d; m]$ LTF LTF $[2^n; 2^n; n^2]$ circuit C with n variables, create a copy of the circuit $C_v := C(v; \cdot)$ for all possible assignments $v \in \{0, 1\}^{2^n}$ to the first 2^n variables of C , and define

$$C^0(x_{n^2+1}; \dots; x_n) := \bigvee_v C_v(x_{n^2+1}; \dots; x_n)$$

Observe that C^0 is satisfiable if and only if C is satisfiable, C^0 has size at most 2^{2^n} , C^0 is also an ACC^0 LTF LTF circuit, and C^0 has only $n - n^2$ variables.

2. Replace the middle LTFs with MAJORITYs (Theorem 8.16). Note that each LTF on the second layer of C^0 has fan-in at most $n^2 + n$, since the number of LTFs on the first layer is n^2 . Applying the low fan-in transformation of Theorem 8.16, we can replace each of the LTFs on the second layer of C^0 with poly(n)-size ACC^0 MAJ circuits where each MAJ has fan-in at most $n^2 = 2$. This generates at most 2^{2^n} new MAJ gates in the circuit C^0 , for some constant $d > 0$, and produces a circuit of type

$$ACC^0 \text{ MAJ LTF}$$

3. Replace those MAJORITYs with (derandomized) probabilistic polynomials over F_2 (Theorem 7.4). We replace each of these new MAJ gates with our low-randomness probabilistic polynomials for the MAJORITY function, as follows. Recall from Theorem 7.4 that we can construct a probabilistic polynomial over F_2 for k -bit MAJORITY with degree $O(\sqrt{k \log(1/\epsilon)})$ and error at most ϵ , using a distribution of $k^{O(\log(k/\epsilon))}$ uniformly chosen F_2 -polynomials. Setting $k := n^2 = 2$ for the fan-in of the MAJ gates, and the error to be $\epsilon := 1/2^{2^{2^n}}$, the degree becomes

$$D := O\left(\sqrt{n^2 = 2 \cdot 2^{2^n}}\right) = O(n^{1/2 + \epsilon/2})$$

and the sample space has size $2^D = n^{O(n)}$. For $\epsilon = 4$, we have $D := O(n^{1/2 + 2})$, and each polynomial in our sample space has at most $n^{1/2 + 8} = 2^{O(n^{1/2 + 8 \log n})}$ monomials. For every choice of the random seed to the probabilistic polynomial, let C_r^0 be the circuit C^0 with the corresponding F_2 polynomial P_r substituted in place of each MAJ gate. That is, each MAJ gate is substituted by an XOR of $2^{O(n^{1/2 + 8 \log n})}$ ANDs of fan-in at most $O(n^{1/2 + 8})$.

We now form a circuit C^{00} which takes a majority vote over all $2^{O(n^{1/2 + 8 \log n})}$ circuits C_r^0 . The new circuit C^{00} therefore has the form

$$\text{MAJ ACC}^0 \text{ XOR AND LTF}$$

where the MAJ ACC⁰ part has size $2^{O(n^{1/2 + 8 \log n})}$, and each XOR AND LTF subcircuit has size $2^{O(n^{1/2 + 8 \log n})}$. Since our probabilistic polynomial computes MAJORITY with

$1=2^{2dn}$ error and there are at most 2^{2dn} MAJ gates in C^0 , the new circuit C^{00} is equivalent to the original circuit C^0 .

4. Apply Beigel Tarui to the top of the circuit, and distribute. It is very important to observe that we cannot apply Beigel Tarui (Theorem 8.18) to the entire circuit C^{00} , as its total size is $2^{(n^{1+8 \log n})}$, and the quasi-polynomial blowup of Beigel Tarui would generate a huge circuit of size (2^n) , rendering our conversion intractable.

However, the top MAJ ACC⁰ part is still small. Invoking the depth reduction lemma of Beigel and Tarui (Theorem 8.18 above), we can replace the MAJ ACC⁰ part in C^{00} of size $2^{O(n \log n)}$ (even though it has $2^{O(n \log n)}$ inputs from the XOR layer!) with a SYM AND circuit of size 2^{n^a} for a constant $a > 1$, where each AND has fan-in at most n^a , and a depends only on the (constant) depth d and (constant) modulus m of the ACC⁰ subcircuit.

The resulting circuit C_3 now has the form

SYM AND XOR AND LTF:

Applying the distributive law to the AND XOR parts, where the ANDs have fan-in at most n^a and the XORs have fan-in $2^{O(n^{1+8 \log n})}$, each AND XOR parts can be converted into an XOR AND circuit of size $2^{O(n^{1+8+a \log n})}$, where the fan-in of ANDs is at most n^a . Letting $a = (ca)$ for sufficiently large $c > 1$, the fan-in of the new XORs is at most $2^{O(n^{1+a})}$. We now have a circuit C_4 of the form

SYM XOR AND LTF:

Note that the fan-in of the SYM gate is at most 2^{n^a} , and the fan-in of the (merged) ANDs is $O(n^{1+8+a})$.

5. Apply modulus-amplifying polynomials to eliminate the XOR layer. We'd like to remove the XOR layer, to further reduce the depth of the circuit. But as the gates of this layer have very high fan-in, we must be careful not to blow the circuit size up to (2^n) . The following construction will take advantage of the fact that we have only $\text{poly}(n)$ total gates in the bottom LTF layer.

We apply one step of Beigel-Tarui's transformation [BT94] (from ACC⁰ to SYM AND) to the SYM XOR AND part of our circuit. In particular, we apply a modulus-amplifying polynomial P (over the integers) of degree $2D^0 = 2n^a$ to each of the XOR AND parts. Construing the XOR AND as a sum of products $\sum_{P,Q} P_i Q_j$, the polynomial P has the property:

$$\begin{aligned} \text{If the } \sum_{P,Q} P_i Q_j &= 1 \pmod{2}, \text{ then } P\left(\sum_{P,Q} P_i Q_j\right) = 1 \pmod{2^{D^0}}. \\ \text{If the } \sum_{P,Q} P_i Q_j &= 0 \pmod{2}, \text{ then } P\left(\sum_{P,Q} P_i Q_j\right) = 0 \pmod{2^{D^0}}. \end{aligned}$$

So, composing P with each XOR AND part, each P outputs either 0 or 1 modulo 2^{n^a} . The key property here is that the modulus exceeds the fan-in of the SYM gate, so the sum of all $P\left(\sum_{P,Q} P_i Q_j\right)$ simply counts the number of XOR ANDs which are true; this is enough to determine the output of the SYM gate. Construing the output of each bottom LTF gate as a variable, there are at most 2^{n^2} variables (n^2 P,Q for each of the 2^n copies of the circuit from step 1 above). Expressing each $P\left(\sum_{P,Q} P_i Q_j\right)$

(expanded as a sum of products) as a multilinear polynomial in the ± 1 variables, the total number of terms is at most

$$2^{n^2} = 2^{O(n^2)} = 2^{O(n^{2a+1})} = 2^{O(n^{2a+1})}.$$

Let $n := \lfloor cn \rfloor$ for a sufficiently large constant $c > 1$ so that $2a^{n^2} + 1 = 8 + n^2$ (it suffices to pick any $c > 16(1 + 1/a)$). We can then merge the sum of a^{n^2} 's into the SYM gate, and obtain a SYM AND circuit where the SYM has fan-in

$$2^{O(n^{2a+1})} = 2^{O(n^2)};$$

and the AND gates have fan-in $O(n^{2a+1}) = O(n^2)$. The result is a circuit C_4 of the form

SYM AND LTF:

6. Replace the bottom threshold gates with DNFs (Theorem 8.5), and distribute. Note that the circuit C_4 has n^2 variables, so our SAT algorithm would follow if we could evaluate C_4 on all of its variable assignments in $2^{n^2} \text{poly}(n)$ time. We are now in a position to apply Lemma 8.5, which lets us reduce the evaluation problem for SYM AND LTF circuits to the evaluation problem for SYM AND SUM AND circuits, with a parameter K that needs setting. Recall the middle AND gates have fan-in $O(n^2)$, and the fan-in of the SUM is $O(\log K)$. Therefore by the distributive law, we can rewrite the circuit as a SYM SUM AND circuit, where each SUM gate has $(\log K)^{O(n^2)}$ ANDs below it, and at most one AND below each SUM is true. Thus we can wire these AND gates directly into the top SYM gate without changing the output.

In more detail, let $A; B = \{0, 1\}^{(n^2)=2}$, and set $N = 2^{(n^2)=2}$ and the integer parameter $K := 2^{bn^2}$ for a sufficiently large constant $b > 1$. By Lemma 8.5, we can reduce the SAT problem for SYM AND LTF circuits of size $2^{O(n^2)}$ on the set $A; B = \{0, 1\}^{n^2}$ to the SAT problem for SYM SUM AND circuits of size

$$2^{O(n^2)} = 2^{2bn^2} = 2^{O(n^2)}$$

on a prescribed set $A^0; B^0$ with $|A^0| = |B^0| = N$ and $A^0; B^0 = \{0, 1\}^{2bn^2 - n^2}$. By the distributive argument from the previous paragraph, we can convert the SYM SUM AND circuit into a SYM AND circuit of size at most

$$2^{O(n^2)} = 2^{O(n^2 \log \log K)} = 2^{O(n^2 \log(n))}.$$

By Lemma 8.5, we know that if the SYM AND SAT problem is solvable in time T on the set $A^0; B^0$, then the SAT problem for C_4 on the set $A; B$ can be solved in time $O(T + N^2 \sum_{Z=K}^N S) \text{poly}(n)$.

7. Evaluate the depth-two circuit on many pairs of points. By applying fast rectangular matrix multiplication in a now-standard way [Wil14c, Wil14b], the resulting SYM AND circuit of $2^{O(n^2)}$ size can be evaluated on all points $A^0; B^0$,

in time $\text{poly}(n) 2^{n^{\epsilon}}$, thus solving its SAT problem. Therefore, the SAT problem for C_4 can be solved in time

$$\text{poly}(n) 2^{n^{\epsilon}} + \frac{2^{n^{\epsilon}} 2^{O(n^{1-\epsilon})}}{2^{bn^{1-\epsilon}}} + 2^{\frac{n-n^{\epsilon}}{2}} 2^{O(n^{1-\epsilon} \log(n))}.$$

Setting $b > 1$ to be sufficiently large, we obtain a SAT algorithm for C_4 (and hence the original circuit C) running in $\text{poly}(n) 2^{n^{\epsilon}}$ time.

8.6.4 Satisfiability for Three Layers of Majority and AC^0

In this section, we give our SAT algorithm for MAJ AC^0 LTF AC^0 LTF circuits with low-polynomial fan-in at the output gate and the middle LTF layer:

Theorem 8.19. For all $\epsilon > 0$ and integers $d \geq 1$, there is a $\delta > 0$ and a randomized satisfiability algorithm for MAJ AC^0 LTF AC^0 LTF circuits of depth d running in $2^{n^{\epsilon + \delta}}$ time, on circuits with the following properties:

- the top MAJ gate, along with every LTF on the middle layer, has $O(n^{6-\epsilon})$ fan-in, and
- there are $O(2^n)$ many AND/OR gates (anywhere) and LTF gates at the bottom layer.

We need one more result concerning probabilistic polynomials over the integers:

Theorem 8.20 ([BRS91, Tar93]) For every AC^0 circuit C with n inputs and size s , there is a distribution of n -variate polynomials D over Z such that every p has degree $\text{poly}(\log s)$ (depending on the depth of C) and for all $x \in \{0, 1\}^n$, $\Pr_{p \in D} [C(x) = p(x)] \geq 1 - 2^{-\text{poly}(\log s)}$.

Proof of Theorem 8.19. The SAT algorithm is somewhat similar in structure to Theorem 8.17, but with a few important changes. Most notably, we work with probabilistic polynomials over Z instead of F_2 .

Start with a circuit C of the required form. Let s be the number of AND/OR gates in C plus the number of LTF gates on the bottom layer. Let $f = n^{6-\epsilon}$ be the maximum fan-in of the top MAJ gate and the LTFs on the middle layer, and recall that we're planning to consider C with size at most 2^n where $\epsilon > 0$ is a sufficiently small constant (depending on $\delta > 0$ and the circuit depth) in the following. Our SAT algorithm runs as follows:

1. By Theorem 8.16, every LTF of fan-in f can be replaced by an AC^0 MAJ of fan-in $f^{1+o(1)}$ and $\text{poly}(f)$ size. Hence we can reduce C to a circuit of similar size, but of the form

$$\text{MAJ } AC^0 \text{ MAJ } AC^0 \text{ MAJ}:$$

The fan-ins of the majority gates in the middle and bottom layer can be made at most $n^{6-\epsilon \delta}$, for any $\delta > 0$ which is smaller than ϵ . To be concrete, let us set $\epsilon \delta := \epsilon/2$.

- Replace the middle majority gates of fan-in $n^{6-5} = n^1$ with probabilistic polynomials (over Z) of degree $n^{3-5} = n^{-2}$ $\text{poly}(\log s)$ and error $1=2^{\text{poly}(\log s)}$ [AW15] (Theorem 7.4 from the previous Chapter). Replace all the C^0 subcircuits of sizes by probabilistic polynomials (over Z) of degree $\text{poly}(\log s)$ and error $1=2^{\text{poly}(\log s)}$, via Lemma 8.20. Note that the latter $\text{poly}(\log s)$ factor depends on the depth of the circuit.
- Replace the majority gate at the output (of fan-in $n^{6-5} = n^1$) with the probabilistic PTF of Corollary 7.4, setting the threshold parameter s^0 (which is called s in the statement of the corollary) to be 2^{2^n} and setting the error (called ϵ in the statement of the corollary) to be $1=2^{\text{poly}(\log s)}$. The resulting polynomial has degree $n^{2-5} = n^{-3} \text{poly}(n)$.

Applying the distributive law to all the polynomials from steps 2 and 3, the new circuit C^0 can be viewed as an integer sum of at most T AND-LTF circuits of at most T size, where

$$T = 2^{n^{3-5} = n^{-2} n^{2-5} = n^{-3} \text{poly}(\log s; n)} = 2^{n^{1-7} = 12} \text{poly}(\log s; n)$$

and all AND gates have fan-in at most $n^{1-7} = 12 \text{poly}(\log s; n)$ (because the resulting polynomial has at most this degree).

Now is a good time to mention our choice of ϵ , as it will considerably clean up the exponents in what follows. We will choose $\epsilon > 0$ to be sufficiently small so that the $\text{poly}(\log s; n)$ factor in the exponent of T is less than n^{12} . That is, we take $\epsilon := \epsilon = c$ and the size parameter $s < 2^n = 2^{n^c}$, for a sufficiently large constant $c \geq 12$ (Note that c depends on the depth of the circuit, since the degree of the $\text{poly}(\log s)$ factor depends on the depth.) Thus we have the size bound

$$T = 2^{n^{1-7} = 12 \text{poly}(\log s; n)} = O(2^{n^{1-7} = 12} n^{12}) = O(2^{n^{1-2}});$$

and all AND gates have fan-in at most n^{1-2} .

- For all assignments α to the first n variables of C^0 , plug α into C^0 , creating a copy C_a^0 . Let C^{00} be the integer sum of all 2^n circuits C_a^0 . By the properties of the polynomial constructed in Theorem 7.6 and the chosen parameter $\epsilon = 2^{-2^n}$, with probability at least 2^{-3} there is a (computable) threshold value $v = 3s = 2$ such that

$$\begin{aligned} C^{00}(x) > v & \text{ when at least one } C_a^0(x) \text{ outputs 1, and} \\ C^{00}(x) < v & \text{ when all } C_a^0(x) \text{ output 0.} \end{aligned}$$

The circuit C^{00} is a Sum-of-AND-LTF circuit; note that C^{00} has n variables.

- We now want to evaluate C^{00} on all of its 2^n possible variable assignments. Applying Lemma 8.5 for an integer parameter $K \geq 2$ (to be determined), $N = 2^{(n-1)K}$, and $Z; S = 2^{n^{1-2}}$, we can convert this evaluation problem for C^{00} into a corresponding evaluation problem for a Sum-of-AND-SUM-AND

circuit C^{000} on an appropriate combinatorial rectangle $A^0 \times B^0$ of 2^{n-n} variable assignments in total. The relative size of the circuit is unchanged, as each SUM-AND has size $O(\log^2 K) = O(n^2)$. The time for conversion of C^{000} into C^{000} is

$$\frac{N^2 Z^2}{K} + N \cdot S \cdot \text{poly}(n) = \frac{2^{n-n} \cdot 2^{2n^{1-\epsilon=2}} \cdot \text{poly}(n)}{K}.$$

Setting $K := 2^{2n^{1-\epsilon=2}}$ makes this time bound 2^{n-n} .

Recall that in the Sum-of-AND SUM-AND circuit C^{000} , the fan-in of the middle ANDs is at most $n^{1-\epsilon=2}$, and each SUM has $O(n)$ fan-in. We can therefore apply the distributive law to each AND-SUM part, and obtain a SUM-AND of size at most $n^{O(n^{1-\epsilon=2})}$. Merging the SUMs into the SYM gate, we obtain a SYM-AND circuit of size at most $n^{O(n^{1-\epsilon=2})}$.

6. Finally, applying rectangular matrix multiplication (Lemma 8.1) we can evaluate the Sum-of-AND C^{000} of $n^{O(n^{1-\epsilon=2})}$ size on the combinatorial rectangle $A^0 \times B^0$ in 2^{n-n} time, by preparing matrices of dimension $2^{n=2-n} \times n^{O(n^{1-\epsilon=2})}$ (for A^0) and $n^{O(n^{1-\epsilon=2})} \times 2^{n=2-n}$ (for B^0), then multiplying them. Note that preparing these matrices takes time no more than $2^{n=2+O(n^{1-\epsilon=2} \log n)}$, which is negligible for us.

After multiplying the matrices, we obtain a value for $C^{00}(x)$ for each assignment x , which is correct with probability at least 2^{-3} . By repeating steps 2-5 for $100n$ times, we obtain correct values on all 2^{n-n} points with high probability.

This completes the proof.

Part III

Probabilistic Rank and Matrix Rigidity

Chapter 9

Background and Overview

Let R be any commutative ring. The rank- r rigidity of a matrix $H \in \mathbb{R}^{N \times N}$, denoted $R_H(r)$, is the minimum Hamming distance between H and any matrix of rank at most r . Ever since Leslie Valiant introduced the notion of matrix rigidity [Val77], it has been a major challenge to construct interesting rigid matrices. Valiant and other complexity theorists have shown that explicit rigid matrices would yield new lower bounds for a number of different models of computation. The two most interesting rigidity parameter regimes for a family $\{M_N\}_{N \geq 2}$ of matrices, where M_N is a $N \times N$ matrix, are as follows:

A family $\{M_N\}_{N \geq 2}$ is called Valiant-rigid if there is a constant $\epsilon > 0$ such that

$$R_{M_N}(N - \log \log N) = \Omega(N^{1+\epsilon});$$

Valiant [Val77] showed that the linear transformations corresponding to Valiant-rigid matrices cannot be computed by $O(N)$ -size $O(\log N)$ -depth arithmetic circuits. There are currently no known lower bounds showing that such circuits cannot compute any explicit families of matrices.

A family $\{M_N\}_{N \geq 2}$ is called Razborov-rigid if there is any super-constant function $f(N) = \omega(1)$ such that

$$R_{M_N}(2^{(\log \log N)^{f(N)}}) = \Omega(N^2);$$

Razborov [Raz89] (see also [Wun12]) showed that if the communication matrix M_f of a Boolean function f is Razborov-rigid, then f is not in PH^{cc} , the communication analogue of the polynomial hierarchy. There are currently no explicit Boolean functions known to be outside PH^{cc} .

We say $\{M_N\}_{N \geq 2}$ is explicit if there is a deterministic algorithm which, on input N , outputs the matrix M_N in $\text{poly}(N)$ time. Aiming for a deterministic algorithm is important, since random matrices are known to be very rigid with high probability. Indeed, a random such matrix R_N has $R_{R_N}(r) = \Omega\left(\frac{(N-r)^2}{\log N}\right)$ for all r with high

probability¹ [Val77].

Despite decades of work and many known applications of rigid matrices, there has not been much success in actually constructing rigid matrices for almost any interesting rank parameter. There are essentially only three known deterministic constructions:

For all ranks r , there is a family of $N \times N$ matrices M_N constructible in P with $R_{M_N}(r) = \frac{N^2}{r} \log(N/r)$ [Fri93, SSS97]. This is proved via a combinatorial argument (untouched minor argument), and it is known that this type of approach cannot be further improved [Lok00].

A counting argument shows there is a family of $N \times N$ matrices R_N over a finite field F_q with $R_{R_N}(r) = O(N^2)$ for all $r = o(N)$. By combining a brute-force search for such rigid matrices with a padding argument (see Lemma 11.5 below), we can construct an $N \times N$ matrix L_N in $\text{TIME}[\exp(r^2)]$ with $R_{L_N}(r) = O(N^2)$.

Goldreich and Tal [GT16] show that random $N \times N$ Toeplitz matrices T_N over a finite field F_q have $R_{T_N}(r) = \frac{N^3}{r^2 \log N}$ for all $r = \frac{N}{\text{poly}(N)}$ with high probability. Their proof is primarily combinatorial and linear algebraic. Since random $N \times N$ Toeplitz matrices over F_2 are defined by $O(N)$ random bits, such rigid matrices can be constructed in E^{NP} .

Over large fields F , there are also approaches to constructing matrices which are rigid by virtue of having very large entries. For instance, an 'algebraic dimension' approach [SS96] can be used to construct rigid matrices over F with algebraically independent entries [Lok00, Lok06]. In this dissertation, we focus mainly on matrix rigidity over constant-size finite fields F_p , where such techniques cannot work.

In this Part of the dissertation, we make new progress on the problem of constructing rigid matrices by using new techniques which haven't yet been used in this area. First, in Chapter 10, we define a generalization of the polynomial method involving a new variant on the rank of a matrix called probabilistic rank, and use it to give a number of new rigidity upper bounds. We will give rigidity upper bounds for matrices which were previously conjectured to be very rigid, and give new connections between rigidity, circuit complexity, and communication complexity. Our probabilistic rank constructions use our probabilistic polynomial constructions from earlier in Chapter 7, together with the connection between sparse polynomials and low-rank matrices we explored in Chapter 8.

Second, in Chapter 11, we give a new construction of rigid matrices. Our construction makes use of ideas from circuit complexity theory which hadn't been used before in this context, and gives the first nontrivial construction of a family of Razborov-rigid matrices (although it is not an 'explicit' family of matrices as we discuss shortly). Interestingly, both our new rigidity upper bounds and lower bounds make use of the polynomial method: our new construction of rigid matrices critically uses a polynomial method algorithm for counting orthogonal vectors.

¹By comparison, one can see that $R_{R_N}(r) = O(N/r)^2$ for all r and all $N \times N$ matrices R_N .

9.1 Our Results

9.1.1 Probabilistic Rank and Matrix Rigidity

Let R be a commutative ring. In analogy with the notion of a probabilistic polynomial, we define a probabilistic matrix over R to be a distribution of matrices $M \in R^{n \times n}$. A probabilistic matrix M computes a matrix $A \in R^{n \times n}$ with error $\epsilon > 0$ if for every entry $(i; j) \in [n]^2$,

$$\Pr_{B \leftarrow M} [A[i; j] = B[i; j]] \geq 1 - \epsilon.$$

In this way, a probabilistic matrix is a worst-case randomized representation of a fixed matrix. A probabilistic matrix M has rank r if the maximum rank of a $M \leftarrow M$ is r .

We define the ϵ -probabilistic rank of a matrix $M \in R^{n \times n}$ to be the minimum rank of a probabilistic matrix computing M with error ϵ . Such probabilistic matrices are of interest and potentially very useful, because some full rank matrices can be represented by probabilistic matrices of rather low rank. For example, every identity matrix has ϵ -probabilistic rank $O(1/\epsilon)$ over any field, by simulating a protocol for EQUALITY using $\log(1/\epsilon) + O(1)$ communication that computes random inner products (cf. Theorem 10.5).

Probabilistic rank is related to the use of probabilistic polynomials in algorithm design. We saw in Chapter 8 how substituting low-degree probabilistic polynomials in place of common subroutines can be very useful for speeding up the best known running times for many core problems. All our algorithmic applications ended up embedding the low-degree polynomial evaluation problem in fast multiplication of two low-rank (rectangular) matrices (see Lemma 8.2 above). That is, this algorithmic work is really using the fact that various circuits and subroutines from core algorithms have low probabilistic rank and is applying low-rank representations to obtain an algorithmic speedup. Because low probabilistic rank is potentially a far broader notion than that of low-degree probabilistic polynomials, it makes more sense to study probabilistic rank directly, in the hopes of finding stronger algorithmic applications.

Probabilistic rank is also very related to matrix rigidity: it is not hard to see that probabilistic rank upper bounds imply rigidity upper bounds. In Chapter 10, we consider complexity-theoretic aspects of probabilistic rank. We demonstrate how probabilistic rank is a powerful notion for understanding matrix rigidity, and some models of communication complexity where knowledge is still sparse.

Hadamard Ain't So Rigid. Among the many attempts to prove arithmetic circuit lower bounds via rigidity, perhaps the most commonly studied explicit matrix has been the Walsh-Hadamard transform [PS88, Alo90, Gri, Nis, KR98, Cod00, Lok01, LTV03, Mid05, dW06b, Ras16]:

Definition 9.1. For vectors $x, y \in \mathbb{R}^d$, let $\langle x; y \rangle$ denote their inner product. Let $v_1; \dots; v_{2^n} \in \{0, 1\}^n$ be the enumeration of all n -bit vectors in lexicographical order. The Walsh-Hadamard matrix H_n is the $2^n \times 2^n$ matrix defined by $H_n(v_i; v_j) := (-1)^{\langle v_i; v_j \rangle}$.

It was believed that H_n is rigid because its rows are mutually orthogonal (i.e. H_n is Hadamard), so in several of the above references, only that property was assumed of the matrices. The best rigidity lower bounds known for H_n have the form $R_{H_n}(r) = \Omega(n \log n)$ for the target rank $r = O(2^n)$; in Valiant's problem, the lower bound is only $\Omega(n \log n)$. It was a folklore theorem that one can modify only $O(n)$ entries of an $n \times n$ Hadamard matrix and make its rank at most $2^n - 2$ [Lok14], but it was believed that for lower rank many more entries would require modification.

We give a good excuse for the weakness of these lower bounds:

Theorem 9.1 (Non-Rigidity of Hadamard Matrices). For every commutative ring R , for every sufficiently small $\epsilon > 0$, and for all n , we have $R_{H_n}(2^{n-\epsilon n}) = \Omega(n^{1+\epsilon})$ over R , for a function $f(\epsilon) = \Omega(\log(1/\epsilon))$.

In fact, we show a strong non-rigidity upper bound: by modifying at most $2^{\epsilon n}$ entries in each row of H_n , the rank of H_n drops to $2^{n-\epsilon n}$. That is, the matrix rigidity approach to arithmetic circuit lower bounds does not apply to Hadamard matrices such as the Walsh-Hadamard transform, since it is not Valiant-rigid. We would have required lower bounds of the form $R_{H_n}(2^{n-\epsilon n}) = \Omega(n^{1+\epsilon})$ for some $\epsilon > 0$ to obtain circuit lower bounds, but the upper bound of Theorem 9.1 shows this is impossible.

We do not (yet) believe that the Walsh-Hadamard transform has $O(2^n)$ -size $O(n)$ -depth circuits; a more appropriate conclusion is that rigidity is too coarse to adequately capture the lower bound problem in this case. Having said that, Theorem 9.1 does imply new circuit constructions: it follows that there is a depth-two unbounded fan-in arithmetic circuit for the Walsh-Hadamard transform with $2^{n+O(\epsilon \log(1/\epsilon))n} + 2^{2\epsilon n - \Omega(\epsilon^2 n)}$ gates; setting $\epsilon > 0$ appropriately, we have a $2^{\epsilon n}$ -size circuit for some $\epsilon < 1$.

We also show non-trivial rigidity upper bounds for H_n in the Razborov-rigidity regime that would be useful for communication complexity, where the rigidity is much closer to 4^n .

Theorem 9.2 (Non-Rigidity of Hadamard Matrices, Part II). For every integer $r \geq 2$, one can modify at most 2^{2n-r} entries of H_n and obtain a matrix of rank $(n \log(r))^{O(\frac{1}{n \log(r)})}$.

While the product of rank and rigidity (a natural measure) of H_n is only known to be at least (4^n) , Theorem 9.2 provides an upper bound of $4^n \cdot n^{O(\frac{1}{n \log(r)})} = r$. This is not small enough to refute the conjectured rigidity lower bounds required for communication complexity applications, but as we show later, these upper bounds still have non-trivial consequences for the communication complexity of IP2 (the Inner Product Modulo 2 function).

New Applications of Explicit Rigid Matrices.

Rigidity has been studied primarily for its connections to communication complexity and to lower bounds on arithmetic circuits computing linear transformations. We show new implications of constructing explicit rigid matrices for Boolean circuit complexity.

First, we show how explicit rigidity lower bounds would yield Boolean circuit lower bounds where only somewhat weak results are known:

Theorem 9.3. Let R be an arbitrary commutative ring, and M_n be a family of Boolean matrices such that (a) M_n is $n \times n$, (b) there is a $\text{poly}(\log n)$ time algorithm A such that $A(n; i; j)$ prints $M_n(i; j)$, and (c) there is a $\epsilon > 0$ such that for infinitely many n ,

$$\text{rank}_R M_n \leq 2^{(\log n)^{1-\epsilon}} \frac{n^2}{2^{(\log n)^2}} \text{ over } R:$$

Then the language $\{ \langle n; i; j \rangle \mid M_n(i; j) = 1 \}$ does not have AC^0 LTF AC^0 LTF circuits of $n^{2-\epsilon}$ -size and $(\log n = \log \log n)$ -depth, for all $\epsilon > 0$.

The theorem is obtained by giving non-trivial probabilistic rank bounds for such circuits, building on Lokam [Lok01]. Therefore, proving rigidity (or probabilistic rank) lower bounds for explicit 0/1 matrices over a commutative ring R would imply nearly-quadratic size lower bounds for AC^0 LTF AC^0 LTF circuits of unbounded depth, a powerful class of Boolean circuits. The best known lower bounds, which we proved above in Corollary 6.1 and Theorem 8.17, are that functions in the huge class E^{NP} do not have such circuits.

Sign-Rank Rigidity. The sign rank of a $n \times n$ matrix M is the lowest rank of a matrix N over R such that $\text{sign}(M[i; j]) = \text{sign}(N[i; j])$, for all $(i; j)$. Lower bounds on the sign-rank of matrices were used 15 years ago to prove exponential lower bounds against LTF MAJ and LTF SYM circuits [For02, FKL⁺01], i.e. restricted versions of depth-two threshold circuits. We extend the sign-rank connection to a circuit class for which strong lower bounds have long been open: explicit matrices with high rigidity under sign-rank would imply strong depth-two threshold circuit lower bounds. (Here, sign-rank rigidity is defined in the natural way, with rank replaced with sign-rank in the rigidity definition.)

A corollary of a theorem of Razborov and Sherstov [RS10] (see Theorem 10.13 below) is that for all n , H_n has sign-rank-rigidity at least $(4^n)^{\epsilon}$, just as in the case of normal rank rigidity. We show that even a somewhat minor improvement would already imply exponential-size lower bounds for depth-two linear threshold circuits with unbounded weights on both layers, a problem open for decades [HM93, KW16]:

Theorem 9.4. Suppose the sign rank-rigidity of H_n is $(4^n)^{\epsilon}$ for some rank bound $r \leq 2^n$ and some $\epsilon > 0$. Then the Inner Product Modulo 2 requires $2^{\Omega(n)}$ -size LTF LTF circuits.

Theorem 10.14 below gives a more general statement. Under the hood is an upper bound: matrices defined by small LTF LTF circuits have low probabilistic sign-rank for every such circuit of s gates, viewing its truth table as a $2^{n \times s} \times 2^{n \times s}$ matrix, there is a distribution of $O(s^2 n^2)$ -rank matrices which sign-represent the truth table in a worst-case probabilistic sense with error ϵ .

Rigidity, Communication, and Probabilistic Rank: An Equivalence. Probabilistic rank arises very naturally in studying generalized models of communication complexity. For a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let M_f be the $2^n \times 2^n$ truth table matrix of f with $M_f[x; y] = f(x; y)$ for all x, y . The following correspondence between probabilistic rank and communication complexity is immediate (one could even take the proposition as definition of BP-MOD_mP communication complexity).

Proposition 9.1. Let $m > 1$ be an integer, let $f : \{0, 1\}^n \rightarrow \{0, 1\}$, and let M_f be its truth table matrix. The BP-MOD_mP communication complexity of f with error ϵ equals the (base-2) logarithm of the probabilistic rank of M_f over Z_m (within additive constants).

Similarly, AM (Arthur-Merlin communication complexity) is equivalent to probabilistic Boolean rank.

It's easy to see that if a matrix has ϵ -probabilistic rank r , then its rank- r rigidity is at most ϵ^{2^n} ; thus rigidity lower bounds imply communication lower bounds. But conversely, it seems easier to prove lower bounds on probabilistic rank compared to rigidity: with probabilistic rank, we need to rule out a distribution of erroneous matrix entries which are required to spread the errors around; with rigidity, we have to rule out any adversarial choice of bad entries.

We show that for every randomly self-reducible function $f : \{0, 1\}^{2^n} \rightarrow R$ in which the self-reduction makes k non-adaptive queries, low rigidity implies low probabilistic rank: the ϵ -probabilistic rank of its corresponding matrix is at most $(kr)^k$ if its rank- r rigidity is at most ϵ^{4^n} . Thus there is a strong relationship between ϵ -probabilistic rank (and communication complexity, by Proposition 9.1) and the rank for which the rigidity is an ϵ -fraction of the matrix. For the Walsh-Hadamard transform, we prove that the probabilistic rank of H_n and the rigidity of H_n are equivalent concepts:

Theorem 9.5. For every commutative ring R and for every n, r , $R_{H_n}(r) \leq \epsilon^{4^n}$ over R if and only if H_n has ϵ -probabilistic rank r over R .

The matrices H_n represent the communication matrices of the widely-studied Inner Product Modulo 2 (IP2) function. By Proposition 9.1, the BP-MOD_pP communication complexity of IP2 and the rigidity of H_n over F_p are really equivalent concepts. Applying this theorem, our earlier rigidity upper bounds also imply some modest but interesting improvements on communication complexity protocols. From the rigidity upper bound of Theorem 10.2, we obtain a communication protocol for IP2 with $O\left(\frac{n}{p} \log\left(\frac{1}{\epsilon}\right) \log\left(\frac{n}{\log\left(\frac{1}{\epsilon}\right)}\right)\right)$ bits and error ϵ in the BP-MOD_pP communication model, for every prime p . (Aaronson and Wigderson gave an MA protocol for IP with $O\left(\frac{n}{p} \log\left(\frac{1}{\epsilon}\right)\right)$ communication complexity and error ϵ [AW09]; ours is more efficient for $\epsilon = 2^{-\log n}$.) Applying Theorem 10.1 yields an IP2 protocol with $n(1 - \epsilon^2 \log(1/\epsilon))$ communication and only $1 - 2^{-n}$ error. We are skeptical that our rigidity upper bounds for H_n are tight; we hope these results will aid future work (to prove rigidity upper bounds, one only has to think about communication protocols for IP2).

9.1.2 Efficient Construction of Rigid Matrices Using an NP Oracle

In Chapter 11, we give a new construction of rigid matrices. Unlike the previous constructions (which we enumerated near the beginning of this Chapter), which primarily use combinatorial and algebraic techniques, our construction primarily uses complexity-theoretic ideas. Our matrices are Razborov-rigid, and we can prove new lower bounds in communication complexity, Boolean circuit complexity, and arithmetic circuit complexity.

Constructions of Rigid Matrices in P^{NP}

Our main result is a construction of a rigid matrix in P^{NP} :

Theorem 9.6 (An Infinitely Often Rigid Matrix Construction in P^{NP}). There is an absolute constant $\epsilon > 0$ such for all prime powers $n = p^r$ and all constants $\epsilon > 0$:

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix H_N such that $R_{H_N}(2^{(\log N)^{1-\epsilon}}) \geq \epsilon N^2$ over F_q .

By comparison, applying previously known techniques to construct rigid matrices M_N for this rank $r = 2^{(\log N)^{1-\epsilon}}$, one either obtains:

M_N constructible in P with only $R_{M_N}(r) \leq \frac{N^2}{2^{(\log N)^{1-\epsilon}}}$, or

M_N only constructible in $\text{TIME}[\exp(\exp((\log N)^{1-\epsilon}))]$ with $R_{M_N}(r) \leq \epsilon N^2$. Note that the time bound here is larger than any quasi-polynomial in N , which can be written as $\exp(\exp(\log \log N))$.

Our construction in Theorem 9.6 is in P^{NP} , and achieves $R_{M_N}(r) \geq \epsilon N^2$.

It is natural to ask whether one can improve the constant ϵ in the rank in Theorem 9.6. We show an interesting win-win theorem: either the constant can be improved from ϵ to $1 - \epsilon$, or $\text{NQP} \subseteq P_{\text{poly}}$ follows.

Theorem 9.7 (Either a Better Construction in P^{NP} or $\text{NQP} \subseteq P_{\text{poly}}$). There is an absolute constant $\epsilon > 0$ such that for all prime powers $n = p^r$ and all constants $\epsilon > 0$, at least one of the following holds:

$\text{NQP} \subseteq P_{\text{poly}}$.

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix H_N such that $R_{H_N}(2^{(\log N)^{1-\epsilon}}) \geq \epsilon N^2$ over F_q .

Theorem 9.7 is interesting from the perspective of proving circuit lower bounds. Recall that a main motivation for constructing rigid matrices is to construct an explicit function which cannot be computed by $O(n)$ -size $O(\log n)$ -depth circuits [Val77]. If

we aim to show that NE (or E^{NP}) does not admit such circuits (which is still open), then we can safely assume $NE = P_{poly}$ before constructing the required rigid matrices. Therefore, if one could further improve the construction in the second bullet of the above Theorem 9.7 to match the Valiant-rigid parameters (which would require N matrices H_N with $R_{H_N}(N = \log \log N) = N^{1+\epsilon}$ for any $\epsilon > 0$, i.e. an improved rank parameter in exchange for a worsened rigidity parameter), it would imply that E^{NP} does not have $O(n)$ -size $O(\log n)$ -depth circuits.

Applications

Using the many different connections between rigid matrices and different areas of complexity theory, including a new connection we prove in Chapter 10, we derive from our construction a number of new lower bounds.

PH^{cc} Lower Bound for $NTIME[2^{(\log n)^{(1)}]}^{NP}$. A longstanding open problem in communication complexity is to prove a PH^{cc} (the communication complexity analogue of the polynomial hierarchy) lower bound for an explicit function [BFS86] (see [GPW18] for a recent reference). In fact, even for the much weaker subclass AM^{cc}, it is a notoriously open question to prove an $\Omega(\log n)$ lower bound for any explicit function [GPW16, CW19a]. Prior to this, it was even open whether $E^{NP} = AM^{cc}$, i.e., whether every function in E^{NP} has an efficient AM communication protocol.

Recall that Razborov-rigid matrices are known to yield lower bounds against PH^{cc}:

Lemma 9.1 ([Raz89], see also [Wun12]) Letting f be a function in PH^{cc}, the $2^n \times 2^n$ communication matrix M_f of f has $R_{M_f}(2^{(\log n)^c}) = 4^n$, where $\epsilon > 0$ is arbitrary and $c > 0$ is a constant depending only on ϵ , but not n .

Using this, our construction of rigid matrices in Theorem 9.6 immediately shows that $E^{NP} \not\subseteq PH^{cc}$, giving the first non-trivial lower bound against PH^{cc}. In fact, our rigidity bound is for a much higher rank than is necessary for applying Lemma 9.1 (setting $n = \log N$ in Theorem 9.6, we give a $2^n \times 2^n$ matrix M with $R_M(2^{n^{1-\epsilon}}) = 4^n$ for infinitely many n).

By a simple modification of our construction, we prove an even stronger lower bound:

Theorem 9.8. For all functions $f(n) = \Omega(1)$ such that $n^{(n)}$ is time-constructible, there is a function $f \in TIME[2^{(\log n)^{f(n)}}]^{NP}$ which is not in PH^{cc}.

Of the three previously-known deterministic constructions of rigid matrices mentioned near the beginning of this Chapter, only the second constructs rigid enough matrices to apply Lemma 9.1. However, it only yields a $2^n \times 2^n$ matrix M with $R_M(2^{(\log n)^{f(n)}}) = 4^n$ in $TIME[\exp(\exp((\log n)^{f(n)}))]$. We obtain exponential time savings using an NP oracle.

Depth-2 Arithmetic Circuit Lower Bounds

Although the rank parameters in our rigidity lower bounds from Theorem 9.6 are not high enough to give log-depth arithmetic circuit lower bounds via Valiant's approach, the rigidity parameters are

high enough that we can prove lower bounds against constant-depth arithmetic circuits. We consider a variant on rigidity which is useful for studying depth-2 arithmetic circuits:

Definition 9.2. For a field F and a matrix $A \in F^{N \times N}$, let

$$w_2(A) := \min \{ \text{nnz}(B) + \text{nnz}(C) \mid A = BC \};$$

where the min is over all pairs $B; C$ of matrices of any dimensions over F whose product is A , and $\text{nnz}(X)$ denotes the number of nonzero entries in the matrix X .

It is not hard to see that $w_2(A)$ equals, up to an additive ϵn , the minimum size (number of wires) of a depth-2 linear circuit over F which computes A , i.e. a depth-2 circuit which takes as input the N entries of a vector $x \in F^N$ and outputs the N entries of the vector Ax , and whose gates compute F -linear combinations of their inputs.

Every matrix $M \in \{0, 1\}^{N \times N}$ has $w_2(M) = O(N^2 \log N)$ over any field, and similar to the situation for rigidity, for any fixed prime power $q = p^r$, a random matrix $A \in F_q^{N \times N}$ has $w_2(A) = O(N^2 \log N)$ with high probability [Lup56]. However, the best known lower bounds on w_2 for explicit families of $N \times N$ matrices are only:

$(N \log N)$ for Boolean Hadamard matrices [AKW90]

$(N \log^2 N = (\log \log N)^2)$ for asymptotically good error-correcting codes [GHK⁺12]

$(N \log^2 N = \log \log N)$ for matrices based on super-concentrator graphs [RTS00]

Connections between rigidity lower bounds and w_2 lower bounds for a number of different parameter settings are known [Pud94]. We apply our rigidity lower bounds using a similar connection in the high rigidity setting to show high w_2 lower bounds for matrices constructible in P^{NP} :

Theorem 9.9. For all prime powers $q = p^r$ and constants $\epsilon > 0$:

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix $H_N \in \{0, 1\}^{N \times N}$ such that $w_2(H_N) = O(N^{2(\log N)^{1-\epsilon}})$ over F_q .

Threshold Circuit Lower Bounds. We next give an application of our construction to Boolean circuit complexity. Using the probabilistic rank upper bound we gave for threshold circuits in Theorem 9.3 above, we give a new lower bound:

Theorem 9.10. For every $\epsilon > 0$ and prime p , there is an $a > 0$ such that the class E^{NP} does not have non-uniform $AC^0[p]$ LTF $AC^0[p]$ LTF circuits of depth $o(\log n = \log \log n)$ where the bottom LTF layer has $2^{O(n^a)}$ gates, the rest of the circuit has polynomial size, and the middle layer LTF gates have fan-in $O(n^{1-\epsilon})$.

² $w_2(A)$ equals the minimum size of a depth-2 linear circuit for A when wires are not allowed to go directly from inputs to outputs. We can convert a circuit where wires do go from inputs to outputs to one where they do not, by adding in n middle-level gates which take the values of the inputs. Hence, the minimum size of a depth-2 linear circuit for A differs from $w_2(A)$ by a negligible additive ϵn .

We briefly compare with prior lower bounds for threshold circuits:

We showed above in Corollary 6.1 and Theorem 8.17 that E^{NP} does not have non-uniform ACC^0 LTF LTF circuits where the bottom LTF layer has 2^n gates and the remaining ACC^0 LTF subcircuit has $2^{n^{o(1)}}$ size. Tamaki [Tam16] also showed similar results for depth-2 circuits with symmetric and threshold gates. Our new lower bound is incomparable to these: we allow for many more LTF gates in the bottom layer, and unbounded depth, but the prior result allowed for larger size above the bottom layer, as well as ACC^0 circuitry rather than just $AC^0[p]$ circuitry.

Kane and Williams [KW16] previously showed there is a function in P which requires MAJ LTF LTF circuits of size $(n^{3-\epsilon} = \log^3 n)$. Our lower bound is for much larger circuits than this, but without a MAJ gate on top, and for a function in E^{NP} instead of P .

9.2 Other Related Work

Toggle Rank. By Yao's minimax principle [Yao83], BP $MOD_m P$ communication complexity (randomized communication with counting modulo m power) equals worst-case distributional $MOD_m P$ communication complexity. In matrix terms, putting an arbitrary distribution P on the pairs $\{0, 1\}^n \times \{0, 1\}^n$, the worst-case P -distributional complexity of M is the lowest rank (over Z_m) of a $2^n \times 2^n$ matrix N with error $\|M - N\|_P$ over P . Wunderlich [Wun12] calls this rank notion the approximate toggle rank. Proposition 9.1 shows that probabilistic rank and approximate toggle rank are very closely related, but they are not the same as the usual rigidity concept, which corresponds to the uniform distribution on pairs. For structured functions like IP_2 , we prove (in Theorem 9.5) that the uniform distribution is the worst case.

Sign-rank Rigidity and AC^0 - MOD_2 circuits. A tantalizing open problem that has gained popularity in recent years [SV12, ABG14, CGJ⁺16] is whether IP_2 has polynomial-size AC^0 MOD_2 circuits: i.e., circuits of $O(1)$ -depth over AND/OR/NOT, but with a layer of gates computing PARITY at the bottom nearest the inputs. Servedio and Viola [SV12] propose an interesting attack: in our terminology, they note that AC^0 MOD_2 circuits of size s have $n^{O(\log^d s \log(1/\epsilon))}$ sign-rank rigidity at most $\epsilon 2^n$ over R , and prove a lower bound on the correlation of signs of sparse polynomials (taken as a proxy for low-rank sign-matrices) with IP_2 . That is, they prove a weak sign-rank rigidity lower bound (note Razborov and Sherstov prove an analogous lower bound for sign-rank rigidity of IP_2 ; see Theorem 10.13). Our results have two consequences for this sort of approach. First, Theorem 10.12 shows that a sign-rank rigidity lower bound would prove something much stronger: a lower bound for depth-two threshold circuits computing IP_2 , a longstanding open problem. Second, our non-trivial upper bounds on the rank rigidity of the IP_2 matrix (which is H_n) suggest that IP_2 may have much lower sign-rank rigidity than expected.

Sign-Rank Rigidity and Margin Complexity. Linial and Shraibman [LS09] prove (in our terminology) that the sign-rank rigidity of an $n \times n$ matrix A is at most n^2 for target rank $O(\text{mc}(A)^2 \log(1/\epsilon))$, where $\text{mc}(A)$ is the margin complexity of A . Thus the margin complexity of a matrix can be used to upper bound sign-rank rigidity. They also study rigidity notion based on mc , conjecture that high mc implies high margin-complexity rigidity, and show that high margin-complexity also implies communication complexity lower bounds (for similar parameters as the standard rank-rigidity setting).

Approximate Rank. A different approximating rank notion has been studied in [BdW01, KS10, ALSV13], with connections to quantum computing and approximation algorithms. The ϵ -approximate rank of $M \in \mathbb{R}^{n \times n}$ is the lowest rank of a matrix A such that $\|M - A\|_1 \leq \epsilon$. That is, we can obtain one matrix from the other by perturbing each entry by at most $\epsilon/n > 0$. The appropriate analogy here seems to be that probabilistic polynomials are to probabilistic rank, as ϵ -approximate polynomials are to approximate rank: both are natural generalizations of polynomial representations to matrix representations, with different properties.

Non-rigidity of conjectured-to-be-rigid matrices. Are there other conjectured-to-be-rigid matrices which are not? One candidate would be the generating matrix of a good linear code over \mathbb{F}_2 . Recently, Goldreich [Dvi16] reported a distribution of matrices in which most of them are the generating matrix of a good linear code that is not rigid, found by Dvir. It would be very interesting to find an explicit code with this property. Since a preliminary version of our proof of the non-rigidity of the Walsh-Hadamard transform appeared [AW17], other researchers have extended our results to even more families of matrices: Dvir and Edelman [DE17] showed the non-rigidity of matrices arising from functions of inner products over finite fields, and Dvir and Liu [DL19] showed the non-rigidity of Fourier and circulant matrices.

Explicit Construction Based on Complexity-Theoretical Ideas. In a recent breakthrough work, Oliveira and Santhanam [OS17] gave an infinite deterministic construction of primes in sub-exponential time (That is, given an input 1^N , the (randomized) algorithm outputs a fixed prime P_N of N bits with high probability, for infinite number of N 's, and it runs in sub-exponential time). Their results are similar to our construction of rigid matrices in Theorem 9.6 in that they construct algebraic objects by building on complexity-theoretic ideas.

Our approach differs from theirs in several ways. First, [OS17] make crucial use of the fact that primes can be recognized in polynomial-time [AKS04], while in contrast, testing whether a matrix is rigid is coNP-complete (cf, Proposition 29 of [Des07]). Second, their results build on hardness vs randomness and a crucial component of their arguments is to use special pseudo-random generators to hit the set of all N -bit primes, while our results build on Williams' algorithmic approach to

lower bounds [Wil13, Wil14c]: we show one can contradict the non-deterministic time hierarchy theorem, assuming there is no P^{NP} construction of rigid matrices.

Conditional Explicit Construction of Rigid Matrices. There are several works achieving deterministic polynomial-time construction of rigid matrices under strong complexity assumptions. They are all based on the hardness-vs-randomness paradigm [NW94]. The observation is that since checking rigidity is in coNP , the ability to fool a non-deterministic algorithm implies the ability to construct rigid matrices.

In [KvM02], it is shown that under the assumption that E has no $2^{o(n)}$ -size SAT-oracle circuits, there is a P -time construction of matrices M_N over $\mathbb{Z}_{p(N)}$ such that $R(M_N)(r) \leq ((n-r)^2 = \log n)$, where $p(N)$ is a prime bounded by a polynomial of N . [MV05] give the same construction under the weaker assumption that E has no $2^{o(n)}$ -size non-deterministic circuits.³ In [GST03], the same construction is achieved with a uniform assumption that E has no $2^{o(n)}$ -time Arthur-Merlin protocols.

Lower Bounds on w_2 . Recall the variant on rigidity, w_2 , from Definition 9.2 above. Recently, Kumar and Volk [KV19] gave a construction of matrices with high w_2 , which is incomparable with our Theorem 9.9. Among other results, they show that there are constants $a, b, c > 0$ and a family $\{A_N\}_{N \in \mathbb{N}}$ such that A_N is an $N \times N$ matrix over an extension of \mathbb{F}_2 of degree $\exp(N^{1-a})$ which can be computed in time $\exp(N^{1-b})$ and with $w_2(A_N) > N^{1+c}$. By comparison, our Theorem 9.9 constructs $N \times N$ matrices H_N in P^{NP} with the worse lower bound $w_2(H_N) \leq (N \cdot 2^{(\log N)^{1-4\epsilon}})$, but our matrices are over \mathbb{F}_2 instead of a large extension field of \mathbb{F}_2 . Their techniques seem very different from ours, although they also use a padding trick, similar to our Lemma 11.5, of taking the Kronecker product of a rigid matrix with a large simple matrix to decrease its computational complexity in terms of the matrix size.

Circuit Lower Bounds via PCPP. In recent work, Chen and Williams [CW19b] applied PCPPs to show that, in order to prove C lower bounds for various non-deterministic time classes such as NEXP or NQP , it suffices to solve CAPP on $\frac{1}{2}$ C circuits (an XOR of two C circuits) in better-than- 2^n time. The proof crucially combines the assumed CAPP algorithm and PCPPs to obtain a non-trivial CAPP algorithm for general circuits. Here, our proof for Theorem 9.7 makes similar, but more sophisticated use of PCPPs. In particular, we actually require the PCPP to be smooth which is not required in [CW19b]. Our proof for Theorem 9.6 also relies on a completely different bootstrapping argument, which is specific for our task of constructing rigid matrices.

Rigidity and Data Structure Lower Bounds. Recent work by Dvir, Golovnev, and Weinstein [DGW19] showed connections between rigidity and static data structure lower bounds. In particular, they posed the challenge of constructing rigid ma-

³Indeed, the requirement is E has no $2^{o(n)}$ -size SV-nondeterministic circuits, which is the non-uniform analogue of $\text{NP} \setminus \text{coNP}$; see [MV05] for details.

trices in P^{NP} or E^{NP} as an avenue toward proving new data structure lower bounds. Unfortunately, the parameters of our new P^{NP} construction do not seem to yield any new data structure bounds using their approach.

9.3 Bibliographic Details

This Part of the dissertation is based on the results in two previously published papers:

‘Probabilistic Rank and Matrix Rigidity’ with Ryan Williams [[AW17](#)], which appeared in STOC 2017, and

‘Efficient Construction of Rigid Matrices Using an NP Oracle’ with Lijie Chen [[AC19](#)], which will appear in FOCS 2019.

Chapter [10](#) presents results from [[AW17](#)], and Chapter [11](#) presents results from [[AC19](#)].

Chapter 10

Probabilistic Rank and Matrix Rigidity

10.1 Polynomials, Rank, and Rigidity

In this Chapter, we study the interplay between the notions of probabilistic rank and matrix rigidity. We begin in this Section with their basic properties.

Definition 10.1. For any $\epsilon \in [0, 1]$, and any commutative ring R , a probabilistic matrix with error ϵ and rank r for the matrix $A \in R^{N \times N}$ is a distribution M on matrices $B \in R^{N \times N}$ of rank at most r over R such that, for every $i, j \in [N]$, we have

$$\Pr_{B \sim M} [A(i; j) \neq B(i; j)] \leq \epsilon:$$

The ϵ -probabilistic rank of A over R is the minimum rank of a probabilistic matrix with error ϵ for M .

Definition 10.2. For any commutative ring R , matrix $A \in R^{N \times N}$, and $r \in \mathbb{N}$, the rank- r rigidity of A , denoted by $R_A(r)$, is the minimum Hamming distance from A to an $N \times N$ matrix of rank r over R . That is, $R_A(r)$ is the number of entries of A that must be modified in order for the rank to drop to r . (The ring R should be clear from context.)

By drawing a 'typical' matrix from the probabilistic rank distribution, we can always obtain a matrix rigidity upper bound:

Proposition 10.1. For any commutative ring R , matrix $M \in R^{N \times N}$, and $\epsilon \in [0, 1]$, if the ϵ -probabilistic rank of R is r , then $R_A(r) \leq \epsilon N^2$.

We now describe a basic connection between probabilistic rank, matrix rigidity, and probabilistic polynomials (recall the definition from Definition 7.3 in Section 7.3 above). We show that probabilistic polynomials for a Boolean function can be used to give upper bounds on the probabilistic rank of the truth table matrix of f :

Definition 10.3. Let R be any commutative ring, and $f : \{0, 1\}^{2n} \rightarrow R$ be any function on $2n$ Boolean variables. The truth table matrix M_f of f is the $2^n \times 2^n$ matrix given by

$$M_f(v_i; v_j) = f(v_i; v_j);$$

where $v_1; \dots; v_{2^n} \in \{0, 1\}^{2n}$ is the enumeration of all $2n$ -bit vectors in lexicographical order.

Given the above definition, it is natural to define the probabilistic rank of a function:

Definition 10.4. The μ -probabilistic rank of a function $f : \{0, 1\}^{2n} \rightarrow R$ is the μ -probabilistic rank of its truth table matrix M_f . The rank of f and the rigidity of f are defined similarly.

We will make use of the following simple mapping from sparse polynomials to low-rank matrices; this is the same connection we used in Section 8.1 above to design algorithms using probabilistic polynomials.

Lemma 10.1. Let R be any commutative ring, and $f : \{0, 1\}^{2n} \rightarrow R$. Let $p : R^{2n} \rightarrow R$ be a polynomial with m monomials such that $p(x; y) = f(x; y)$ for any $x; y \in \{0, 1\}^{2n}$. Then the rank of f is at most m .

Proof. Let $a_1; \dots; a_m; b_1; \dots; b_m \in R^{2n} \rightarrow R$ be monomials such that $p(x; y) = \sum_{i=1}^m a_i(x) b_i(y)$ is the monomial expansion of p . For $1 \leq i \leq m$, define vectors $\tilde{a}_i; \tilde{b}_i \in R^{2n}$ by $\tilde{a}_i(x) = a_i(x)$ and $\tilde{b}_i(y) = b_i(y)$ for each $x; y \in \{0, 1\}^{2n}$. Then $M_f = \sum_{i=1}^m \tilde{a}_i \tilde{b}_i^T$, where $\tilde{a}_i \tilde{b}_i^T$ denotes the outer product of vectors. Thus $\text{rank}(M_f) \leq m$. \square

As a corollary, the probabilistic rank of f is at most the sparsity of a probabilistic polynomial for f :

Corollary 10.1. Let R be any commutative ring, and $f : \{0, 1\}^{2n} \rightarrow R$. If f has a probabilistic polynomial P with at most m monomials and error ϵ , then the μ -probabilistic rank of f is at most m .

Proof. Let p be a polynomial in the support of the distribution P . Since p has at most m monomials, by Lemma 10.1 the truth table matrix M_p of p (restricted to the domain $\{0, 1\}^{2n}$) has rank at most m . The distribution of M_p over p drawn from P is therefore an μ -probabilistic rank- m distribution for M_f , since $M_f(x; y) = M_p(x; y)$ if and only if $f(x; y) = p(x; y)$. \square

It follows that we can obtain a matrix rigidity upper bound from a sparse probabilistic polynomial.

Corollary 10.2. Let R be any commutative ring, and $f : \{0, 1\}^{2n} \rightarrow R$ be any function on $2n$ Boolean variables. If f has a probabilistic polynomial P with at most m monomials and error ϵ , then one can modify $\epsilon 2^{2n}$ entries of the truth table matrix M_f and obtain a matrix of rank at most m .

10.2 Non-Rigidity of Walsh-Hadamard

In this Section, we present our new rigidity upper bounds for the Walsh-Hadamard Transform.

Definition 10.5. Let $v_1; \dots; v_{2^n} \in \{0, 1\}^n$ be the enumeration of all-bit vectors in lexicographical order. The Walsh-Hadamard matrix H_n is the $2^n \times 2^n$ matrix defined by $H_n(v_i; v_j) := (-1)^{\langle v_i, v_j \rangle}$.

10.2.1 Rigidity Upper Bound for Low Error

We first prove that the Walsh-Hadamard matrices are not rigid enough for Valiant's program:

Theorem 10.1. For every field K , for every sufficiently small $\epsilon > 0$, and for all n , we have $R_{H_n} \leq 2^{n - \epsilon n} = 2^{n(1 - \epsilon)}$ over K , for a function f where $f(\epsilon) = (-\epsilon \log_2(1 - \epsilon))$.

We recall some notation from the Preliminaries. For a vector $v \in \{0, 1\}^n$, let $|v|$ be the number of ones in v . Let $H : [0, 1] \rightarrow [0, 1]$ denote the binary entropy function

$$H(p) = -p \log_2 p - (1 - p) \log_2(1 - p):$$

We also gave the following estimates on binomial coefficients (in Proposition 2.3 and Corollary 2.1). For $n \geq 0$, $0 \leq k \leq n$:

$$\binom{n}{k} \leq 2^{n H(k/n)}; \text{ and} \tag{10.1}$$

$$\sum_{k=0}^n \binom{n}{k} \leq 2^n; \tag{10.2}$$

Our first (simple) lemma uses a polynomial to compute a large fraction of H_n 's entries with a low-rank matrix. However, this fraction won't be high enough; we'll need another idea to correct many entries later.

Lemma 10.2. For every commutative ring R , and for every $n \geq 0$, there is a multilinear polynomial $p(x_1, \dots, x_n; y_1, \dots, y_n)$ over R with at most 2^{n+1} monomials, such that for all $x; y \in \{0, 1\}^n$ with $|x| \leq n/2$, $|y| \leq n/2$,

$$p(x; y) = (-1)^{\langle x, y \rangle}:$$

The proof uses properties of multivariate polynomial interpolation over the integers. To be concrete, we will apply the following Lemma from Chapter 7:

Reminder of Lemma 7.1 For any integers $n; r; k$ with $n \geq r + k$ and any integers c_1, \dots, c_r , there is a multivariate polynomial $p : \{0, 1\}^n \rightarrow \mathbb{Z}$ of degree $\leq r - 1$ with integer coefficients such that $p(z) = c_i$ for all $z \in \{0, 1\}^n$ with Hamming weight $|z| = k + i$.

Proof of Lemma 10.2. By Lemma 7.1 with $k = 2^n - 1$, $r = (1 - 2^{-\epsilon})n + 1$, and $c_i = (1 - 2^{-\epsilon})^{k+i}$, one can construct a multivariate polynomial $q : \{0, 1\}^n \rightarrow \mathbb{Z}$ with integer coefficients, of degree $(1 - 2^{-\epsilon})n$, such that for all $z \in \{0, 1\}^n$ with $\|z\|_1 \in [2^n - \epsilon n, (1 - 2^{-\epsilon})n]$, we have $q(z) = (1 - 2^{-\epsilon})^{\|z\|_1}$. As discussed in Section 2.1, q can be viewed as a polynomial over \mathbb{R} . Then our desired polynomial is

$$p(x_1, \dots, x_n; y_1, \dots, y_n) = q(x_1 y_1, x_2 y_2, \dots, x_n y_n) :$$

We can upper-bound the number of monomials in p as follows. First, since we only care about the value of p on $\{0, 1\}^{2n}$, we can make p multilinear by applying the equation $v^2 = v$ to all variables. Second, observe that for all $i = 1, \dots, n$, x_i and y_i appear in exactly the same monomials. So if we introduce a variable z_i in place of each $x_i y_i$ in p , the number of monomials in our new n -variate polynomial p^0 equals the number of monomials in p .

Since p^0 is multilinear and degree $(1 - 2^{-\epsilon})n + 1$, the number of monomials is at most $\sum_{i=0}^n \binom{n}{i} (1 - 2^{-\epsilon})^{n+1}$, which by (10.2) is at most $2^n c_2^{n+1}$ for some constant $c_2 > 0$. \square

Our second lemma says: fixing a vector x with about $(1 - 2^{-\epsilon})n$ ones, there is a strong upper bound on the number of vectors which has about $(1 - 2^{-\epsilon})n$ ones but has small (integer) inner product with x ; we'll use this to upper bound the number of erroneous entries at the very end.

Lemma 10.3. For every vector $x \in \{0, 1\}^n$ with $\|x\|_1 \in [(1 - 2^{-\epsilon})n, (1 - 2^{-\epsilon} + \epsilon)n]$, and any parameters $a, b \in (0, 1 - 5^{-\epsilon})$, the probability that a uniformly random vector y from $\{0, 1\}^n$ satisfies both

$$\|y\|_1 \in [(1 - 2^{-\epsilon} - a)n, (1 - 2^{-\epsilon} + a)n], \text{ and}$$

$$\left| \sum_{k=1}^n x_k y_k - bn \right| \leq bn$$

is at most $(2an + 1)(bn + 1) 2^{-\epsilon f(a; b) n}$, where f is a function such that $f(a; b) > 0$ as $a, b > 0$.

The usual toolbox of small-deviation estimates does not seem to yield the lemma; we give a direct proof.

Proof. For all x of the above form, every $k \in [(1 - 2^{-\epsilon} - a)n, (1 - 2^{-\epsilon} + a)n]$, and every $s \leq bn$, we count the number of $y \in \{0, 1\}^n$ with $\|y\|_1 = k$ and $\sum_{k=1}^n x_k y_k = s$. A vector y satisfies these properties if and only if:

there are exactly s integers i with $y[i] = 1$ and $x[i] = 1$, and

there are exactly $k - s$ integers i with $y[i] = 1$ and $x[i] = 0$.

So there are $\sum_{s=0}^k \binom{n}{s} x^s$ such choices of \mathbf{y} . The total probability is hence

$$\begin{aligned} & \frac{1}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \\ &= \frac{1}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \\ & \frac{1}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \\ & \frac{1}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \end{aligned}$$

Recall that if $k_1 < k_2 < n=2$ then $\binom{n}{k_1} < \binom{n}{k_2}$, and if $k_3 > k_4 > n=2$, then $\binom{n}{k_3} < \binom{n}{k_4}$. Step (*) therefore follows since $bn < \frac{1}{2}(1-a)n$ and $k_s \binom{n}{k_s} > \frac{1}{2}(1-a)n$ whenever $0 < a; b < 1=2$. Let $g(n) = (2an+1)(bn+1)$. Simplifying further, the above expression is at most

$$\begin{aligned} & \frac{g(n)}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \\ & \frac{g(n)}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \\ & \frac{g(n)}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \\ & \frac{g(n)}{2^n} \sum_{k=(1-a)n}^n \sum_{s=0}^k \binom{n}{s} x^s \sum_{j=0}^s \binom{s}{j} x^j \end{aligned}$$

where $f(a; b) = (1-a)(4b \log(1-2b) + (8a+4b) \log(1-(4a+2b)))$. □

Our third lemma is a simple linear-algebraic observation: given a low-rank matrix M that computes another matrix N on all but a small number of rows and columns, N must also have relatively low rank.

Lemma 10.4. Let M^0 be a matrix of rank r which is equal to M except in at most k columns and ℓ rows. Then the rank of M is at most $r + k + \ell$.

Proof. We will start with M^0 , and add at most $k + \ell$ rank-one matrices to M^0 so that it equals M .

Consider a column c on which M does not equal M^0 . We can add to M^0 a correction

matrix C_c given by

$$C_c(i; j) = \begin{cases} M(i; v) - M^0(i; v) & \text{if } j = v; \\ 0 & \text{otherwise.} \end{cases}$$

Then, $M^0 + C_c$ equals M on column c , and is unchanged in any other column. Moreover, since C_c is only nonzero on a single column, it has rank one. So all we have to do is add the correction matrix C_c for each column c on which M and M^0 differ. The rows of M^0 can be corrected analogously. \square

Corollary 10.3. Let T be any $2^n \times 2^n$ matrix. Let $a \in (0; 1/2)$, and let M be a $2^n \times 2^n$ matrix of rank r , indexed by n -bit vectors. There is a $2^n \times 2^n$ matrix M^0 of rank at most $r + 4 \cdot n \cdot 2^{n(1-a^2)}$ such that $M^0(v_i; v_j) = T(v_i; v_j)$ on all $v_i; v_j \in \{0; 1\}^n$ where at least one of the following holds:

$$|v_i| \geq [(1-2^{-a})n; (1-2^{-a} + a)n],$$

$$|v_j| \geq [(1-2^{-a})n; (1-2^{-a} + a)n], \text{ or,}$$

$$M(v_i; v_j) = T(v_i; v_j).$$

Proof. The number of $v_i \in \{0; 1\}^n$ with $|v_i| \geq [(1-2^{-a})n; (1-2^{-a} + a)n]$ is at most

$$\sum_{i=0}^{(1-2^{-a})n} \binom{n}{i} + \sum_{i=(1-2^{-a}+a)n}^n \binom{n}{i} = 2 \sum_{i=0}^{(1-2^{-a})n} \binom{n}{i} \leq 2^n \cdot 2^{n(1-a^2)};$$

by (10.2). Applying Lemma 10.4 to M and M^0 with k and ϵ set to $2^n \cdot 2^{n(1-a^2)}$, the result follows. \square

Let us outline how we'll use all of the above. First, we construct a matrix M of rank about $2^{n(1-a^2)}$ approximating H_n , using the polynomial from Lemma 10.2 in a straightforward way. This matrix M has far more erroneous entries than what we desire. But by Lemma 10.3, we can infer that the errors in M are highly concentrated on a relatively small number of rows and columns. Applying Corollary 10.3, the rows and columns can be corrected in a way that increases the rank of M by only $2^n \cdot 2^{n(1-a^2)}$. By Lemma 10.3, each row of the matrix left over will have $2^{O(n \log(1-a))}$ erroneous entries.

Proof of Theorem 10.1. In fact, we prove that one only has to modify $2^{O(n \log(1-a))}$ entries in each row of H_n , to obtain the desired rank.

Let $\epsilon > 0$ be given. By Lemma 10.2, there is a polynomial $p(x; y)$ in $2n$ variables with $m = 2^{n(1-a^2)}$ monomials which computes $(-1)^{x \cdot y}$ correctly, on all $(x; y) \in \{0; 1\}^{2n}$ such that $|x; y| \in [2^n; (1-2^{-a})n]$.

Construct a $2^n \times 2^n$ matrix M of rank m as in Corollary 10.1, so that $M(x; y) = p(x; y)$. By definition, M equals H_n on all $(x; y) \in \{0; 1\}^{2n}$ satisfying $|x; y| \in [2^n; (1-2^{-a})n]$.

Applying Corollary 10.3 to M with $T = H_n$ and $a = \epsilon$, we obtain a matrix M^0 of rank $m + 4 \leq 2^n - \binom{2^n}{\epsilon n}$ which is correct on all $(x; y)$ where either $|x_j| \geq [(1-\epsilon)n; (1+\epsilon)n]$, $|y_j| \geq [(1-\epsilon)n; (1+\epsilon)n]$, or $|x_i| \geq [2^n; (1+\epsilon)n]$.

Fix a row of H_n indexed by $x \in \{0, 1\}^n$ with $|x_j| \geq [(1-\epsilon)n; (1+\epsilon)n]$ (note the other rows are already correct). To show that M^0 differs from H_n on a small number of entries, we need to bound the number of y such that none of the above conditions hold, i.e.,

1. $|y_j| \geq [(1-\epsilon)n; (1+\epsilon)n]$ and
2. $|x_i| \geq [2^n; (1+\epsilon)n]$.

Note for our given x , it is never true that $|x_i| > (1+\epsilon)n$. Therefore we only need to bound the number N of y such that $|y_j| \geq [(1-\epsilon)a; (1+\epsilon)a]$ and yet $|x_i| < 2^n$. By Lemma 10.3 with $a = \epsilon n$ and $b = \epsilon n$, the probability that a random y satisfies $|x_i| < 2^n$ and $|y_j| \geq [(1-\epsilon)n; (1+\epsilon)n]$, is at most $O(n^2) \cdot 2^{-(\epsilon n)^2}$, where $\epsilon > 0$ as $n \rightarrow \infty$. Therefore $N \leq 2^n \cdot O(n^2) \cdot 2^{-(\epsilon n)^2} = O(n^2) \cdot 2^{-\epsilon^2 n}$.

Now for sufficiently large n and $\epsilon \in (0, 1/2)$, M^0 has rank at most $m + 4 \leq 2^n - \binom{2^n}{5\epsilon n} \leq 5\epsilon n \cdot 2^n - \binom{2^n}{5\epsilon n}$. Furthermore, on every row, M^0 differs from H_n in at most $n^2 \cdot 2^{-\epsilon^2 n} = 2^{O(-\epsilon^2 n)}$ entries. \square

10.2.2 Rigidity Upper Bound for High Error

In this section, we prove upper bounds on the rigidity of the Walsh-Hadamard transform in the regime where the error is constant, or much larger than ϵ ; this setting is of interest for communication complexity lower bounds.

Theorem 10.2. For every integer $r \leq 2^{2^n}$, over any commutative ring R , one can modify at most $2^{2^n} = r$ entries of H_n and obtain a matrix of rank at most $(n \log(r))^{O(n \log(r))}$.

The proof follows from applying our optimal-degree probabilistic polynomial for symmetric functions from Chapter 7:

Reminder of Theorem 7.3 There is a probabilistic polynomial over any commutative ring for any symmetric Boolean function on n variables, with error ϵ and degree $O(\frac{n}{\epsilon \log(1/\epsilon)})$.

Proof of Theorem 10.2. Set $\epsilon = 1/r$, and define the Boolean function $IP_2 : \{0, 1\}^{2^n} \rightarrow \{0, 1\}$ by $IP_2(x; y) = \prod_{i=1}^n x_i y_i$ for all $x, y \in \{0, 1\}^n$. We can see that H_n is the truth table matrix M_{IP_2} . By Corollary 10.2, it is sufficient to construct a probabilistic polynomial for IP_2 with error ϵ and $(n \log(1/\epsilon))^{O(n \log(1/\epsilon))}$ monomials. Consider the Boolean function $PARITY(z_1, \dots, z_n) = \prod_{i=1}^n z_i$ for all $z \in \{0, 1\}^n$, and note that $IP_2(x_1, \dots, x_n; y_1, \dots, y_n) = PARITY(x_1 y_1, x_2 y_2, \dots, x_n y_n)$. Since $PARITY$ is symmetric, by Theorem 7.3 it has a probabilistic polynomial P of error ϵ and degree $= O(\frac{n}{\epsilon \log(1/\epsilon)})$. Hence, the distribution of $p(x_1 y_1, \dots, x_n y_n)$ over p drawn from P is a probabilistic polynomial for IP_2 . Since we are only interested in the value of $p(z)$ when $z \in \{0, 1\}^n$, we can make p multilinear by applying the

equation $v^2 = v$ to all variables. Then the number of monomials of p is at most $\sum_{i=0}^n \binom{n}{i} = 2^n$. Since in $p(x_1, y_1, \dots, x_n, y_n)$ we are substituting in a monomial for each variable, its expansion has the same number of monomials as p , as desired.

10.2.3 Generalization To SYM-AND circuits

Here we generalize Theorems 10.1 and 10.2 to SYM AND circuits. In the proof of Theorem 10.1, the key property of the IP_2 function required is that it has the form

$$IP_2(x_1, \dots, x_n, y_1, \dots, y_n) = f(x_1 \wedge y_1, \dots, x_n \wedge y_n);$$

where f is a symmetric Boolean function (in our case, computes parity). The same proof yields the following generalization:

Theorem 10.3. For every symmetric function $f : \{0, 1\}^n \rightarrow \mathbb{R}$, define the function $IP_f : \{0, 1\}^{2n} \rightarrow \mathbb{R}$ by $IP_f(x; y) = f(x_1 \wedge y_1, \dots, x_n \wedge y_n)$. For all sufficiently small ϵ , there is a $\delta < 1$ and a matrix of rank 2^n which differs from the truth table matrix M_{IP_f} in at most $2^{(1+\delta)n}$ entries.

The proof of Theorem 10.2 only requires a probabilistic polynomial construction in Corollary 10.2. Our probabilistic matrix distribution simply substitutes monomials into the probabilistic polynomial of Theorem 7.3 for any symmetric function. Since each monomial can be viewed as an AND, the same argument will work for any SYM AND circuit.

Theorem 10.4. For any Boolean function $f : \{0, 1\}^{2n} \rightarrow \mathbb{R}$ which can be written as a SYM AND circuit with s AND gates, and for every integer $2 \leq r \leq 2^{2n}$, one can modify $2^{2n} - r$ entries of the truth table matrix M_f and obtain a matrix of rank at most $(s \log r)^{O(\frac{2n}{s \log r})}$.

10.3 Equivalence Between Probabilistic Rank and Rigidity

In this section, we show that the probabilistic rank of H_n and the rigidity of H_n are the same concept over fields. It is easy to see that if the probabilistic rank of H_n is k over a field K , then the rank- k rigidity of H_n is at most ϵ^{2^n} over K . Exploiting the random self-reducibility of the H_n function, we can show a converse: lower bounds on probabilistic rank imply proportionate rigidity lower bounds. This is of interest because probabilistic rank lower bounds appear to be fundamentally easier to prove than rigidity lower bounds.

Theorem 10.5. For every commutative ring R and for every n , $R_{H_n}(r) \leq \epsilon^{2^n}$ over R if and only if H_n has ϵ -probabilistic rank r over R .

First let us give some definitions. Let \otimes denote the outer product of vectors. For vectors $a \in \mathbb{K}^{2^n}$ whose entries are indexed by $y_1, \dots, y_{2^n} \in \{0, 1\}^n$, and $x \in \{0, 1\}^n$, let $a^{(x;y)}$ denote the vector in \mathbb{K}^{2^n} given by

$$a^{(x;y)}[v_i] = (-1)^{h_{v_i;y_i}} a[v_i \oplus x];$$

This permutes the entries of a , then negates half of the entries.

Proof. One direction is straightforward: low probabilistic rank implies low rigidity, by simply drawing a typical matrix from the distribution. For the other direction, suppose a_1, \dots, a_r and b_1, \dots, b_r are vectors in \mathbb{R}^{2^n} such that the $2^n \times 2^n$ matrix

$$M := \sum_{k=1}^r a_k \otimes b_k \tag{10.3}$$

differs from H_n in at most $\epsilon 2^{2n}$ entries. Pick vectors $x, y \in \{0, 1\}^n$ uniformly at random, and consider the $2^n \times 2^n$ matrix

$$M^0 = \sum_{k=1}^r (-1)^{h_{x;y_i}} a_k^{(x;y)} \otimes b_k^{(y;x)}. \tag{10.4}$$

In this form it is clear that M^0 has rank at most r . We claim that each entry of M^0 is equal to the corresponding entry of H_n with probability at least $1 - \epsilon$, over the choice of x and y , which will complete the proof.

Consider a given entry $M^0(v_i; v_j)$. It is sufficient to show that if $M(v_i \oplus x; v_j \oplus y) = H_n(v_i \oplus x; v_j \oplus y)$ then $M^0(v_i; v_j) = H_n(v_i; v_j)$, since $(v_i \oplus x; v_j \oplus y)$ is a uniformly random pair of vectors in $\{0, 1\}^n$. Suppose this is the case, meaning $M(v_i \oplus x; v_j \oplus y) = (-1)^{h_{v_i \oplus x; v_j \oplus y}} \sum_{k=1}^r a_k \otimes b_k$. Applying definition (10.3) and then (10.4) we see that

$$\begin{aligned} (-1)^{h_{v_i \oplus x; v_j \oplus y}} &= \sum_{k=1}^r a_k[v_i \oplus x] b_k[v_j \oplus y] \\ &= (-1)^{h_{v_i;y_i} + h_{v_j;x_j}} \sum_{k=1}^r (-1)^{h_{v_i;y_i}} a_k[v_i \oplus x] (-1)^{h_{v_j;x_j}} b_k[v_j \oplus y] \\ &= (-1)^{h_{v_i;y_i} + h_{v_j;x_j}} \sum_{k=1}^r a_k^{(x;y)}[v_i] b_k^{(y;x)}[v_j] \\ &= (-1)^{h_{v_i;y_i} + h_{v_j;x_j}} (-1)^{h_{x;y_i}} M^0(v_i; v_j); \end{aligned}$$

Rearranging, we see as desired that

$$M^0(v_i; v_j) = (-1)^{h_{v_i;x_j} + h_{v_j;y_i} + h_{v_i;y_i} + h_{v_j;x_j} + h_{x;y_i}} = (-1)^{h_{v_i;y_i}};$$

where the last step follows from the bilinearity of the inner product; \square

Therefore, proving communication lower bounds for the IP2 function against (for

example) the class $\text{BP} \text{ MOD}_m \text{P}$ is equivalent to proving rigidity lower bounds for H_n over the ring Z_m . Applying our rigidity upper bounds for H_n (Theorems 10.1 and 10.2), we obtain surprisingly low probabilistic rank bounds for H_n (and therefore communication-efficient protocols as well):

Corollary 10.4. For every commutative ring R , for every sufficiently small $\epsilon > 0$, and for all n , H_n has $(1-\epsilon)^{2^n}$ -probabilistic rank at most $2^{n \cdot (1-\epsilon)^{2^n}}$, and ϵ -probabilistic rank at most $(1-\epsilon)^{O(n \log n)}$, over R .

10.3.1 Generalization to Random Self-Reducible Functions

In fact, our reduction from rigidity to probabilistic rank works for any (non-adaptive) random self-reducible function [FF93] that makes a small number of oracle calls. Our notion of random self-reducibility is adapted for the communication complexity setting (for example, we do not care about the feasibility of the reduction).

Definition 10.6. A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is k -random self-reducible if there are random sampling procedures S_1, S_2 and a function $g : \{0, 1\}^k \rightarrow \{0, 1\}$ such that:

- (a) S_1 takes $x \in \{0, 1\}^n$ and a random string r and outputs $x_1, \dots, x_k \in \{0, 1\}^n$ such that for all n -bit strings z , $\Pr_r[x_i = z] = 1/2^n$ for all i ,
- (b) S_2 takes $y \in \{0, 1\}^n$ and a random string s and outputs $y_1, \dots, y_k \in \{0, 1\}^n$ such that for all n -bit strings z , $\Pr_s[y_i = z] = 1/2^n$ for all i , and
- (c) $f(x; y) = g(f(x_1; y_1), \dots, f(x_k; y_k))$.

Requirements (a) and (b) in the definition ensures that each x_i and y_i are uniform random variables; requirement (c) says that we can reconstruct $f(x; y)$ from the values $f(x_1; y_1), \dots, f(x_k; y_k)$.

Theorem 10.6. Let $r, n \in \mathbb{N}$ and $\epsilon \in (0, 1)$, and let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be k -random self-reducible. Suppose M_f has rank r rigidity at most $\epsilon 4^n$ over a commutative ring R . Then the (k^n) -probabilistic rank of M_f over R is at most $O\left(\frac{kr}{k}\right)$.

Proof. Suppose there is an $2^n \times r$ matrix A and $r \times 2^n$ matrix B , such that M_f and $A \cdot B$ differ in at most $\epsilon 4^n$ entries. We construct a distribution of low-rank matrices for M_f as follows.

Let $P(z_1, \dots, z_k)$ be the unique multilinear polynomial over R that represents the function g from the random self-reduction for f . Given k rows $X_1, \dots, X_k \in K^r$ of A , and k columns $Y_1, \dots, Y_k \in K^r$ of B , define a polynomial in $2kr$ variables:

$$Q(X_1, \dots, X_k; Y_1, \dots, Y_k) = P(\langle X_1, Y_1 \rangle, \dots, \langle X_k, Y_k \rangle)$$

Treating each term of the form $X_i[j] \cdot Y_i[j]$ as a variable, Q can be written as a sum of $\binom{kr}{k}$ total terms. Call the terms m_1, \dots, m_t , each of which are over $2kr$ variables.

Let r be a random string for S_1 and s be a random string for S_2 . For $x \in \{0, 1\}^n$, let $x_1, \dots, x_k \in \{0, 1\}^n$ be the outputs of $S_1(x)$ with randomness r . We define the x th row of a new $2^n \times t$ matrix A_r to be

$$m_1(A[x_1; :], \dots, A[x_k; :], \mathbf{1}; \dots; \mathbf{1}); \dots; m_t(A[x_1; :], \dots, A[x_k; :], \mathbf{1}; \dots; \mathbf{1}) :$$

For $y \in \{0, 1\}^n$, let y_1, \dots, y_k be the outputs of $S_2(x)$ with randomness s . Define the y th column of a new $t \times 2^n$ matrix B_s to be

$$m_1(\mathbf{1}; \dots; \mathbf{1}; B[:, y_1]; \dots; B[:, y_k]); \dots; m_t(\mathbf{1}; \dots; \mathbf{1}; B[:, y_1]; \dots; B[:, y_k])^T :$$

Then, for all $(x; y) \in \{0, 1\}^n \times \{0, 1\}^n$, the inner product of the x th row of A_r and the y th column of B_s is

$$\begin{aligned} & \sum_i m_i(A[x_1; :], \dots, A[x_k; :], B[:, y_1]; \dots; B[:, y_k]) \\ &= Q(A[x_1; :], \dots, A[x_k; :], B[:, y_1]; \dots; B[:, y_k]) \\ &= P(\mathbf{h}A[x_1; :], B[:, y_1] \mathbf{i}; \dots; \mathbf{h}A[x_k; :], B[:, y_k] \mathbf{i}) \end{aligned}$$

Since $(A \cdot B)$ differs from M_f on an ϵ -fraction of entries, for uniform random $x_i, y_j \in \{0, 1\}^n$ we have $\mathbf{h}A[x_i; :], B[:, y_j] \mathbf{i} \neq f(x_i, y_j)$ with probability at most ϵ . So with probability at least $1 - \epsilon$, $f(x_i, y_j) = \mathbf{h}A[x_i; :], B[:, y_j] \mathbf{i}$ for all $i = 1, \dots, k$. Thus the polynomial $P(\mathbf{h}A[x_1; :], B[:, y_1] \mathbf{i}; \dots; \mathbf{h}A[x_k; :], B[:, y_k] \mathbf{i})$ being implemented by $A_r \cdot B_s$ outputs $f(x; y)$ with probability at least $1 - \epsilon$. Hence all matrices $C_{r,s} = A_r \cdot B_s$ in our defined distribution have rank at most $O(\frac{k}{\epsilon})$, and for every $(x; y) \in \{0, 1\}^n \times \{0, 1\}^n$, $\Pr_{r,s}[C_{r,s}[x; y] = f(x; y)] \geq 1 - \epsilon$. \square

10.4 Explicit Rigid Matrices and Threshold Circuits

In this section, we show how explicit rigidity lower bounds would also imply circuit lower bounds where we currently only know weak results (e.g., we know that some functions in E^{NP} do not have such circuits).

Theorem 10.7. For every constant $\epsilon > 0$ and every AC^0 LTF AC^0 LTF circuit C of size $s = n^2$ and depth $d = o(\log(n) = \log \log(n))$, there exists a $\delta > 0$ such that the truth table of C as a $2^{n-2} \times 2^{n-2}$ matrix M_C has rigidity $R_{M_C} \geq 2^{n-1} \log(1/\delta) \geq \delta 2^n$, for all $\epsilon \geq (1/\delta)^2$, over any commutative ring.

Our proof will use a technique by Maciel and Thérien for converting each middle layer LTF gate into an equivalent AC^0 MAJ circuit:

Theorem 10.8 ([MT98] Theorem 3.3, [ACW16] Theorem 7.1) For every $\epsilon > 0$, every LTF on n inputs can be computed by a polynomial-size AC^0 MAJ circuit where the fan-in of each MAJ gate is $n^{1+\epsilon}$ and the circuit has depth $O(\log(1/\epsilon))$.

We will also use Tarui's probabilistic polynomial for AC^0 :

Theorem 10.9 ([Tar93] Theorem 3.6) Every circuit in AC^0 with depth d has a probabilistic polynomial over Z (and hence, any commutative ring) of degree $O(\log^d(n))$ and error $1=2^{\log^{O(1)}(n)}$.

Proof of Theorem 10.7. By Lemma 10.7, each LTF gate in the bottom layer has s -probabilistic rank $O(n^2s)$. We will design a probabilistic polynomial for the upper AC^0 LTF AC^0 circuitry, which will give the desired result when composed with this probabilistic rank expression.

First, each LTF gate in the middle layer has fan-in at most $s = n^2$. Applying Theorem 10.8 with $\epsilon = 1/2$ to each, the upper AC^0 LTF AC^0 circuit becomes a AC^0 MAJ AC^0 where each MAJ gate has fan-in at most $n^{(2)}(1+\epsilon) = n^{2+2}$, and the depth is still $O(d)$.

We can now apply the probabilistic polynomial for AC^0 from Theorem 10.9 with degree $O(\log^d(n))$ error $1=2^{\log^{O(1)}(n)}$ to the AC^0 circuits, and the probabilistic polynomial for symmetric functions on n^{2+2} bits from Theorem 7.3 with error ϵ and degree $O(n^{1+2+4} \log(s))$ to the MAJ gates in the middle. This results in a probabilistic polynomial of degree $O(n^{1+2+4} \log^{O(d)}(n))$. For $d = o(\log(n) = \log \log(n))$, this is $O(n^1)$ for any $\epsilon < 1/2$ ($2=4$).

We can view the terms in the probabilistic rank expression for the LTF gates in the bottom layer as variables that we substitute into this probabilistic polynomial; the number of monomials in this expansion will upper bound the rank, as in Lemma 10.1. Since there are at most s such gates, and each probabilistic rank expression has $O(n^2s)$ terms, we are substituting $O(n^2s^2)$ terms into our polynomial. Hence, the number of monomials will be upper bounded by

$$(n^2s^2)^{O(n^1)} = 2^{O(n^1) \log(n^2s^2)}$$

for any $\epsilon < 1/2$. This is of the desired form, where we can pick any positive value $\epsilon < 1/2$. The correctness follows by union bounding over all s probabilistic substitutions we make, each of which has error probability at most ϵ .

From the above theorem, setting the error ϵ appropriately, we infer a new consequence of explicit rigid matrices:

Theorem 10.10. Let R be any commutative ring, and M_n be a family of Boolean matrices such that (a) M_n is $n \times n$, (b) there is a $\text{poly}(\log n)$ time algorithm A such that $A(n; i; j)$ prints $M_n(i; j)$, and (c) there is a $\epsilon > 0$ such that for infinitely many n ,

$$R_{M_n} \geq 2^{(\log n)^{\epsilon}} \frac{n^2}{2^{(\log n)^{\epsilon}}} \text{ over } R:$$

Then the language $\{ (n; i; j) \mid M_n(i; j) = 1 \}$ does not have AC^0 LTF AC^0 LTF circuits of n^2 -size and $o(\log n = \log \log n)$ -depth, for all $\epsilon > 0$.

Therefore, proving strong rigidity lower bounds for explicit matrices has consequences for Boolean circuit complexity as well. Indeed, the desired circuit lower bounds could be derived from lower-bounding probabilistic rank.

10.5 Sign-Rank Rigidity and Depth-Two Threshold Circuits

Given a matrix $A \in \mathbb{R}^{n \times n}$, its sign rank is the minimum rank of any $B \in \{-1, 1\}^{n \times n}$ such that $\text{sign}(A[i; j]) = \text{sign}(B[i; j])$ for all $i, j \in [n]$. The ϵ -probabilistic sign-rank of A is defined analogously as with probabilistic rank. We say A has sign rank r -rigidity t if a minimum of t entries of A need to be modified in order for A to have sign rank at most r .

First, we observe that in the sign-rank setting, random $\{-1, 1\}$ matrices are still rigid: for example, with high probability, a random $\{-1, 1\}$ matrix has sign-rank $(n = \log^2 n)$ rigidity at least (n^2) . The proof follows readily from recent work:

Theorem 10.11 (Follows from Alon-Moran-Yehudayo [AMY16]). Let $r(n) = \alpha(n = \log n)$. For all sufficiently large n , a random $n \times n$ matrix with $\{-1, 1\}$ entries has sign-rank $r(n)$ rigidity at least (n^2) , with high probability.

Proof. There are 2^{n^2} different $n \times n$ matrices over $\{-1, 1\}$. The number of distinct matrices with sign rank at most r is bounded by $2^{O(r \log n)}$ [AMY16]. For a fixed $\{-1, 1\}$ matrix M , the number of $\{-1, 1\}$ matrices within Hamming distance t of M is at most $O(\binom{n^2}{t})$. Thus the number of matrices for which up to t entries can be changed to obtain a matrix of sign rank at most r , is upper-bounded by

$$2^{O(r \log n)} \binom{n^2}{t} = n^{O(r)} (en^2/t)^t.$$

Suppose we set $t = \epsilon n^2$. Then the above quantity is at most

$$n^{O(r)} (e/\epsilon)^{\epsilon n^2}.$$

For $r = \alpha(n = \log n)$ and $\epsilon \log_2(e/\epsilon) < 1$, a random matrix is not among these matrices with high probability. Therefore a random matrix has sign rank $(n = \log n)$ rigidity (n^2) with high probability. \square

Even though most $\{-1, 1\}$ matrices have high sign-rank rigidity, we show that the truth table of a small LTF-LTF circuit is always close to a matrix of low sign-rank. For even n , we say a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ has ϵ -probabilistic sign rank r if the truth table of C construed as an $2^{n/2} \times 2^{n/2}$ matrix has ϵ -probabilistic sign-rank r .

Theorem 10.12. For every function f with a LTF-LTF circuit of size s , and every $\epsilon > 0$, the ϵ -probabilistic sign-rank of f is $O(s^2 n^2 / \epsilon)$. Moreover, we can sample a low-rank matrix from the distribution of matrices in $2^{n/2} \text{poly}(s; n)$ time.

We will prove this theorem in a few steps. Let $EQ_n : \{-1, 1\}^{2n} \rightarrow \{-1, 1\}$ be the equality function, i.e., $EQ_n(x; x) = [x = y]$ (using Iverson bracket notation). Similarly, let $LEQ_n : \{-1, 1\}^{2n} \rightarrow \{-1, 1\}$ be the function $LEQ_n(x; x) = [x < y]$ where x and y are interpreted as integers in $\{-1, 1\}^n$.

Lemma 10.5. For every n , EQ_n has ϵ -probabilistic rank at most $O(1/\epsilon)$ over any commutative ring.

Proof. We mimic a well-known randomized communication protocol for EQ_n . Pick $k = \lceil \log_2(1/\epsilon) \rceil$ uniformly random subsets $S_1, \dots, S_k \subseteq \{0, 1\}^n$, and define the hash functions $h_1, \dots, h_k : \{0, 1\}^n \rightarrow \{0, 1\}$ by $h_i(x) = \prod_{j \in S_i} x_j$. Note that $h_i(x) \neq h_i(y)$ with $1/2$ chance if $x \neq y$. Hence, the following expression equals $\text{EQ}(x; y)$ with error probability at most ϵ :

$$\prod_{i=1}^k (h_i(x)h_i(y) + (1 - h_i(x))(1 - h_i(y))) \quad (10.5)$$

When expanded out, (10.5) is a sum of $2^k = O(1/\epsilon)$ terms of the form $f(x)g(y)$ for some functions f and g , each of which has rank one. \square

Lemma 10.6. For every n , LEQ_n has ϵ -probabilistic rank at most $O(n^2/\epsilon)$ over any commutative ring.

Proof. We express LEQ_n in terms of EQ predicates which check for the i th bit in which x and y differ, as

$$\text{LEQ}_n(x_1; \dots; x_n; y_1; \dots; y_n) = \prod_{i=1}^n (1 - x_i) y_i \text{EQ}_{i-1}(x_1; \dots; x_{i-1}; y_1; \dots; y_{i-1}) \quad (10.6)$$

We then get the desired rank bound by replacing each EQ with the probabilistic rank expression from Lemma 10.5 with error ϵ/n . By the union bound, all n of the EQ predicates will be correct with probability at least $1 - \epsilon$, and hence we will correctly compute LEQ_n . \square

Lemma 10.7. For every n , every linear threshold function $f : \{0, 1\}^{2n} \rightarrow \{0, 1\}$ has ϵ -probabilistic rank $O(n^2/\epsilon)$ over any commutative ring.

Proof. A linear threshold function f is defined as $f(x_1; \dots; x_n; y_1; \dots; y_n) = [\sum_{i=1}^n v_i x_i + \sum_{i=1}^n w_i y_i \geq k]$, where all v_i 's, w_i 's, and k are reals. We want to show that the $2^n \times 2^n$ matrix indexed by x_i -assignments on the rows and y_i -assignments on the columns has low probabilistic rank. We will exploit the fact that the linear forms on x_i 's and y_i 's can be preprocessed separately in a rank decomposition.

Define $a : \{0, 1\}^n \rightarrow \mathbb{R}$ by $a(x_1; \dots; x_n) = \sum_{i=1}^n v_i x_i$, and $b : \{0, 1\}^n \rightarrow \mathbb{R}$ by $b(y_1; \dots; y_n) = k - \sum_{i=1}^n w_i y_i$. Hence

$$f(x; y) = [a(x) \geq b(y)]$$

Let L be the list, sorted in increasing order, of all values $a(x)$ and $b(y)$, for all $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$. Then define the function $\sigma : \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$ where $\sigma(x)$ equals the earliest index i such that $a(x) \geq L_i$, interpreted as an $(i+1)$ bit

number. Define $f: \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$ similarly. Then

$$f(x; y) = \text{LEQ}_{n+1}(f(x); f(y)):$$

So the ϵ -probabilistic rank of M_f is at most that of $M_{\text{EQ}_{n+1}}$, which we upper-bounded in Lemma 10.6. \square

Now we are ready to upper-bound the probabilistic sign-rank of depth-two threshold circuits:

Proof of Theorem 10.12. We interpret our LTF circuit C as a function on two groups of $n/2$ bits, $x_1, \dots, x_{n/2}$, and $y_1, \dots, y_{n/2}$. Let $w_1, \dots, w_s; k \in \mathbb{R}$ be the weights of the output gate, so that

$$C(x_1, \dots, x_{n/2}; y_1, \dots, y_{n/2}) = \sum_{i=1}^s w_i f_i(x_1, \dots, x_{n/2}; y_1, \dots, y_{n/2}) + k;$$

for s different LTF functions f_i . By Lemma 10.7, the truth table matrix M_{f_i} of each f_i has ϵ -probabilistic rank $r = O(n^2/\epsilon)$. Our probabilistic distribution of matrices for M_C can be constructed as follows: for all $i = 1, \dots, s$, draw a random rank r matrix P_i from the distribution for M_{f_i} , and set

$$Q_C = (k \mathbf{J}) + \sum_i (w_i P_i);$$

where \mathbf{J} is the all-1s matrix. Q_C has rank at most $sr + 1 = O(n^2/\epsilon)$ and for all $(x; y)$, $\Pr[\text{sign}(Q_C[x; y]) \neq C(x; y)] \leq \epsilon$.

Are there explicit matrices with non-trivial sign-rank rigidity? We observe that the best-known rank rigidity lower bounds for H_n extend to sign-rank rigidity:

Theorem 10.13 (Follows from Razborov and Sherstov [RS10]) For all n , and $r \geq 2$, the sign-rank r rigidity of H_n is at least $(4^n - r)$.

Proof. Theorem 5.1 of [RS10] gives the following lower bound on sign-rank: given any matrix $A \in \{0, 1\}^{n \times n}$, suppose that all but h entries of matrix A have absolute value at least ϵ . Then

$$\text{sign-rank}(A) \geq \frac{n^2}{\|A\|_F + h};$$

where $\|A\|_F$ is the spectral norm of A . For the case of H_n , if we modify $h := 4^n - r$ entries arbitrarily, all but h entries have absolute value equal to 1. Thus

$$\text{sign-rank}(H_n) \geq \frac{4^n}{\|H_n\|_F + 4^n - r};$$

As $\|H_n\|_F = O(2^{3n/2})$ [For02], we have $\text{sign-rank}(H_n) \geq (4^n - (2^{3n/2} + 4^n - r)) = (2^{n/2} + r)$. \square

Can the above lower bound be improved slightly? Combining the previous two theorems, it follows that any minor improvement in the above rank/rigidity trade-off would begin to imply lower bounds for LTF-LTF:

Theorem 10.14. Suppose there is an $\epsilon > 0$ such that for infinitely many n , the sign rank r -rigidity of H_n is $(4^{n-\epsilon r^1})$, for some $r = \Omega(n^{2-\epsilon} s(n)^{2-\epsilon})$. Then the Inner Product Modulo 2 does not have LTF-LTF circuits of $s(n)$ gates.

Proof. Suppose the sign rank r -rigidity of H_m is $(4^{n-\epsilon r^1})$. Let $\epsilon = 1-\epsilon r^1$. It follows that the ϵ -probabilistic sign-rank of H_n is greater than r . But for a LTF-LTF function with s gates, its matrix always has ϵ -probabilistic rank $O(s^2 n^{2-\epsilon}) = O(s^2 n^{2-\epsilon r^1})$, by Theorem 10.12. Thus we have a contradiction when $O(s^2 n^{2-\epsilon r^1})$ is asymptotically less than r , i.e.,

$$r = \Omega(n^{2-\epsilon} s^{2-\epsilon});$$

corresponding to an s -gate lower bound against LTF-LTF circuits. Since H_n is just a linear translation of the matrix for Inner Product Modulo 2, the proof is complete. \square

For instance, proving the sign-rank 2^n -rigidity of H_n is at least $4^{n-2^{999n}}$ for some $\epsilon > 0$ would imply exponential-gate lower bounds for depth-two threshold circuits computing IP2.

10.6 Equivalence Between Probabilistic Rank Modulo m and BP-MOD m Communication Complexity

We conclude this Part by expanding on the connection between probabilistic rank and communication complexity. We sketch how probabilistic rank over \mathbb{Z}_m is equivalent to BP-MOD m communication complexity:

Proposition 10.2. Let $m > 1$ be an integer, let $f : \{0, 1\}^n \rightarrow \{0, 1\}$, and let M_f be its truth table matrix. Let $C_\epsilon(f)$ be the BP-MOD m communication complexity of f with error ϵ , and let ϵ -rank $_{\mathbb{Z}_m}(M_f)$ be the ϵ -probabilistic rank of M_f over \mathbb{Z}_m . Then $C_\epsilon(f) = \log_2(\epsilon$ -rank $_{\mathbb{Z}_m}(M_f) + 1) = 2C_\epsilon(f)$.

This proposition is different from the one quoted in the introduction (giving constant-factor equivalences between the log of the rank and the communication complexity) because we are assuming a more stringent communication model here. However, the more general model is often taken as the definition, in which case the probabilistic rank and communication complexity truly coincide.

First, given a distribution of low-rank matrices for M_f , it is easy to construct a protocol for f : Alice and Bob publicly randomly sample a matrix from the distribution, which is a product of two matrices A and B . Alice takes the row of A of length r corresponding to her input, Bob takes the column of B of length r corresponding

to his, and they then compute the inner product of these two vectors over \mathbb{Z}_m with $\log_2(r + 1)$ communication in the $\text{MOD}_m P$ model.

To construct a distribution of matrices from communication protocols, we do a simple modification of the BPP^P communication model. In fact, sometimes the literature defines the BPP^P communication model in this modified way [GPW16]. After the public randomness is chosen, Alice and Bob can, along with their nondeterministic bits, also sum over all possible transcripts of at most c bits between them. For each choice of randomness and nondeterminism there is a unique accepting transcript, so this extra choice does not alter the number of accepting communication patterns. But in this modified version, now Alice and Bob do not even have to communicate: they only have to send a single bit indicating whether they would accept or not, given the transcript and the nondeterminism. From such a protocol, it is straightforward to construct a $2^n \times 2^c$ matrix A representing Alice's protocol and a $2^c \times 2^n$ matrix B representing Bob, for any given string of public randomness.

Chapter 11

Efficient Construction of Rigid Matrices Using an NP Oracle

11.1 Construction Overview

In this Chapter, we give our new P^{NP} construction of rigid matrices:

Theorem 11.1 (An Infinitely Often Rigid Matrix Construction in P^{NP}). There is an absolute constant $\epsilon > 0$ such for all prime powers $q = p^r$ and all constants $\delta > 0$:

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix $H_N \in \mathbb{F}_q^{N \times N}$ such that $R_{H_N}(2^{(\log N)^{1-\delta}}) \geq \delta N^2$ over \mathbb{F}_q .

Along the way, we also give the following conditional construction which achieves better parameters:

Theorem 11.2 (Either a Better Construction in P^{NP} or $NQP \not\subseteq P_{=poly}$). There is an absolute constant $\epsilon > 0$ such that for all prime powers $q = p^r$ and all constants $\delta > 0$, at least one of the following holds:

$NQP \not\subseteq P_{=poly}$.

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix $H_N \in \mathbb{F}_q^{N \times N}$ such that $R_{H_N}(2^{(\log N)^{1-\delta}}) \geq \delta N^2$ over \mathbb{F}_q .

We begin, in this Section, with an overview of both of these constructions. In the rest of the Chapter we will give the formal proof and all the details, and then discuss some applications. For simplicity, we only consider the field \mathbb{F}_2 in this overview.

11.1.1 Either $NQP \not\subseteq P_{=poly}$ or a Construction of Rigid Matrices

We begin with a proof overview of Theorem 11.2. Note that Theorem 11.2 is equivalent to saying that there is a rigid matrix construction in P^{NP} under the assumption

$\text{NQP} = \text{P}_{=\text{poly}}$. Here we outline a conditional construction under the stronger assumption $\text{NE} \not\subseteq \text{P}_{=\text{poly}}$ for simplicity. We will then show how to get rid of the assumption using an additional bootstrapping argument.

Theorem 11.3 (Either a Better Construction in P^{NP} or $\text{NE} \not\subseteq \text{P}_{=\text{poly}}$). There is an absolute constant $\epsilon > 0$ such that for all constants $\delta > 0$, at least one of the following holds:

$\text{NE} \not\subseteq \text{P}_{=\text{poly}}$.

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix $H_N \in \mathbb{F}_q^{N \times N}$ such that $\text{rank}(H_N) \geq (2^{\log N})^{1-\delta} N^2$ over \mathbb{F}_q .

Low-Rank Matrices as a Circuit Class, and Corresponding Circuit Analysis Algorithms.

We begin with the observation that we can view low-rank matrices over \mathbb{F}_2 as a special type of 'circuit' defined by a pair of matrices. That is, supposing $M \in \mathbb{F}_2^{N \times N}$ is a matrix with rank r (think of $r \leq N$), then there are matrices $A \in \mathbb{F}_2^{N \times r}$ and $B \in \mathbb{F}_2^{r \times N}$ such that $M = AB$. Assuming N is a power of 2 for simplicity, M can be interpreted as (the truth-table of) a Boolean function $f: \{0, 1\}^N \rightarrow \{0, 1\}$, which has a special type of circuit of size $\Theta(N \cdot r)$ defined by A and B .

In this way, our task of constructing rigid matrices can equivalently be viewed as the task of proving a certain average-case lower bound against this special class of circuits. This is how Williams' algorithmic approach [Wil13, Wil14c], which exploits circuit analysis algorithms to prove such lower bounds, comes into play. When given the matrices $A; B$, the corresponding circuit analysis questions are:

1. Satisfiability (# SAT), which asks whether AB is the all zero matrix,
2. Derandomization (CAPP), which asks for an estimate of the probability that a random entry of AB is 1, and
3. Counting (# SAT), which asks for the exact number of ones in AB .

In fact, we observe that given the pair $(A; B)$, we can solve the hardest of these three problems, # SAT, in better-than- 2^n time (note $n = 2 \log N$). More formally, let a_i denote the i -th row of A , and let b_j denote the j -th column of B . The goal of # SAT is to count the number of pairs such that $a_i \cdot b_j = 0$ (the number of ones is N^2 minus the number of zeros). This is exactly an instance of Counting OV over \mathbb{F}_2 (\mathbb{F}_2 -# OV), with N vectors of r dimensions; compared to the usual OV problem, our inner product here is over \mathbb{F}_2 instead of \mathbb{Z} . An algorithm by Chan and Williams [CW16] solves this problem in deterministic $N^{2 \cdot (1 + \log(r = \log N))}$ time, for all $r \leq N^{\alpha(1)}$ (see Section 11.6 for the details). This algorithm will play a crucial part in our construction.

Williams' Algorithmic Approach to Circuit Lower Bounds, and a First Attempt. In a seminal work [Wil13], Williams demonstrated an algorithmic approach to proving circuit lower bounds. At a high level, the approach works as follows: Assuming a circuit lower bound is false, one combines the resulting small circuits with other algorithmic ideas to get a better-than- 2^n non-deterministic algorithm for $\text{NTIME}[2^n]$, therefore contradicting the non-deterministic time hierarchy theorem [Zák83].

A first attempt at using this approach in our situation proceeds as follows. Let L be a unary language in $\text{NTIME}[2^n] \setminus \text{NTIME}[2^{n-1}]$ [Zák83]. Fix an efficient PCP verifier V for L (such as [BV14]). That is, for a function $\ell := \ell(n) = n + O(\log n)$, $V(1^n)$ takes ℓ random inputs, runs in $\text{poly}(n)$ time, and is given access to an oracle $O : \{0, 1\}^\ell \rightarrow \{0, 1\}$ (O corresponds to the length- ℓ proof for V , but we will interpret it as an ℓ -bit Boolean function to help with intuition later on), and satisfies the following conditions:

1. (PCP Completeness) if $1^n \in L$, then there exists an oracle O such that $V(1^n)^O$ always accepts;
2. (PCP Soundness) if $1^n \notin L$, then for all possible oracles O , the probability $V(1^n)^O$ accepts is $\leq \epsilon$.

Intuitively, we are going to show that the truth table of the oracle O which makes V always accept (in the PCP Completeness case) has to be a rigid matrix. More precisely, letting $N = 2^{\ell} = 2^{n+O(\log n)}$, we can fix a P^{NP} machine M_{rigid} such that, on input 1^N , $M_{\text{rigid}}(1^N)$ outputs the lexicographically first oracle O_n which makes $V(1^n)$ always accept. M_{rigid} runs in P^{NP} (on input 1^N , which has length 2^{ℓ}), since it can guess the oracle outputs bit by bit, using its NP oracle to verify its guesses. The output of $M_{\text{rigid}}(1^N)$, and hence O_n itself, can be viewed as a matrix from $\{0, 1\}^{N \times N}$ which we want to show is rigid.

Assume toward a contradiction that $R_{M_{\text{rigid}}(1^N)}(r) \leq N^2$ for a small constant (one can think of $r := 2^{(\log N)^c}$ for a small constant $c > 0$) for all N . It follows that O_n can be $(1 - \epsilon)$ -approximated by a matrix of rank at most r . We can thus attempt to solve L as follows:

Given an input 1^n , we guess matrices $A \in \mathbb{F}_2^{N \times r}$ and $B \in \mathbb{F}_2^{r \times N}$ in $\Theta(r \cdot 2^{n/2})$ time, with the hope that $M := A \cdot B$ approximates O_n .

We estimate

$$p_{\text{acc}}(M) = \Pr_{z \in \{0, 1\}^\ell} [V(1^n)^M(z) = 1];$$

and accept only if $p_{\text{acc}}(M) \geq \epsilon$.

Following Williams' approach, the hope is that we can estimate $p_{\text{acc}}(M)$ in $2^{n/2}$ time (i.e. faster than iterating over all choices of the randomness) by taking advantage of the given low-rank approximation of M , combined with the # SAT algorithm for low rank matrices. If this were possible, it would put L in $\text{NTIME}[2^{n/2}]$, and contradict the non-deterministic time hierarchy theorem, completing our proof.

Two Issues with the First Attempt. Unfortunately, there are two main issues with this attempt. The first issue is that $V(1^n)^M(\cdot)$ can no longer be written as a low-rank matrix, even if M can. Ideally we would like $V(1^n)^M(\cdot)$ to be a low-rank matrix so that our #SAT algorithm applies to estimate $p_{\text{acc}}(M)$; without this condition, it's unclear how the low-rank matrix M is helpful. From [BV14], one can actually take $V(1^n)$ to be a 3-CNF, but this is still not enough, since a 3-CNF of low-rank matrices is not necessarily a low-rank matrix.

The second issue is more subtle. If we can estimate $p_{\text{acc}}(M)$ with a high enough accuracy, clearly we will always reject when $1^n \geq L$, by the soundness of the PCP. But, in order to accept when $1^n \leq L$, even if we have guessed a M which $(1 - \epsilon)$ -approximates O_n , it still could be the case that $p_{\text{acc}}(M)$ is small. For instance, what if $V(1^n)$ always queries positions on which M and O_n differ?

We will ultimately resolve the second issue by making the verifier smooth (meaning each query is uniformly distributed), which we will explain later. To resolve the first issue, we use a recent idea from Chen and Williams [CW19b], together with easy-witness lemmas [IKW02, MW18].

The Easy Witness Lemma. Assuming $NE = P_{\text{poly}}$, by [IKW02], we know that all NE verifiers have polynomial-size witness circuits, including the verifier $V(1^n)$ discussed above. In other words, when $1^n \leq L$, before we were only able to assume there is an oracle $O : \{0, 1\}^n \rightarrow \{0, 1\}$ such that $V(1^n)^O$ always accepts, but now we can further assume that there is such an oracle which is computed by a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ of size k for a constant k . Let us set C_{best} to be the lexicographically first circuit having this property. Now we can modify our algorithm from the first attempt by guessing C , and trying to estimate $p_{\text{acc}}(C)$ instead. Notice that with this modification, there are no longer any low-rank matrices involved in our current approach. We will instead use low-rank approximations of the proof for a different PCP, which we describe next.

Smooth PCP of Proximity (PCPP). We are now going to make use of a very recent construction of a smooth PCPP [Par19]. Using PCPPs in conjunction with Williams' algorithmic approach to circuit lower bounds in this way was a key idea from [CW19b]. For a polynomial-size circuit $F : \{0, 1\}^n \rightarrow \{0, 1\}$ (we are eventually going to pick F to be a modification of the circuit C from above), a smooth PCPP verifier $V_{C\text{-EVAL}}(F)$ for F ¹ takes as input a proof π of length $\text{poly}(n)$ and $O(\log n)$ random bits, and makes a constant number of uniformly distributed, non-adaptive queries to the proof and the input (i.e. which bits are queried depend only on the random bits, and each bit has an equal probability of being queried). Moreover, for some small constant $\rho > 0$:

(PCPP Completeness) If $F(\cdot) = 1$, then there is a proof π such that $V_{C\text{-EVAL}}(F)$ always accepts. Moreover, there is a polynomial-time algorithm which computes π given F and \cdot .

¹More precisely, $V_{C\text{-EVAL}}(F)$ is a smooth PCPP for the Circuit-Eval problem, in which we have fixed the circuit to be F .

(PCPP Soundness) If $V_C(F) = 0$ for every $F \in \{0, 1\}^n$ which differs from f in at most a ρ fraction of entries, then $\Pr_{u \in \{0, 1\}^{O(\log n)}} [V_{C\text{-EVAL}}(F)(u) = 1] \leq 3^{-1}$ for all possible proofs $\pi \in \{0, 1\}^{\text{proof}}$.

The fact that the queries are both smooth and non-adaptive will be crucial to our construction later on. This 'proximity' aspect of the soundness condition is necessary for these properties to hold. For instance, if one tried to be the parity function, then it is not hard to see that such a construction without the 'proximity' aspect (i.e. with $\rho = 0$) is impossible. Our goal is to apply such a smooth PCPP to C , but since we don't have any guarantees about which input w should reject, we will first need to make some modifications to C to deal with the 'proximity' aspect.

Defining $D_C(w) := V(1^n)^C(w)$, which is still a polynomial-size circuit, our goal is to design a fast algorithm to estimate

$$p_{\text{acc}}(C) := \Pr_{w \in \{0, 1\}^n} [V(1^n)^C(w) = 1] = \Pr_{w \in \{0, 1\}^n} [D_C(w) = 1]:$$

In preparation for using the PCP of proximity, we next apply an error correcting code to C . Specifically, let ECC be a constant-rate F_2 -linear error correcting code with efficient encoder $\text{Enc}: \{0, 1\}^n \rightarrow \{0, 1\}^{c_1}$ and decoder $\text{Dec}: \{0, 1\}^{c_1} \rightarrow \{0, 1\}^n$ which can recover error up to a δ_{dec} fraction. We define another circuit $E_C: \{0, 1\}^{c_1} \rightarrow \{0, 1\}^n$, as $E_C(w) := D_C(\text{Dec}(w))$. That is, E_C treats the input as a codeword of ECC , decodes it, and feeds the result into the circuit D_C . Now our goal is to estimate

$$p_{\text{acc}}(C) = \Pr_{w \in \{0, 1\}^{c_1}} [E_C(w) = 1]:$$

Now we will apply the PCP of Proximity to simplify the estimation of $p_{\text{acc}}(C)$. More precisely, we use a $q = O(1)$ query smooth PCPP, $V_{C\text{-EVAL}}(E_C)$, for the circuit E_C , which has proximity parameter $\leq \delta_{\text{dec}}$, proof length $|\pi_{\text{proof}}| = \text{poly}(\text{SIZE}(E_C)) = \text{poly}(n)$, and number of random bits $m = O(\log |\pi_{\text{proof}}|) = O(\log n)$. The crucial observation here is that we have dealt with the 'proximity' aspect of the smooth PCPP by using the error correcting code: if $D_C(w) = 0$, then $\text{Enc}(w)$ is δ_{dec} -far from any yes-inputs to E_C . This is because, for any $w' \in \{0, 1\}^{c_1}$ which is δ_{dec} -close to $\text{Enc}(w)$, w' decodes to w and $E_C(w') = D_C(w) = 0$.

Summarizing, so far we have the following:

(PCPP Completeness) If $D_C(w) = 1$, then there is a proof $\pi \in \{0, 1\}^{\text{proof}}$ such that $V_{C\text{-EVAL}}(E_C)^{\text{Enc}(w), \pi}$ always accepts. Moreover, given E_C and π , there is a polynomial-time computable function $(E_C; \pi) \in \{0, 1\}^{\text{proof}}$ to compute the proof π .

(PCPP Soundness) If $D_C(w) = 0$, then $\Pr_{u \in \{0, 1\}^m} [V_{C\text{-EVAL}}(E_C)^{\text{Enc}(w), \pi}(u) = 1] \leq 3^{-1}$ for all possible proofs $\pi \in \{0, 1\}^{\text{proof}}$.

The P^{NP} Machine M_{rigid} . Finally, we are ready to define our rigid matrix. It will be the concatenation, over all $w \in \{0, 1\}^n$, of the proof $(E_{C_{\text{best}}}; \pi)$ from the

PCPP Completeness condition above. More precisely, let $c_{\text{best}}(j)$ be the j -th bit of $(E_{C_{\text{best}}})$. Note that C_{best} is a Boolean function on $n := n + O(\log n)$ bits. Letting $N = 2^n$, we define our P^{NP} machine M_{rigid} as the function which, on input 1^N , outputs the truth-table of C_{best} , which we interpret as a matrix in $\{0, 1\}^{N \times N}$. M_{rigid} runs in P^{NP} since, similar to before, one can guess C_{best} bit-by-bit and verify with the NP oracle.

Again, assume toward a contradiction that $R_{M_{\text{rigid}}(1^N)}(r) = N^2$ for a small constant r (recall that one can think of $r := 2^{(\log N)^{1-\epsilon}}$ for a small constant $\epsilon > 0$) for all N . That is, we know $C_{\text{best}}(j)$ can be $(1 - \epsilon)$ -approximated by a matrix M of rank at most r . We guess a low-rank decomposition of that matrix $M = AB$, in $O(N^r)$ time, and now we wish to estimate

$$p_{\text{acc}}(M) := \Pr_{u \in \{0, 1\}^m; v \in \{0, 1\}^n} [V_{C\text{-EVAL}}(E_C)^{\text{Enc}(u)} M(v) = 1]:$$

Recall that v is the randomness to the old PCP verifier V , of length $n = n + O(\log n)$, and u is the randomness to the new smooth PCPP verifier $V_{C\text{-EVAL}}(E_C)$, of length $m = O(\log n)$.

Fast Algorithm for Computing $p_{\text{acc}}(M)$. We now use the fact that the queries made by $V_{C\text{-EVAL}}(E_C)$ only depend on u . Our algorithm will simply iterate over all poly(n) choices of u . Hence, $x \in \{0, 1\}^m$, and suppose V queries $M(j_1); M(j_2); \dots; M(j_{q_1})$ in $M(j;)$, and $e_1; e_2; \dots; e_{q_2}$ in $\text{Enc}(u)$. Now we want to estimate

$$\Pr_{u \in \{0, 1\}^m} [F_u(M(j_1); M(j_2); \dots; M(j_{q_1}); \text{Enc}(u)_{e_1}; \text{Enc}(u)_{e_2}; \dots; \text{Enc}(u)_{e_{q_2}}) = 1]$$

for a Boolean function F_u on $q = q_1 + q_2 = O(1)$ inputs. Next, using a standard trick from the analysis of Boolean functions, we observe that since we are aiming to compute the expected value of F_u , we can assume that F_u is a parity function. In other words, it is sufficient to quickly estimate

$$\Pr_{u \in \{0, 1\}^m} [M(j_1) + M(j_2) + \dots + M(j_{q_1}) + \text{Enc}(u)_{e_1} + \text{Enc}(u)_{e_2} + \dots + \text{Enc}(u)_{e_{q_2}} = 1];$$

where the sum is taken mod 2. The parity of $M(j_1) + M(j_2) + \dots + M(j_{q_1})$, which is a sum of a constant number of low-rank matrices, can itself be written as a low rank matrix. Since Enc is a linear function over F_2 , incorporating $\text{Enc}(u)_{e_1} + \text{Enc}(u)_{e_2} + \dots + \text{Enc}(u)_{e_{q_2}}$, which is a linear function of the indices of the matrix, can only increase the rank by an additive constant. Hence, our goal is exactly to compute the number of 1s in a low rank matrix. This is an instance of the previously discussed SAT problem for low-rank matrices which, as discussed, can be solved in $N^{2(1-\epsilon \log r)}$ time as described by [CW16].

Notice that:

If $1^n \leq L$, and we guessed the circuit C_{best} and a matrix M which $(1 - \epsilon)$ -approximates C_{best} , then $p_{\text{acc}}(M) \geq 1 - \epsilon$ since $V_{C\text{-EVAL}}(E_C)$'s queries are

smooth (meaning, uniformly distributed over the proof).

If $1^n \geq L$, then for all possible guesses $p_{acc}(M) = 1/2$, by the soundness of PCPP and PCP.

Putting everything together, it follows that L is in non-deterministic time

$$\text{poly}(n) \cdot N^{2 \cdot (1/\epsilon \log r)} = 2^{n \cdot (1/\epsilon \log r)} = 2^{n \cdot (1/\epsilon)};$$

contradicting the non-deterministic time hierarchy. This completes the proof overview for Theorem 11.3.

11.1.2 Unconditional Construction of Rigid Matrices

Getting Rid of the Easy-Witness Assumption: A Boot-Strapping Scheme.

We now move on to a proof overview of Theorem 11.1. Note that in the above argument, the only consequence of $\text{P} = \text{poly}$ used is the fact that $V(1^n)$ has a succinct witness circuit. In order to get rid of the assumption $\text{P} = \text{poly}$, we next show how to construct a succinct witness for $V(1^n)$ solely based on the assumption that all P^{NP} machines have non-rigid output matrices.

The key idea is based on a bootstrapping argument. Observe that an $N^{o(1)}$ -rank decomposition of a matrix $M \in \{0, 1\}^{N \times N}$ actually compresses the N^2 bits of M into an $N^{1+o(1)}$ bit representation. If we can further treat those bits after the compression as a low-rank matrix, and compress it again, and so on, we can further reduce the number of bits required to represent the matrix.

A key property of low-rank decompositions we will use is that they are locally decodable. That is, if $A; B$ are the two matrices of a rank- r expression for M , then one can compute a particular entry M_{ij} , by looking at only $O(r)$ entries of the matrices A and B (the i th row of A and the j th column of B).

High-Level Idea. Recall from the proof above that O_n is the lexicographically first oracle which makes $V(1^n)$ always accept. The high level idea for constructing a succinct witness for O_n is as follows. We first interpret O_n as a matrix $M_1 \in \{0, 1\}^{N_1 \times N_1}$. Letting $(A_1; B_1)$ be its low-rank decomposition, we then interpret the concatenation $(A_1; B_1)$ as a matrix $M_2 \in \{0, 1\}^{N_2 \times N_2}$. We will show that M_2 also has a low-rank decomposition $(A_2; B_2)$. We then interpret this as a matrix $M_3 \in \{0, 1\}^{N_3 \times N_3}$, and repeat until we have a small enough matrix $M_k \in \{0, 1\}^{N_k \times N_k}$. Note that for all i , we have $N_i = N_{i-1}^{1+2^{-i}}$; that is, each time we compress the bits by about a square-root.

Why do all these matrices have low-rank approximations? This follows from our assumption that all P^{NP} machines' output matrices are non-rigid, and hence have low-rank approximations. First, similar to before, we know that there is a P^{NP} algorithm M that, on input $1^{j \cdot O_n}$, outputs $O_n = M_1$. Then, we can recursively show that each of the matrices $M_2; \dots; M_k$ can be constructed by an NP oracle machine: for each i , to construct M_i , we use the oracle to find the lexicographically first low-rank approximation of M_{i-1} .

Our succinct witness for O_n is M_k for a large constant k . M_k is small enough that we can construct a circuit for it by brute-force. The idea is to then repeatedly use the local decoding scheme we discussed earlier to construct circuits for $M_{k-1}; M_{k-2}; \dots; M_1$, since each corresponds to a low-rank approximation of the next. However, having a low-rank approximation of M_i is not enough to recover M_i exactly. To circumvent this issue, we apply locally-decodable codes to the matrices. Indeed, if our low-rank decomposition $A_i B_i$ gives a $(1 - \epsilon)$ -approximation to the matrix $\text{Enc}(M_i)$ (the encoding of M_i using a suitable locally-decodable code), rather than to M_i , then we can use the local decoder to compute M_i exactly.

Locally-Decodable Codes and the Actual Compression Scheme $f_i(\cdot)$. We now give more details of the construction. We fix a locally-decodable code $\text{ECC}_{\text{local}}$, with message length $n^{1+\epsilon_{\text{enc}}}$ (ϵ_{enc} can be made an arbitrarily small constant), and a $\text{polylog}(n)$ -time local decoder. Let the encoder be $\text{Enc}: \{0, 1\}^n \rightarrow \{0, 1\}^{n^{1+\epsilon_{\text{enc}}}}$. The local decoder implies that, for $S \in \{0, 1\}^n$, if we have a T -size circuit which approximates the string $\text{Enc}(S)$, then there is a $\text{polylog}(n) \cdot T$ -size circuit which computes S exactly.

We now define two functions to describe how to go from a matrix to its low-rank decomposition. First define the function $\text{rk}(N) = 2^{(\log N)^b}$ for a constant $b > 0$. Then, for a string S , define $\text{comp}(S)$ as follows: Let $N = \lceil \sqrt{jSj} \rceil$ (we will pretend here that jSj is the square of an integer; in the real proof we use a slight padding to make sure of this), and let $A; B$ be two matrices in $\{0, 1\}^{N \times \text{rk}(N)}$ and $\{0, 1\}^{\text{rk}(N) \times N}$, respectively, such that AB equals S on the most possible positions (viewing S as a matrix in $\{0, 1\}^{N \times N}$). If there are multiple equally good options for $A; B$, then pick the lexicographically first one. We then define $\text{comp}(S) = AB$, as the concatenation of matrices A and B .

Next, we define a series of functions which recursively give compressions of a given string S :

$$f_i(S) := \begin{cases} \text{Enc}(S) & i = 1, \\ \text{Enc}(\text{comp}(f_{i-1}(S))) & i \geq 2. \end{cases}$$

Note that A and B (the outputs of $\text{comp}(S)$) can be computed from S in $\text{TIME}[\text{poly}(jSj)]^{\text{NP}}$. Now, we set $\hat{n}_{ij} = \lceil \sqrt{jf_i(O_n)j} \rceil$. We can then pick our NP oracle machine to, on input $1^{\hat{n}_{ij}}$, output the corresponding matrix for $f_i(O_n)$. (In the full proof below we use some simple tricks to make sure the \hat{n}_{ij} 's are all distinct.) Therefore, by assumption, we know that each $f_i(O_n)$ can be approximated by a $\text{rk}(\hat{n}_{ij})$ -rank matrix.

Finally, we are ready to implement our bootstrapping. We know that, for a parameter j , $f_j(O_n)$ has an \hat{n}_{ij} -size circuit which computes it exactly. Suppose we have a T -size circuit C_j which $(1 - \epsilon)$ -approximates $f_j(O_n)$. From this we can construct a circuit C_{j-1} which $(1 - \epsilon)$ -approximates $f_{j-1}(O_n)$ as follows:

First, applying the local decoder for $\text{ECC}_{\text{local}}$, we can construct a $\text{polylog}(\hat{n}_{ij-1})$ T -size circuit C_{comp} which exactly computes $\text{comp}(f_{j-1}(O_n))$.

Let A, B be the matrices corresponding to $\text{comp}(f_{j-1}(O_n))$. By our assumption, $A \cdot B$ is a $(1 - \epsilon)$ -approximation for $f_{j-1}(O_n)$. We know that $(A \cdot B)_{x,y}$ can be computed in $\text{rk}(O_n)^{1+o(1)}$ time, given oracle access to \mathbb{C}_{comp} . It follows from the locally-decodable property of $\text{ECC}_{\text{local}}$ that we get a circuit of size $\text{rk}(O_n)^{1+o(1)}$ $\text{polylog}(\text{rk}(O_n))$ which approximates $f_{j-1}(O_n)$.

From this construction, we can show that $f_1(O_n) = \text{End}(O_n)$ can be approximated by a small circuit, which in turn shows O_n has a small exact circuit. It is not hard to see, in particular, that

$$\text{SIZE}(O_n) \leq \sum_{i=1}^{\log N} \text{rk}(O_n)^{1+o(1)} \cdot \text{poly}(\text{rk}(O_n)) = 2^{O(n^b)} \cdot \text{poly}(n)$$

The Constant $\epsilon = 1/4$. Supposing that O_n has a 2^{n^a} -size witness, and the rank we consider is $\text{rk}(N) = 2^{(\log N)^b}$, the running time of our algorithm is

$$2^{O(n^a)} \cdot 2^{n \cdot (\log N)^b} = 2^{n + O(n^a) \cdot (\log N)^b}.$$

In order to make the above faster than 2^n and get a contradiction, we want to pick $b < 1 - a$. From the bound on $\text{SIZE}(O_n)$, we can see that the bootstrapping scheme can only achieve $a > b$. Therefore, we set $a = 1 - 2\epsilon$ and $b = 1 - 2\epsilon^2$, for a small constant $\epsilon > 0$.

We now consider the running time of M_{comp} . Since we only aim to compress O_n to a witness of size $2^{O(n^a)}$, we can stop if we find a witness of size 2^{n^a} , as there is no need to further compress. Let $M := M_{\text{comp}}$. On input 1^M , M_{comp} needs $\text{poly}(\text{rk}(O_n)) = 2^{O(n)} \cdot M^{\log M}$ time to compute $f_1(O_n)$. Therefore, M_{comp} runs in $\text{TIME}[n^{\log n}]^{\text{NP}}$, and hence yields a rigid matrix constructible in $\text{TIME}[n^{\log n}]^{\text{NP}}$ for rank $2^{(\log N)^{1-2\epsilon}}$. In other words, the time is slower than we hoped for, but the rank is higher than we hoped for.

Finally, we use a tensor product argument (Lemma 11.5 below) to transform this into a P^{NP} construction, which is rigid for a worse rank of $2^{(\log N)^{1-4\epsilon}}$. The idea is to take the tensor product of our rigid matrix with a large all-1s matrix. The resulting matrix is still rigid for the same rank, but has larger dimensions. Equivalently, in terms of the dimension N of the matrix, the complexity to compute the matrix has gone down, but it is also rigid for a lower rank.

11.2 Tools from Complexity Theory

Our construction of rigid matrices makes use of a number of tools from the complexity theory literature; in this Section we precisely define the tools from prior work which we will use.

The Circuit Evaluation Problem (Circuit-Eval) is the language of pairs $(C; w)$ where C is a general fan-in-2 circuit, and w is an input such that $C(w) = 1$. For two strings $a; b$ we use $a \cdot b$ to denote their concatenation.²

²The symbol \cdot is also used for circuit composition; its meaning will always be clear from context.

Probabilistic Checkable Proofs of Proximity

Our proof will make heavy use of probabilistically checkable proofs of proximity.

Definition 11.1 (Probabilistic Checkable Proofs of Proximity (PCP of proximity, or PCPP)). For $s, \delta : \mathbb{N} \rightarrow [0; 1]$ and $r, q : \mathbb{N} \rightarrow \mathbb{N}$, a verifier V is a PCP of proximity system for a pair language L with proximity parameter δ , soundness parameter s , number of random bits r and query complexity q if the following holds for all x, y :

(Completeness) If $(x, y) \in L$, then there is a proof π such that $V(x)$ accepts oracle y with probability 1.

(Soundness) If y is $(j(x))$ -far from $L(x) := \{z : (x, z) \in L\}$, then for all proofs π , $V(x)$ accepts oracle y with probability at most $s(j(x))$.

$V(x)$ tosses $r(j(x))$ random coins, and makes at most $q(j(x))$ non-adaptive queries.

Lemma 11.1 ([BGH⁺06, Theorem 3.3]) For any constants $0 < \delta; s < 1$, there is a PCP of proximity system for Circuit-Eval with proximity δ , soundness s , number of random bits $r = O(\log n)$ and query complexity $q = O(1)$. Moreover, given the pair $(C; w) \in \text{Circuit-Eval}$ a proof π which makes $V(C)$ always accept can be constructed in $\text{poly}(|C| + |w|)$ time.

Remark 11.1. The last ('Moreover') sentence is not explicitly stated in [BGH06], but it is evident from their construction.

In this paper, we need a stronger PCPP construction which is additionally smooth meaning, every position in the proof is queried with equal probability (assuming without loss of generality that all queries are non-adaptive and distinct). Such a construction can be found in [Par19].

Lemma 11.2 ([Par19]). For any constants $0 < \delta; s < 1$, there is a smooth PCP of proximity system for Circuit-Eval with proximity δ , soundness s , number of random bits $r = O(\log n)$ and query complexity $q = O(1)$. Moreover, given the pair $(C; w) \in \text{Circuit-Eval}$ a proof π making $V(C)$ always accepts can be constructed in $\text{poly}(|C| + |w|)$ time.

Error Correcting Codes

We also need standard constructions of two different types of codes: constant-rate linear error correcting codes, and ϵ -rate codes with $\text{polylog}(n)$ time local decoders.

Lemma 11.3 ([Spi96]). There is a constant-rate linear error correcting code ECC with a linear-time encoder Enc and a linear-time decoder Dec recovering error up to a universal constant ϵ .

Lemma 11.4 (cf, Section 2.3 of [Yek12]) For any constant $\epsilon > 0$, there is an ϵ -rate error correcting code ECC with a $\text{poly}(n)$ -time encoder Enc and a $\text{polylog}(n)$ -time local-decoder Dec which recovers up to $\epsilon/10$ fraction of errors.

³[Par19]'s construction actually ensures that this holds for every query position in the second input y as well. This additional property is not required by our proof.

A Simple Fact About Matrix Rigidity

We use 1_N to denote the all-ones matrix of size $N \times N$, and \otimes to denote the Kronecker product of matrices.

Lemma 11.5. For any field F and any matrix $A \in F^{M \times M}$, we have

$$R_{1_N \otimes A}(r) = R_A(r) \cdot N^2.$$

Proof. We first show $R_{1_N \otimes A}(r) \leq R_A(r) \cdot N^2$. Assume to the contrary that there is a way to change $k < R_A(r) \cdot N^2$ entries of $1_N \otimes A$ to make its rank r . The matrix $1_N \otimes A$ consists of N^2 disjoint copies of A , so by the pigeonhole principle, there were at most $k/N^2 < R_A(r)$ entries changed in one of those copies. Thus, that submatrix still has rank greater than r after the change, a contradiction.

We next show that $R_{1_N \otimes A}(r) \geq R_A(r) \cdot N^2$. Let B be a matrix of rank r whose Hamming distance from A is $R_A(r)$. Thus, $1_N \otimes B$ has rank $\text{rank}(1_N) \cdot \text{rank}(B) = r$, and its Hamming distance from $1_N \otimes A$ is $R_A(r) \cdot N^2$. \square

F_{p^r} -# OV

One crucial component of our construction is the algorithm for F_{p^r} -# OV from [CW16].

Definition 11.2. For a prime power $q = p^r$, in an F_q -# $OV_{n,d}$ instance, we are given two collections of vectors from F_q^d , $A = \{a_1, a_2, \dots, a_n\}$ and $B = \{b_1, b_2, \dots, b_n\}$, and want to compute the number of pairs such that $\langle a_i, b_i \rangle = 0$ over F_q .

We use the following algorithm for F_q -# $OV_{n,d}$.

Theorem 11.4 ([CW16]). For all fixed prime powers $q = p^r$, there is an $n^{2 \cdot (1 + \log(d/\log n))}$ time deterministic algorithm for F_q -# $OV_{n,d}$, when $d = n^{o(1)}$.

The original paper [CW16] only states an algorithm for # OV (the problem when A and B are collections of vectors from $\{0, 1\}^d$ and the inner product is over \mathbb{Z}). We make two small modifications to their algorithm to get the result stated in Theorem 11.4 above; see Section 11.6 for details.

11.3 Construction of Rigid Matrices Assuming an Easy Witness Lemma

We now move on to the formal construction and proof. We begin in this Section by proving Theorem 11.2.

We say an algorithm is a matrix-constructing algorithm if on input 1^N , it outputs a matrix in $\{0, 1\}^{N \times N}$. We say a function $f : \mathbb{N} \rightarrow \mathbb{N}$ is a typical resource bound function if it is strictly increasing, and satisfies $f(n) = \Theta(f(n+1))$. We first prove the following lemma, which says that if certain non-deterministic time classes have easy witnesses, then there is a P^{NP} construction of rigid matrices.

Lemma 11.6. There is an absolute constant $\epsilon > 0$ such that, for all prime powers $n = p^r$, and any three typical resource bound functions $T, S, R : \mathbb{N} \rightarrow \mathbb{N}$ with $T(n); S(n) \leq n$ for all n , the following three conditions cannot hold simultaneously.

- (1) All polynomial-time verifiers⁴ for unary $\text{NTIME}[T(n)]$ languages have $S(n)$ -size witness circuits.
- (2) For all P^{NP} matrix-constructing algorithms M , $R_{M(1^N)}(R(N)) \leq N^2$ for almost all N .
- (3) $\log T(n) = \log R(N) = \Omega(\log \log T(n) + \log S(n))$, where $N = 2^{n^{\epsilon}}$, for $n = \log T(n) + O(\log \log T(n)) + O(\log S(n))$, and $R(N) = N^{\Omega(1)}$.

Remark 11.2. In the following proof, we will actually only need the first assumption to hold for the special PCP verifier $V(1^n)$ of the language L we consider. This remark will be useful in the proof in the next Section.

Proof. Let $\epsilon > 0$ be a constant to be decided later. We first only consider the case when the field is F_2 , and then show how to generalize the argument for other finite fields. We will assume that all three items are true, and derive a contradiction.

Unary Language L and PCP. Let L be a unary language in $\text{NTIME}[T(n)] \cap \text{NTIME}[T(n)=n]$. Using the non-deterministic time hierarchy theorem [Zak83], such an L exists because $\bar{T}(n)$ is a typical resource function. Let V be an efficient PCP verifier for L from [BV14]. That is, there is a function $\ell = \ell(n) = \log T(n) + O(\log \log T(n))$, such that $V(1^n)$ takes an oracle $O : \{0, 1\}^{\ell} \rightarrow \{0, 1\}$ and ℓ random bits as input, runs in $\text{poly}(n)$ time, and:

1. (PCP Completeness) If $1^n \in L$, then there exists a circuit $C : \{0, 1\}^{\ell} \rightarrow \{0, 1\}$ of size $S(n)$ such that $\Pr_{r \in \{0, 1\}^{\ell}} [V(1^n)^C(r) = 1] = 1$. (This follows from the first assumption of the Lemma.)
2. (PCP Soundness) If $1^n \notin L$, then for all oracles $O : \{0, 1\}^{\ell} \rightarrow \{0, 1\}$, we have $\Pr_{r \in \{0, 1\}^{\ell}} [V(1^n)^O(r) = 1] \leq \epsilon$.

We will next show how to put $L \in \text{NTIME}[T(n)=n]$ by using the second and the third assumptions of the Lemma, which will give us the contradiction we want.

The Plan. Let C_{best} be the circuit of size $S(n)$ such that $\Pr_{r \in \{0, 1\}^{\ell}} [V(1^n)^{C_{\text{best}}}(r) = 1] = 1$, and if there are multiple such circuits, we break the tie by choosing the lexicographically first one. Note that such a circuit doesn't exist when $1^n \notin L$, and in that case we set C_{best} to be a trivial circuit which always outputs 0.

In our non-deterministic algorithm to solve L , given an input 1^n , we first guess a circuit C of size at most $S(n)$, and wish to ensure that the following two conditions hold:

⁴That is, for $L \in \text{NTIME}[T(n)]$, the verifier V takes two inputs $x; y$ with $|x| = n$ and $|y| = \text{poly}(T(n))$, runs in $\text{poly}(|x| + |y|)$ time, and has the property that $x \in L$ if and only if there is a y such that $V(x; y) = 1$.

1. When $1^n \leq L$ and $C = C_{\text{best}}$, we accept, and
2. When $1^n \not\leq L$, we always reject.

If our algorithm satisfies these two conditions and runs in $T(n) = n$ non-deterministic time, then we have $L \leq \text{NTIME}[T(n) = n]$ and arrived at the desired contradiction.

Implementation. Now suppose we have guessed a circuit C of size at most $S(n)$. Toward achieving the two conditions above, we want to estimate

$$p_{\text{acc}}(C) := \Pr_{r \in \{0,1\}^n} [V(1^n)^C(r) = 1]:$$

Define another circuit $D_C : \{0,1\}^n \rightarrow \{0,1\}$ as $D_C(r) := V(1^n)^C(r)$. We thus equivalently have that

$$p_{\text{acc}}(C) = \Pr_{r \in \{0,1\}^n} [(D_C; r) \in \text{Circuit-Eval}];$$

by the definition of Circuit-Eval

Applying Error Correcting Codes. Fix an F_2 -linear error correcting code ECC with rate c_1 and recovering error ϵ_1 , whose existence is guaranteed by Lemma 11.3. Let $\text{Enc} : \{0,1\}^n \rightarrow \{0,1\}^{c_1 n}$ and $\text{Dec} : \{0,1\}^{c_1 n} \rightarrow \{0,1\}^n$ be the corresponding linear-time encoder and decoder.

We now define yet another circuit $E_C : \{0,1\}^{c_1 n} \rightarrow \{0,1\}^n$ as $E_C(w) := D_C(\text{Dec}(w))$. Then it suffices to estimate

$$p_{\text{acc}}(C) = \Pr_{r \in \{0,1\}^{c_1 n}} [(E_C; \text{Enc}(r)) \in \text{Circuit-Eval}]:$$

Notice that $\text{SIZE}(E_C) = \text{poly}(n) \leq S(n)$, since the verifier $V(1^n)$ runs in $\text{poly}(n)$ time, and the decoder Dec runs in linear time.

Applying the PCPP. Now we use a $q_{\text{PCPP}} = O(1)$ -query smooth PCPP for Circuit-Eval from Lemma 11.2 with constant soundness s_{PCPP} and proximity parameter ρ_{PCPP} to be specified later. Let $V_{\text{C-EVAL}}(E_C)$ be the verifier for this smooth PCPP with the circuit fixed to E_C . Hence, $V_{\text{C-EVAL}}(E_C)$ uses proof length $\ell_{\text{proof}} = \text{poly}(\text{SIZE}(E_C)) = \text{poly}(S(n))$ and $m = O(\log \ell_{\text{proof}})$ random bits.

Claim 11.1. $V_{\text{C-EVAL}}(E_C)$ satisfies the following three properties by setting $q_{\text{PCPP}} < \frac{1}{3}$ and $s_{\text{PCPP}} = \frac{1}{3}$.

1. (PCPP Completeness) If $(D_C; r) \in \text{Circuit-Eval}$ there is a proof $u \in \{0,1\}^{\ell_{\text{proof}}}$ that

$$\Pr_{u \in \{0,1\}^m} [V_{\text{C-EVAL}}(E_C)^{\text{Enc}(r)}(u) = 1]:$$

2. (From PCPP Smoothness) Suppose $(D_C; r) \in \text{Circuit-Eval}$ and let $\text{proof} \in \{0, 1\}^{\ell}$ be a proof satisfying the previous property. If $\text{proof} \in \{0, 1\}^{\ell}$ is a $(1 - \epsilon)$ -approximation to $\text{Dec}(r)$ for some $\epsilon \in [0, 1]$, then

$$\Pr_{u \in \{0, 1\}^m} [V_{\text{C-EVAL}}(E_C)^{\text{Enc}(r)}(u) = 1 - \epsilon] \geq \epsilon_{\text{PCPP}}$$

3. (PCPP Soundness) If $(D_C; r) \notin \text{Circuit-Eval}$ then for all proofs $\text{proof} \in \{0, 1\}^{\ell}$, we have

$$\Pr_{u \in \{0, 1\}^m} [V_{\text{C-EVAL}}(E_C)^{\text{Enc}(r)}(u) = 1] \leq \epsilon_{\text{PCPP}}$$

Property (1) of Claim 11.1 follows from the completeness property of the PCPP system, and property (2) follows from the smoothness of the PCPP system combined with a simple union bound.

For property (3), note that if $(D_C; r) \notin \text{Circuit-Eval}$ then $\text{Enc}(r)$ is ϵ_{dec} -far from the set $\{w \in \{0, 1\}^{\ell} : (E_C; w) \in \text{Circuit-Eval}\}$. This is because for any string $w \in \{0, 1\}^{\ell}$ which is ϵ_{dec} -close to $\text{Enc}(r)$, we know $\text{Dec}(w) = r$ and hence $(E_C; w) \in \text{Circuit-Eval}$. Therefore, by setting $\epsilon_{\text{PCPP}} < \epsilon_{\text{dec}}$ and $\epsilon_{\text{PCPP}} = 1 - \epsilon_{\text{dec}}$, property (3) follows from the soundness of the PCPP system.

The Function $C_{\text{best}}(r; j)$. Note that by Lemma 11.2, there is a polynomial time computable function $(E_C; \text{Enc}(r)) \in \{0, 1\}^{\ell}$, such that when $(E_C; \text{Enc}(r)) \in \text{Circuit-Eval}$ we have

$$\Pr_{u \in \{0, 1\}^m} [V_{\text{C-EVAL}}(E_C)^{\text{Enc}(r)}(u) = (E_C; \text{Enc}(r))(u)] = 1$$

Define the Boolean function $C(r; j)$, for $j \in [\ell]$, to be the j -th bit of $(E_C; r)$ (suppose ℓ is a power of 2 for simplicity).

The function $C_{\text{best}}(r; j)$ is computable in E^{NP} , by using the following procedure. First we show how to compute the circuit C_{best} in E^{NP} . We are given two inputs $r; j$ with length $|r| = \ell$ and $|j| = \log \ell = O(\log S(n))$. In $O(2^{\ell})$ time with an NP oracle, we can first decide whether (1^n) always accepts a circuit of size at most $S(n)$. If not, then we just output a trivial circuit. If so, then we guess that circuit bit by bit to construct the lexicographically first one, again using the NP oracle to check each guess. In this way, we can compute C_{best} in E^{NP} . We can then construct the circuit $E_{C_{\text{best}}}$ from C_{best} , and then (using the fact that $(E_C; \text{Enc}(r))$ can be computed in polynomial time) compute $C_{\text{best}}(r)$ and output its j -th bit. The whole procedure runs in E^{NP} .

The E^{NP} Machine M_{rigid} . Note that C_{best} has input length $n = \ell + O(\log S(n))$. We can thus construct a E^{NP} machine M_{rigid} such that, given an input $1^{2^n - 2}$, it outputs the truth-table of C_{best} as a matrix. Therefore, by the second assumption of the Lemma, C_{best} as a matrix can be ϵ -approximated by a matrix of rank $R(2^n - 2)$.

Putting L in $\text{NTIME}[T(n)=n]$. Finally, consider the following algorithm for solving L . We first guess a circuit C of size $S(n)$, with the hope that it is C_{best} . Then, letting $N = 2^n = 2$, we guess a matrix $M : N \times N$ of rank $R(N)$, with the hope that it ϵ -approximates C_{best} . More specifically, we guess two matrices U, V of size $N \times R(N)$ and $R(N) \times N$, and set (implicitly, without explicitly computing it) $M = UV$.

Now we try to calculate

$$p_{\text{acc}}(M) := \Pr_{r \in \{0,1\}^n; u \in \{0,1\}^{m^n}} [V_{C\text{-EVAL}}(E_C)^{\text{Enc}(r)} M(r)(u) = 1]$$

Fix u , and suppose that for randomness s , the verifier $V_{C\text{-EVAL}}(E_C)$ queries $M(r; j_1); M(r; j_2); \dots; M(r; j_{q_1})$ in $M(r; \cdot)$, and $e_1; e_2; \dots; e_{q_2}$ in $\text{Enc}(r)$ (note that $V_{C\text{-EVAL}}(E_C)$'s query positions only depend on the randomness s). Now we want to estimate

$$\Pr_{r \in \{0,1\}^n} [F_u(M(r; j_1); M(r; j_2); \dots; M(r; j_{q_1}); \text{Enc}(r)_{e_1}; \text{Enc}(r)_{e_2}; \dots; \text{Enc}(r)_{e_{q_2}}) = 1]$$

for a Boolean function F_u on $q_{\text{PCPP}} = q_1 + q_2$ inputs. First, we can write F_u in the basis of XOR functions:

$$F_u(z_1; z_2; \dots; z_{q_{\text{PCPP}}}) = \sum_{S \subseteq [q_{\text{PCPP}}]} \alpha_S \prod_{i \in S} z_i$$

(Here, we consider the XOR function to be outputting a $\{0,1\}$ value, and the coefficients α_S and the sum are over \mathbb{R} , not over \mathbb{F}_2 .) Since our goal is to compute the expected value of F_u , by linearity of expectation, it suffices to separately compute the expected value of each of the (constant number of) parity functions. Therefore, it suffices to consider the case where F_u is just an XOR function.

Also, note that since ECC is a linear code, it follows that $\text{Enc}(r)_k$ is an XOR function on a subset of coordinates of r . Thus, if $r = a \oplus b$ where $|a| = n/2$ and $|b| = n/2$ (note that $n > n/2$ by the third assumption of the Lemma), we have $\text{Enc}(r)_e = \text{Enc}_L(a)_e \oplus \text{Enc}_R(b)_e$, where $\text{Enc}_L(a)_e$ and $\text{Enc}_R(b)_e$ are the corresponding contributions of a and b to $\text{Enc}(r)_e$.

Next, we define

$$E_L(a)_e := \begin{cases} (1; 0) & \text{if } \text{Enc}_L(a)_e = 0, \\ (0; 1) & \text{if } \text{Enc}_L(a)_e = 1, \end{cases} \quad \text{and} \quad E_R(b)_e := \begin{cases} (0; 1) & \text{if } \text{Enc}_R(b)_e = 0, \\ (1; 0) & \text{if } \text{Enc}_R(b)_e = 1. \end{cases}$$

It is easy to verify that $\langle E_L(a)_e; E_R(b)_e \rangle = \text{Enc}_L(a)_e \oplus \text{Enc}_R(b)_e = \text{Enc}(r)_e$.

Constructing F_2 -# OV Instance. We can now simplify the quantity we want to compute as

$$\Pr_{a \in \{0,1\}^n; b \in \{0,1\}^n} \sum_{i=1}^m \sum_{j=1}^m \langle U_{a,i}; V_{b,j} \rangle = \Pr_{a \in \{0,1\}^n; b \in \{0,1\}^n} \sum_{i=1}^{q_1} U_{a,i} \cdot \sum_{j=1}^{q_2} V_{b,j} = 0 :$$

In above, we used U_i and V_j to denote the i -th row of U and j -th column of V respectively, so that $\langle U_i; V_j \rangle = M_{ij}$. By duplicating each of the b 's 2^n times, the above can be reduced to a counting F_2 -# OV $_{N;d}$ instance, with $N = 2^n$ vectors of $d = O(R(N))$ dimensions. By Theorem 11.4, this can be solved in time

$$\begin{aligned} N^{2(1+\log d)} &= N^{2(1+\log R(N))} \\ &= 2^{n(2+\log R(N))} \\ &= 2^{\log T(n) + O(\log \log T(n)) + O(\log S(n))} \quad (\log T(n) = \log R(N)) \\ &= (n = \log T(n) + O(\log S(n)) = \log T(n) + \log \log T(n) + S(n)) \end{aligned}$$

Since we also need $O(S(n))$ time for enumerating all possible $a \in \{0,1\}^m$, the overall running time for calculating $p_{acc}(M)$ is

$$2^{\log T(n) + O(\log \log T(n)) + O(\log S(n))} \quad (\log T(n) = \log R(N))$$

By our third assumption, we know the above running time is $2^{\log T(n) + O(\log S(n))}$ $T(n)=n$, since $S(n) = n$.

Analysis of the Algorithm. Consider first when $1^n \in L$. We know that on the correct guess of $C = C_{best}$ and the appropriate $M = M_{C_{best}}$, we have that M $(1 - \epsilon)$ -approximates C_{best} . That is, for a random $r \in \{0,1\}^n$, the average relative distance between $M(r)$ and $C_{best}(r)$ is at most ϵ . Hence, by Property (2) of Claim 11.1 and by linearity of expectation, we know that $p_{acc}(M) > 1 - \epsilon_{PCPP}$ in this case.

Otherwise, if $1^n \notin L$, then for every guess of C and M , by the soundness property of PCP, we know that

$$\Pr_{r \in \{0,1\}^n} [(D_C; r) \in \text{Circuit-Eval}] \leq \epsilon$$

Then by Property (3) of Claim 11.1, we have that

$$p_{acc}(M) \leq \epsilon + \epsilon = 2\epsilon$$

Therefore, when we set ϵ to be small enough so that $1 - \epsilon_{PCPP} > 2\epsilon$, we can distinguish the above two cases. By the above argument, this puts $\#P \subseteq NTIME[T(n)=n]$, a contradiction.

Adaptation for the field F_q . Let $q = p^r$ be a prime power. In the following, we sketch the adaptation to deal with F_q . The only thing we need to modify is how to reduce the computation of $\text{p}_{\text{acc}}(M)$ to F_q -# OV. Again, we guess a $\text{ran} \mathbb{R}(N)$ matrix $M = UV$ over F_q , and we want to calculate

$$\Pr_{r \in \{0,1\}^g} [F_u(M(r; j_1)^{q-1}; \dots; M(r; j_{q_1})^{q-1}; \text{End}(r)_{e_1}; \dots; \text{End}(r)_{e_{q_2}}) = 1]$$

for a Boolean function F_u on $q_{\text{PCPP}} = q_1 + q_2$ inputs. Note that in the above, we raise all the $M(r; j_i)$ inputs to the $(q-1)$ -th power to make them Boolean. Now, we can write F_u as a real sum of $2^{q_{\text{PCPP}}}$ AND functions, each one for a subset of the inputs of F_u . Hence, like before, it suffices to consider the case when F_u is an AND function, and in this case we want to calculate

$$\Pr_{r \in \{0,1\}^g} \left[\prod_{i=1}^{q_1} M(r; j_i)^{q-1} \prod_{i=1}^{q_2} \text{End}(r)_{e_i} = 1 \right];$$

which is equivalent to

$$\begin{aligned} & \Pr_{a \in \{0,1\}^{g_1}} \Pr_{b \in \{0,1\}^{g_2}} \left[\prod_{i=1}^{q_1} U_a(j_i)^{q-1} \prod_{i=1}^{q_2} E_L(a)_{e_i}; E_R(b)_{e_i} = 1 \right] \\ &= \Pr_{a \in \{0,1\}^{g_1}} \Pr_{b \in \{0,1\}^{g_2}} [h[a; b] = 0]; \end{aligned}$$

where

$$a := \left(\prod_{i=1}^{q_1} U_a(j_i)^{q-1} \prod_{i=1}^{q_2} E_L(a)_{e_i} \right) \cdot 1$$

and

$$b := \left(\prod_{i=1}^{q_1} V_b(j_i) \prod_{i=1}^{q_2} E_R(b)_{e_i} \right) \cdot 1$$

The final equality follows from the fact that for vectors $a_1; b_1; a_2; b_2$, we always have $h[a_1; b_1] h[a_2; b_2] = h[a_1 a_2; b_1 b_2]$. Finally, the above can be reduced to an F_q -# OV instance with $2^{g_1+g_2}$ vectors of $\mathbb{R}(N)^{O(1)}$ dimensions. One can see that this polynomial blowup in the dimension is acceptable, and we can still proceed as in the case of F_2 . \square

Now we are ready to prove Theorem 11.2 (restated below). Notice that here we use the stronger condition $\text{NQP } \# P_{\text{poly}}$ instead of $\text{NE } \# P_{\text{poly}}$.

Reminder of Theorem 11.2 There is an absolute constant $\epsilon > 0$ such that, for all prime powers $q = p^r$ and all $n > 0$ at least one of the following holds:

$\text{NQP } \# P_{\text{poly}}$.

There is a P^{NP} machine M such that, for infinitely many N 's, on input 1^N , M outputs an $N \times N$ matrix $H_N \in \{0,1\}^{N \times N}$ such that $R_{H_N} \left(2^{(\log N)^{1-\epsilon}} \right) \geq N^2$ over F_q .

Proof of Theorem 11.2. Let $\epsilon > 0$ be a constant to be chosen later.

Assume that $\text{NQP} = \text{P}_{\text{poly}}$. By [MW18], this in particular implies that for a constant b to be specified later and $T(n) := 2^{\log^b n}$, all polynomial-time verifiers for unary languages in $\text{NTIME}[2^{\log^b n}]$ have $S(n) := n^k$ -size witness circuits, for a constant $k = k(b)$.

Set $R(N) := 2^{(\log N)^{1-\epsilon}}$. We will now apply Lemma 11.6 with $R; S; T$ as above. Note that $n = \log T(n) + O(\log \log T(n)) + O(\log S(n)) = \log^b n + O(\log n)$ and $N = 2^{n-\epsilon} = 2^{\log^b n - 2 + O(\log n)}$. We thus calculate that

$$\begin{aligned} \log T(n) &= \log R(N) & \log^b n &= \log^{b(1-\epsilon)} n \\ & & &= \log^{b-\epsilon} n \\ & & &= \Omega(\log n) \\ & & &= \Omega(\log \log T(n) + \log S(n)); \end{aligned}$$

if we set $b > 2/\epsilon$.

Therefore, since Conditions (1) and (3) of Lemma 11.6 hold, we conclude that Condition (2) of Lemma 11.6 does not hold, and this completes the proof.

11.4 Unconditional Construction of Rigid Matrices

In this Section, we prove Theorem 11.1, giving our main construction of rigid matrices, by using an additional bootstrapping argument.

For an integer $n \geq N$, we write $n_{[k]}$ to denote the smallest integer $m \leq n$ such that $m \equiv 2^k - 1 \pmod{2^{k+1}}$. Notice that $n_{[k]} \leq n$ and $n_{[k]} \equiv 2^k - 1 \pmod{2^{k+1}}$. Moreover, for all integers $n; m; i; j \geq N$ with $i \neq j$, we have that $n_{[i]} \neq m_{[j]}$.

We first prove the following lemma, which gives an (unconditional) construction of a matrix which is rigid for a higher rank than the construction in Theorem 11.1, but with a slower construction time of $\text{TIME}[n^{\log n}]^{\text{NP}}$.

Lemma 11.7. There is an absolute constant $\epsilon > 0$ such for all prime powers $q = p^r$ and all constants $\epsilon > 0$:

There is a $\text{TIME}[n^{\log n}]^{\text{NP}}$ machine M such that, for infinitely many N 's, on input 1^N , M outputs an $N \times N$ matrix $H_N \in \mathbb{F}_q^{N \times N}$ such that $\text{R}_{H_N}(2^{(\log N)^{1-\epsilon}}) < \epsilon N^2$ over \mathbb{F}_q .

Proof. Let ϵ be a constant to be specified later. For simplicity, we only consider the finite field \mathbb{F}_2 in the following. It is not hard to see that our proof also works for all finite fields \mathbb{F}_p with a straightforward modification.

Assume toward a contradiction that for all $\text{TIME}[n^{\log n}]^{\text{NP}}$ machines M , and for almost all input lengths N , the output matrix $H_N \in \mathbb{F}_2^{N \times N}$ of M satisfies $\text{R}_{H_N}(2^{(\log N)^{1-\epsilon}}) < \epsilon N^2$. (By padding with zeros or only keeping the first N^2 output bits, we can always assume that M outputs exactly N^2 bits on inputs of length N .)

Notation. Throughout the proof, we will often identify a matrix from $\{0, 1\}^{N \times N}$ with a string from $\{0, 1\}^{N^2}$ (reading the matrix from top row to bottom row, and from leftmost column to rightmost column to construct the corresponding string).

Define the functions $\text{rk}(N) := 2^{\lfloor \log N \rfloor - 2}$ and $\text{comp}(N) := 2^{\lfloor \log N \rfloor - \text{rk}(N)}$.

Set $\epsilon_{\text{enc}} = 0.01$, and $\ell_{\text{enc}}(N) := N^{1 + \epsilon_{\text{enc}}}$. Define the function $\ell_{\text{PCP}}(N) := N \log^{C_{\text{PCP}}} N$ for a constant C_{PCP} to be specified later.

Applying Lemma 11.4, we fix a locally-decodable error correcting code $\text{ECC}_{\text{local}}$ with a poly(N)-time encoder $\text{Enc} : \{0, 1\}^N \rightarrow \{0, 1\}^{\ell_{\text{enc}}(N)}$, which has a $(\log N)^{C_{\text{enc}}}$ -time randomized decoder that decodes any position with probability at least $1 - \epsilon_{\text{enc}}$ when given oracle access to a codeword which is corrupted in less than a fraction of its entries.

The Compression Scheme $f_i(\cdot)$. Now, given a string $S \in \{0, 1\}^N$, we define the function $\text{comp}(S)$ as follows. Let N^0 be the smallest square number $\geq N$ and let $A, B \in \{0, 1\}^{N^0 \times \text{rk}(N^0)}$ be the two matrices such that $S \cdot 0^{N^0 - N}$ (interpreted as a $\{0, 1\}^{N^0 \times N^0}$ matrix) agrees with AB^T (over F_2) on the greatest number of positions. In case of a tie, make the choice resulting in $A \leq B$ being the lexicographically earliest string. We define $\text{comp}(S) := A \cdot B$.

Given a string $S \in \{0, 1\}^N$, we further define the sequence of functions $\mathfrak{f}_i(S) := \text{Enc}(S)$ and $f_i(S) := \text{Enc}(\text{comp}(f_{i-1}(S)))$ for $i > 1$.

We now aim to apply Lemma 11.6 from the previous section. Fix a unary language $L \in \text{NTIME}[2^n]$ such that $L \not\in \text{NTIME}[2^n - n]$ [Zák83]. Fix an efficient PCP verifier V for L from [BV14], such that $V(1^n)$ takes $\log_{\text{PCP}}(2^n)$ random bits and oracle access to a string of length $\ell_{\text{PCP}}(2^n)$. In order to apply Lemma 11.6, we need to show $V(1^n)$ has small witness circuits.

The Construction of the $\text{TIME}[n^{\log n}]^{\text{NP}}$ Machine M_{comp} : A Bootstrapping Argument. Let $O_n \in \{0, 1\}^{\ell_{\text{PCP}}(2^n)}$ be the lexicographically first string which makes $V(1^n)$ always accept, if such a string exists, and $\perp^{\ell_{\text{PCP}}(2^n)}$ otherwise.

Our $\text{TIME}[n^{\log n}]^{\text{NP}}$ machine M_{comp} works as follows. For n and $1 \leq i \leq 3 \log n$, let $\ell_{n,i} := \lfloor \frac{\ell_{\text{PCP}}(2^n)}{i} \rfloor$. If $\ell_{n,i} \geq 2^{n^{1-2^{-i}}}$ for a constant ϵ_1 to be specified later, then M_{comp} on input $1^{\ell_{n,i}}$ computes $f_i(O_n)$, padded with $\ell_{n,i} - |f_i(O_n)|$ zeros. Otherwise it outputs an all-zero matrix.

We first claim that M_{comp} is well-defined, meaning there exists a constant N_0 such that for all $n \geq N_0$ and $1 \leq i \leq 3 \log n$, the $\ell_{n,i}$'s are distinct. To prove this, it suffices to show that $\ell_{n,i} < \ell_{m,i}$ whenever $i \leq 3 \log n$ and $n < m$, but this follows from the definitions of the function $f_i(\cdot)$'s.

Next, we note that M_{comp} indeed runs in $\text{TIME}[n^{\log n}]^{\text{NP}}$: on input $1^{\ell_{n,i}}$ of length $m = \ell_{n,i} \geq 2^{n^{1-2^{-i}}}$, the algorithm runs in time $\text{poly}(\ell_{n,i}) = 2^{O(n)} \cdot m^{\log m}$.

$V(1^n)$ Has Succinct Witness. We first show from our assumption (that $\text{TIME}[n^{\log n}]^{\text{NP}}$ does not have rigid matrices) that $V(1^n)$ has a succinct witness circuit if there is an oracle which always satisfies it.

When this is the case, notice that for all $i \leq 2=3 \log n$, the output of $M_{\text{comp}}(1^{n_i})$ can be ϵ -approximated by a matrix of rank $\text{rk}(\hat{c}_{n_i})$. We can calculate that $\hat{c}_{n_i; 2=3 \log n} < 2^{n^{1=2}}$; let j be the largest integer such that $\hat{c}_{n_i; j} < 2^{n^{1=2+ \epsilon}}$, and note that $\hat{c}_{n_i; j} < 2^{3n^{1=2+ \epsilon}}$. Hence, $M_{\text{comp}}(1^{n_i})$ can be implemented as a circuit of size $2^{O(n^{1=2+ \epsilon})}$.

Next, if there is a size S circuit which $(1-\epsilon)$ -approximates $M_{\text{comp}}(1^{n_i})$, then $M_{\text{comp}}(1^{n_i-1})$ can be $(1-\epsilon)$ -approximated by a $\text{rk}(\hat{c}_{n_i-1}) \text{ poly}(n) S$ size circuit by using the local decoder of the corresponding locally decodable codes. Therefore, $M_{\text{comp}}(1^{n-1})$ can be $(1-\epsilon)$ -approximated by a circuit of size

$$\sum_{i=1}^n \text{rk}(\hat{c}_{n_i}) n^{O(\log n)} 2^{O(n^{1=2+ \epsilon})} = 2^{O(n^{1=2+ \epsilon})}.$$

Since $M_{\text{comp}}(1^{n-1}) = \text{End}(O_n)$, it follows that O_n can be computed exactly by a $2^{O(n^{1=2+ \epsilon})}$ -size circuit.

Applying Lemma 11.6. Toward applying Lemma 11.6, we set $T(n) = 2^n$, $S(n) = 2^{O(n^{1=2+ \epsilon})}$ and $R(N) = 2^{(\log N)^{1=2+ \epsilon}}$, where $\epsilon := \epsilon/2 > 0$. The two parameters in Condition (3) of Lemma 11.6 are bounded by $n = \log T(n) + O(\log \log T(n)) + O(\log S(n)) = n + O(n^{1=2+ \epsilon})$ and $N = 2^{n=2+ O(n^{1=2+ \epsilon})}$. We thus calculate that

$$\log T(n) = \log R(N) = (n = n^{1=2+ \epsilon}) = \Omega(n^{1=2+ \epsilon}) = \Omega(\log \log T(n) + \log S(n)).$$

Therefore, Conditions (1) and (3) of Lemma 11.6 are satisfied, and it follows that Condition (2) must be violated, which completes the proof. \square

Finally, we prove Theorem 11.1 (restated below) by using a simple padding argument.

Reminder of Theorem 11.1 There is an absolute constant $\epsilon > 0$ such for all prime powers $q = p^r$ and all constants $\epsilon > 0$:

There is a P^{NP} machine M such that, for infinitely many N 's, on input 1^N , M outputs an $N \times N$ matrix $H_N \in \mathbb{F}_q^{N \times N}$ such that $R_{H_N}(2^{(\log N)^{1=2+ \epsilon}}) = \Omega(N^2)$ over \mathbb{F}_q .

Proof of Theorem 11.1. We have shown, from Lemma 11.7, that there is an absolute constant $\epsilon > 0$ such that for all constants $\epsilon > 0$:

There is a $\text{TIME}[n^{\log n}]^{\text{NP}}$ machine M such that, for infinitely many N 's, on input 1^N , M outputs an $N \times N$ matrix $H_N \in \mathbb{F}_2^{N \times N}$ such that $R_{H_N}(2^{(\log N)^{1=2+ \epsilon}}) = \Omega(N^2)$ over \mathbb{F}_2 .

Let $N^0 = N^{\log N}$, and consider the P^{NP} machine M^0 which, given an input 1^{N^0} , outputs a matrix $H_{N^0}^0 := 1_{N^{\log N}} \oplus H_N$. By Lemma 11.5, we have

$$R_{H_{N^0}^0}(2^{(\log N^0)^{1=2+ \epsilon}}) = \Omega(N^0)$$

for infinitely many N . This rigidity bound is equivalent to

$$R_{H^0}^{N^0}(2^{(\log N)^{1-\epsilon}}) = \Omega(N^\epsilon);$$

as desired.

11.5 Applications

Rigid matrices are known to have applications in many areas of complexity theory. In this section, we give three applications of our construction, to communication complexity, arithmetic circuit complexity, and Boolean circuit complexity.

11.5.1 PH^{CC} Communication Lower Bound

In this Section we apply our construction of rigid matrices to prove a PH^{CC} communication lower bound for functions in $\text{TIME}[2^{(\log n)^{1-\epsilon}}]^{\text{NP}}$. Our main tool will be a known connection between rigid matrices and PH^{CC} :

Lemma 11.8 ([Raz89], see also [Wun12]) Letting f be a function in PH^{CC} , the $2^n \times 2^n$ communication matrix M_f of f has $R_{M_f}(2^{(\log n)^c}) = \Omega(4^n)$ over F_2 , where $\epsilon > 0$ is arbitrary and $c > 0$ is a constant depending only on ϵ , but not n .

We will also use the following simple Lemma.

Lemma 11.9. For any field F and any matrix $A \in F^{N \times N}$, and for $M > N$, define $P_{A;M} \in F^{M \times M}$ to be the matrix such that the top-left $N \times N$ sub-matrix is A , and the rest of entries are all zeros. For all r , we have

$$R_{P_{A;M}}(r) = R_A(r);$$

Theorem 11.5. For all functions $f(n) = \Omega(1)$ such that $n^{(n)}$ is time-constructible, there is a function $f \in \text{TIME}[2^{(\log n)^{1-\epsilon}}]^{\text{NP}}$ which is not in PH^{CC} .

Proof. By Theorem 11.1, we know that there is a P^{NP} machine M such that $R_{M(n)}(2^{(\log N)^{1-\epsilon}}) = \Omega(N^2)$ over F_2 , for a constant $\epsilon > 0$ and infinitely many N 's. For simplicity, we can assume $n^{(n)} = \log n$ (e.g., by setting $n^{(n)} = \min(n^{(n)}; \log n)$).

The Definition of f . Now we define a function $f \in \text{TIME}[2^{(\log n)^{1-\epsilon}}]^{\text{NP}}$ as follows:

Given as input $x \in \{0, 1\}^n$, the function f outputs zero immediately if 4 does not divide n . Otherwise let $m = n/4$.

It treats the first $2m$ bits of the input as an integer N in $[2^{2m}]$, and if $N > 2^{(\log m)^{1-\epsilon}}$, it outputs zero.

Otherwise, it constructs the matrix $H = M(1^N)$. Let $S = 2^m$, and $Q = P_{1_{bS=N}^c} H; S$. It treats the next $2m$ bits of the input as a pair of integers $(i; j) \in [S] \times [S]$, and outputs $Q_{i;j}$.

$Q_{i;j}$ can be computed easily given H , so f can be computed in $\text{TIME}[2^{(\log n)^c}]^{NP}$.

f is not in PH^{cc} . We will now show that f , when interpreted as a communication problem, is not in PH^{cc} . We distribute the input bits of f among the two players as follows: When 4 divides n , setting $m = n/4$, then Alice holds the bits $x_1; x_2; \dots; x_m$ and $x_{2m+1}; \dots; x_{3m}$, and Bob holds the bits $x_{m+1}; x_{m+2}; \dots; x_{2m}$ and $x_{3m+1}; x_{3m+2}; \dots; x_{4m}$.

Assume to the contrary that $f \in \text{PH}^{cc}$. This means that for all assignments to $x_1; x_2; \dots; x_{2m}$, the restricted function $f : \{0, 1\}^m \times \{0, 1\}^m \rightarrow \{0, 1\}$ is still in PH^{cc} . That is, there exists a constant c , such that for all $N = 2^{(\log m)^c}$, $S = 2^m$, and $Q = P_{1_{bS=N}^c} M(1^N); S$, we have

$$R_Q(2^{(\log m)^c}) = 2 S^2.$$

By Lemma 11.9, this implies

$$R_{1_{bS=N}^c} M(1^N)(2^{(\log m)^c}) = 2 S^2 = 2^3 (bS=Nc N)^2.$$

By Lemma 11.5, this further implies

$$R_{M(1^N)}(2^{(\log m)^c}) = 2^3 N^2.$$

Now, let N be a sufficiently large integer such that

$$R_{M(1^N)}(2^{(\log N)^{1=5}}) = N^2.$$

Let m be the smallest integer such that $2^{(\log m)^{4m}} \geq N$. Since (n) is unbounded, we can pick N to be large enough such that $(4m - 4) \geq 20c$. By definition of m , we have $2^{(\log(m-1))^{4m-4}} < N$, meaning $2^{(\log(m-1))^{20c}} < N$, and so $2^{(\log m)^{10c}} < N$. But then by the above discussion, we have

$$R_{M(1^N)}(2^{(\log N)^{1=10}}) = 2^3 N^2;$$

a contradiction. □

11.5.2 Depth-2 Arithmetic Circuit Lower Bound

In this section we prove Theorem 11.6 (restated below). Recall first the definition of w_2 :

Definition 11.3. For a field F and a matrix $A \in F^{N \times N}$, let

$$w_2(A) := \min \{ \text{nnz}(B) + \text{nnz}(C) \mid A = BC \};$$

where the min is over all pairs $B; C$ of matrices of any dimensions over F whose product is A , and $\text{nnz}(X)$ denotes the number of nonzero entries in the matrix X .

Theorem 11.6. For all prime powers $q = p^r$ and constants $\epsilon > 0$:

There is a P^{NP} machine M such that, for infinitely many N , on input 1^N , M outputs an $N \times N$ matrix $H_N \in F_q^{N \times N}$ such that $w_2(H_N) \leq (N^{2(\log N)^{1-\epsilon}})$ over F_q .

We first prove the following folklore lemma.

Lemma 11.10. For any field F , and any matrix $A \in F^{N \times N}$, let $r = w_2(A) \leq N$. Then, for any constant $\epsilon > 0$, we have

$$R_A(\epsilon r^2) \leq N^{2\epsilon};$$

for some constant ϵ depending only on ϵ .

In other words, if $R_A(\epsilon r^2) > N^{2\epsilon}$, then we have $w_2(A) \geq \epsilon r$.

Proof. For some integer M , let B and C be matrices over F of dimensions $N \times M$ and $M \times N$, respectively, such that $A = BC$ and $\text{nnz}(B) + \text{nnz}(C) = r \leq N$. For $i, j \in [N]$, let $b_i \in F^M$ be the i -th row of B , and $c_j \in F^M$ be the j -th column of C . Hence, $A_{ij} = \sum_k b_{ik} c_{kj}$.

Now, let ϵ be a function of ϵ to be specified later, and set $m = \epsilon r^2$. Pick a hash function $P : [M] \rightarrow [m]$ uniformly at random. Next, for each b_i , we define a vector \tilde{b}_i by setting, for each $j \in [m]$:

$$(\tilde{b}_i)_j := \sum_{k \in P^{-1}(j)} b_{ik};$$

We similarly define \tilde{c}_j . Now, let \tilde{B} be the $N \times m$ matrix with the \tilde{b}_i 's as rows, and \tilde{C} be the $m \times N$ matrix with the \tilde{c}_j 's as columns. We will now argue that $\tilde{B}\tilde{C}$ approximates A well.

First, from definition, we have

$$E_{(i,j) \in [N] \times [N]} \text{nnz}(\tilde{b}_i) + \text{nnz}(\tilde{c}_j) = \frac{\text{nnz}(A) + \text{nnz}(B)}{N} = r;$$

Hence, by Markov's inequality, for at least a $1 - \epsilon/2$ fraction of the pairs $(i, j) \in [N] \times [N]$, we have $\text{nnz}(\tilde{b}_i) + \text{nnz}(\tilde{c}_j) \leq r/\epsilon$.

Fix such a pair of (i, j) , and let $I = \{k \in [M] : (b_i)_k \neq 0 \text{ or } (c_j)_k \neq 0\}$, which has size $|I| \leq r/\epsilon$. Note that if all the elements of I have distinct images under the mapping P , then $\sum_k \tilde{b}_{ik} \tilde{c}_{kj} = \sum_k b_{ik} c_{kj} = A_{ij}$. By a union bound, this happens with probability at least $1 - \epsilon/2$ over the random choice of P .

Setting $\epsilon = (\epsilon/2)^3$, we have $1 - \epsilon/2 \geq 1 - \epsilon/2$. Thus, by the probabilistic method, there is a fixed P for which $\tilde{B}\tilde{C}$ agrees with A on a $1 - \epsilon/2$ fraction of inputs, and hence $R_A(\epsilon r^2) \leq N^{2\epsilon}$. \square

Theorem 11.6 then follows by combining Lemma 11.10, Theorem 11.1, and Theorem 11.2.

11.5.3 Threshold Circuit Lower Bound for E^{NP}

We conclude this Section with a new threshold circuit lower bound for E^{NP} .

Theorem 11.7. For every $\epsilon > 0$ and prime p , there is an $\alpha > 0$ such that the class E^{NP} does not have non-uniform $AC^0[p]$ LTF $AC^0[p]$ LTF circuits of depth $\alpha(\log n = \log \log n)$ where the bottom LTF layer has $2^{O(n^\alpha)}$ gates, the rest of the circuit has polynomial size, and the middle LTF gates have fan-in $O(n^{1-2\epsilon})$.

Theorem 11.7 follows from the connection between rigid matrices and threshold circuits which we proved in Theorem 10.7 in the previous Chapter.

11.6 Algorithm for Counting Orthogonal Vectors over Finite Fields

Finally, in this Section, we give a sketch of the algorithm for F_p -# OV which we stated in Theorem 11.4 and which is needed by our construction above. The algorithm is a minor modification of the deterministic algorithm for # OV by Chan and Williams [CW16]. It makes use of the polynomial method in algorithm design, the same algorithmic technique we used earlier in Chapter 8.

11.6.1 Reduction to Prime Fields

We begin by sketching a reduction from F_{p^r} -# OV to F_p -# OV. More precisely, for a prime power $q = p^r$, we give a reduction from one instance of F_q -# $OV_{n;d}$ to a constant number of different instances of F_p -# $OV_{n;d \cdot O_r(1)}$. The reduction builds on ideas from [LPT⁺17] and [Wil18].

We first define an intermediate problem F_p -# AND- $OV_{n;d;r}$: given as input two size- n collections $A, B \subseteq (F_q^d)^r$, with $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ (so, for instance, each a_i is an r -tuple of vectors from F_q^d), the goal is to compute the number of pairs $(i, i') \in [n]^2$ such that $\langle a_{i,j}, b_{i',j} \rangle = 0$ for all $j \in [r]$.

F_q -# $OV_{n;d} \approx F_p$ -# AND- $OV_{n;d;r}$. We first show how to reduce an F_q -# $OV_{n;d}$ instance to an F_p -# AND- $OV_{n;d;r}$ instance in nearly linear time. Pick a degree- r irreducible polynomial P ; we know that F_q isomorphic to $F_p[X]/(P)$. In the calculations below, we perform the arithmetic mod P .

Suppose we have two vectors $u, v \in F_q^d$. Let $u_i = \sum_{j=0}^{r-1} u_{i,j} X^j$, and $v_i =$

$\prod_{j=0}^{r-1} u_j X^j$ for coefficients $u_j; v_j \in \mathbb{F}_p$. We have that

$$\begin{aligned} \sum_{i=1}^d u_i v_i &= \sum_{i=1}^d \sum_{j=0}^{d-1} u_j X^j \sum_{k=0}^{d-1} v_k X^k \\ &= \sum_{i=1}^d \sum_{j=0}^{d-1} \sum_{k=0}^{d-1} u_j v_k X^{j+k} \end{aligned}$$

Define the coefficients $w_{j+k} \in \mathbb{F}_p$ so that $\sum_{j+k=\ell} u_j v_k = w_\ell \pmod{P}$. The above simplifies to

$$\sum_{j=0}^{d-1} \sum_{k=0}^{d-1} u_j v_k X^{j+k} = \sum_{\ell=0}^{2d-2} w_\ell X^\ell \pmod{P}.$$

We therefore see that $u_i v_i = 0$ if and only if

$$\sum_{j=0}^{d-1} \sum_{k=0}^{d-1} u_j v_k X^{j+k} = 0 \tag{11.1}$$

for all $0 \leq \ell < 2d-2$. For each ℓ , we can build vectors $u_i^{(\ell)}$ and $v_i^{(\ell)}$ in \mathbb{F}_p^{2d} so that $(u_i^{(\ell)}; v_i^{(\ell)})$ equals the left hand side of (11.1). This transformation reduces an \mathbb{F}_p -# $OV_{n;d}$ instance to an \mathbb{F}_p -# $AND-OV_{n;dr^2;r}$ instance as desired.

\mathbb{F}_p -# $AND-OV$) \mathbb{F}_p -# OV . Now, given an \mathbb{F}_p -# $AND-OV_{n;d;r}$ instance with input collections $A; B$, we show how to reduce it to r different \mathbb{F}_p -# $OV_{n;dr+1}$ instances, again in nearly linear time.

Let $a; b \in (\mathbb{F}_p^d)^r$. For a random vector $u \in \mathbb{F}_p^r$, observe that:

If $u_i a_i; b_i = 0$ for all $i \in [r]$, then $\prod_{i=1}^r u_i a_i; b_i$ is always zero.

Otherwise, $\prod_{i=1}^r u_i a_i; b_i = 1$ with probability $1/p$.

For our reduction, we iterate over all vectors $u \in \mathbb{F}_p^r$, and sum the number of pairs $(a; b) \in A \times B$ such that

$$\prod_{i=1}^r u_i a_i; b_i = \prod_{i=1}^r u_i a_i; b_i = 1:$$

For each u , this can be written as an \mathbb{F}_p -# $OV_{n;dr+1}$ instance (via $u_i a_i; b_i = 1, u_i a_i; b_i = 0$).

For a pair $(a; b) \in A \times B$, if $u_i a_i; b_i = 0$ for all $i \in [r]$, then $(a; b)$ is never counted in the above sum. Otherwise, it is counted p^{r-1} times. Therefore, by summing up the results of all these \mathbb{F}_p -# OV instances after the reduction, dividing the result by p^{r-1} ,

and then finally subtracting the resulting number from $|A| - |B|$, we can compute the answer to the given $F_{n,d,r}$ -# AND-OV instance.

11.6.2 Algorithm for Prime Fields

In this subsection, we give a self-contained exposition of the F_p -# OV algorithm which is implicit in [CW16]. We will make use of the polynomial method in algorithm design, and in particular, we will use Lemma 8.2 from Chapter 8 for quickly evaluating a sparse polynomials on many inputs by using fast matrix multiplication. In [CW16], the deterministic # OV algorithm works by combining two key technical tools: small-biased sets and modulus-amplifying polynomials. We won't need small-biased sets here as we only aim to solve F_p -# OV. We first recall the definition of modulus-amplifying polynomials.

Lemma 11.11 (Modulus-Amplifying Polynomial [Yao90, BT94]). For all integers $\ell \geq 1$, there is a polynomial F_ℓ over \mathbb{Z} of degree $(2^\ell - 1)$ with $O(\ell)$ -bit coefficients such that for all integers $m \geq 1$ and all $a \in \mathbb{Z}$:

- (1) if $a \equiv 0 \pmod{m}$, then $F_\ell(a) \equiv 0 \pmod{m^\ell}$, and
- (2) if $a \equiv 1 \pmod{m}$, $F_\ell(a) \equiv 1 \pmod{m^\ell}$.

Now we are ready to prove Theorem 11.4 when the modulus is a prime. The case when q is a prime power then follows using the reduction from Section 11.6.1.

Theorem 11.8. For all primes p , there is an $n^{2^{(1 + \log(d = \log n))}}$ time deterministic algorithm for F_p -# $OV_{n,d}$, when $d = n^{o(1)}$.

Proof. Let ℓ be a parameter to be specified later. Let X, Y be two collections of $\ell = 4$ vectors from F_p^d . We define the polynomial

$$P(X; Y) := \sum_{(x;y) \in X \times Y} (1 - F_\ell(\langle hx; yi \rangle^{p-1}));$$

where F_ℓ is the modulus-amplifying polynomial from Lemma 11.11. Hence,

$$1 - F_\ell(\langle hx; yi \rangle^{p-1}) \equiv \begin{cases} 1 \pmod{p^\ell} & \text{when } \langle hx; yi \rangle \equiv 0 \pmod{p}; \\ 0 \pmod{p^\ell} & \text{when } \langle hx; yi \rangle \not\equiv 0 \pmod{p}; \end{cases}$$

Let us count the number M of monomials in $F_\ell(\langle hx; yi \rangle^{p-1}) = F_\ell((x_1 y_1 + x_2 y_2 + \dots + x_d y_d)^{p-1})$ when it is expanded and simplified. F_ℓ is a polynomial of degree $(2^\ell - 1)(p - 1)$ in $x; y \in F_p^d$. In particular, since we are working over F_p , we may simplify F_ℓ so that each of the $2d$ input variables has individual degree at most $p - 1$ in any given monomial. Thus, using the simple bound that no monomial depends on more variables than the degree of the polynomial, combined with the fact that the

Bibliography

- [ABFR94] James Aspnes, Richard Beigel, Merrick Furst, and Steven Rudich. The expressive power of voting polynomials. *Combinatorica*, 14(2):135–148, 1994.
- [ABG⁺14] Adi Akavia, Andrej Bogdanov, Siyao Guo, Akshay Kamath, and Alon Rosen. Candidate weak pseudorandom functions in AC0MOD2. In *ITCS*, pages 251–260, 2014.
- [AC09] Nir Ailon and Bernard Chazelle. The fast Johnson Lindenstrauss transform and approximate nearest neighbors. *SIAM J. Comput.*, 39(1):302–322, 2009.
- [AC19] Josh Alman and Lijie Chen. Efficient construction of rigid matrices using an np oracle. In *FOCS*, to appear 2019.
- [ACR⁺10] Andris Ambainis, Andrew M Childs, Ben W Reichardt, Robert Špalek, and Shengyu Zhang. Any and-or formula of size n can be evaluated in time $n^{1-2^{-\Omega(1)}}$ on a quantum computer. *SIAM J. Computing*, 39(6):2513–2530, 2010.
- [ACW16] Josh Alman, Timothy Chan, and Ryan Williams. Polynomial representations of threshold functions and algorithmic applications. In *FOCS*, pages 467–476, 2016.
- [AFLG15] Andris Ambainis, Yuval Filmus, and François Le Gall. Fast matrix multiplication: limitations of the Coppersmith-Winograd method. In *STOC*, pages 585–593, 2015.
- [AG94] Eric Allender and Vivek Gore. A uniform circuit lower bound for the permanent. *SIAM J. Computing*, 23(5):1026–1049, 1994.
- [AI06] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *FOCS*, pages 459–468, 2006.
- [AINR14] Alexandr Andoni, Piotr Indyk, Huy L Nguyen, and Ilya Razenshteyn. Beyond locality-sensitive hashing. In *SODA*, pages 1018–1028, 2014.

- [AIP06] Alexandr Andoni, Piotr Indyk, and Mihai Patrascu. On the optimality of the dimensionality reduction method. InFOCS, pages 449 458, 2006.
- [AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. Primes is in P. Annals of mathematics pages 781 793, 2004.
- [AKW90] Noga Alon, Mauricio Karchmer, and Avi Wigderson. Linear circuits over $GF(2)$. SIAM Journal on Computing 19(6):1064 1067, 1990.
- [Alm19a] Josh Alman. An illuminating algorithm for the light bulb problem. In SOSA, pages 2:1 2:11, 2019.
- [Alm19b] Josh Alman. Limits on the universal method for matrix multiplication. In CCC, pages 12:1 12:24, 2019.
- [Alo90] Noga Alon. On the rigidity of an Hadamard matrix. Manuscript. See [Juk01, Section 15.1.2], 1990.
- [ALSV13] Noga Alon, Troy Lee, Adi Shraibman, and Santosh Vempala. The approximate rank of a matrix and its algorithmic applications. InSTOC, pages 675 684, 2013.
- [AM17] Josh Alman and Dylan McKay. Theoretical foundations of team match-making. In AAMAS, pages 1073 1081, 2017.
- [AMY16] Noga Alon, Shay Moran, and Amir Yehudayo. Sign rank versus VC dimension. InCOLT, pages 47 80, 2016.
- [And05] Alexandr Andoni. Approximate nearest neighbor problem in high dimensions. Master's thesis, MIT, 2005.
- [AR15] Alexandr Andoni and Ilya Razenshteyn. Optimal data-dependent hashing for approximate near neighbors. InSTOC, pages 793 801, 2015.
- [ASU13] Noga Alon, Amir Shpilka, and Christopher Umans. On sun owers and matrix multiplication. Computational Complexity 22(2):219 243, 2013.
- [AW09] Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. ACM Transactions on Computation Theory 1(1):2:1 2:54, 2009.
- [AW15] Josh Alman and Ryan Williams. Probabilistic polynomials and Hamming nearest neighbors. InFOCS, pages 136 150, 2015.
- [AW17] Josh Alman and Ryan Williams. Probabilistic rank and matrix rigidity. In STOC, pages 641 652, 2017.
- [AW18a] Josh Alman and Virginia Vassilevska Williams. Further limitations of the known approaches for matrix multiplication. InITCS, pages 25:1 25:15, 2018.

- [AW18b] Josh Alman and Virginia Vassilevska Williams. Limits on all known (and some unknown) approaches to matrix multiplication. InFOCS, pages 580–591, 2018.
- [AWY15] Amir Abboud, Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. InSODA, pages 218–230, 2015.
- [BCC⁺ 17a] Jonah Blasiak, Thomas Church, Henry Cohn, Joshua A Grochow, Eric Naslund, William F Sawin, and Chris Umans. On cap sets and the group-theoretic approach to matrix multiplication. Discrete Analysis 2017(3):1–27, 2017.
- [BCC⁺ 17b] Jonah Blasiak, Thomas Church, Henry Cohn, Joshua A Grochow, and Chris Umans. Which groups are amenable to proving exponent two for matrix multiplication? arXiv preprint arXiv:1712.02302, 2017.
- [BCS13] Peter Bürgisser, Michael Clausen, and Mohammad A Shokrollahi. Algebraic complexity theory Springer Science & Business Media, 2013.
- [BdW01] Harry Buhrman and Ronald de Wolf. Communication complexity lower bounds by polynomials. InCCC, pages 120–130, 2001.
- [Bei95] Richard Beigel. The polynomial method in circuit complexity. InStructure in Complexity Theory Conference, pages 82–95, 1995.
- [BFS86] László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. InFOCS, pages 337–347, 1986.
- [BGH⁺ 06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. SIAM J. Comput., 36(4):889–974, 2006.
- [BGL06] Nayantara Bhatnagar, Parikshit Gopalan, and Richard J. Lipton. Symmetric polynomials over \mathbb{Z}_m and simultaneous communication protocols. J. Comput. Syst. Sci, 72(2):252–285, 2006.
- [BH74] James R Bunch and John E Hopcroft. Triangular factorization and inversion by fast matrix multiplication. Mathematics of Computation, 28(125):231–236, 1974.
- [BI13] Peter Bürgisser and Christian Ikenmeyer. Explicit lower bounds via geometric complexity theory. InSTOC, pages 141–150, 2013.
- [Bin80] Dario Bini. Border rank of a $p \times q \times 2$ tensor and the optimal approximation of a pair of bilinear forms. InICALP, pages 98–108, 1980.
- [Blä13] Markus Bläser. Fast matrix multiplication. Theory of Computing, Graduate Surveys 5:1–60, 2013.

- [BR02] Omer Barkol and Yuval Rabani. Tighter lower bounds for nearest neighbor search and related problems in the cell probe model. *JCSS*, 64(4):873–896, 2002.
- [Bro97] Andrei Z Broder. On the resemblance and containment of documents. In *SEQUENCES*, pages 21–29, 1997.
- [BRS91] Richard Beigel, Nick Reingold, and Daniel Spielman. The perceptron strikes back. In *Structure in Complexity Theory Conference*, pages 286–291, 1991.
- [BS92] Jehoshua Bruck and Roman Smolensky. Polynomial threshold functions, AC^0 functions, and spectral norms. *SIAM J. Comput.*, 21(1):33–42, 1992.
- [BT94] Richard Beigel and Jun Tarui. On ACC. *Computational Complexity* 4:350–366, 1994.
- [BV14] Eli Ben-Sasson and Emanuele Viola. Short PCPs with projection queries. In *ICALP*, pages 163–173, 2014.
- [CCGL99] Amit Chakrabarti, Bernard Chazelle, Benjamin Gum, and Alexey Lvov. A lower bound on the complexity of approximate nearest-neighbor searching on the hamming cube. In *STOC*, pages 305–311, 1999.
- [CFL85] Ashok K. Chandra, Steven Fortune, and Richard J. Lipton. Unbounded fan-in circuits and associative functions. *JCSS*, 30(2):222–234, 1985.
- [CGJ⁺16] Mahdi Cheraghchi, Elena Grigorescu, Brendan Juba, Karl Wimmer, and Ning Xie. AC^0 MOD₂ lower bounds for the boolean inner product. In *ICALP*, pages 35:1–35:14, 2016.
- [Cha02] Moses S Charikar. Similarity estimation techniques from rounding algorithms. In *STOC*, 2002.
- [Cha18] Timothy M. Chan. Applications of Chebyshev polynomials to low-dimensional computational geometry. *Journal of Computational Geometry*, 9(2):3–20, 2018.
- [Che99] Pafnuty L. Chebyshev. Sur l'interpolation. In A. Marko and N. Sonin, editors, *Oeuvres de P. L. Tchebychev*, volume 1, pages 539–560. Commissionnaires de L'Académie Impériale des Sciences, 1899.
- [CIP06] Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. A duality between clause width and clause density for SAT. In *CCC*, pages 252–260, 2006.
- [CKL13] Ho Yee Cheung, Tsz Chiu Kwok, and Lap Chi Lau. Fast matrix rank algorithms and applications. *J. ACM*, 60(5):31:1–31:25, 2013.

- [CKSU05] Henry Cohn, Robert Kleinberg, Balazs Szegedy, and Christopher Umans. Group-theoretic algorithms for matrix multiplication. In FOCS, pages 379–388, 2005.
- [CLP17] Ernie Croot, Vsevolod F Lev, and Péter Pál Pach. Progression-free sets in \mathbb{Z}_4^n are exponentially small. *Annals of Mathematics* 185(1):331–337, 2017.
- [CLS18] Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In STOC, pages 938–942, 2018.
- [Cod00] Bruno Codenotti. Matrix rigidity. *Linear Algebra and its Applications* 304(1-3):181–192, 2000.
- [Cop82] Don Coppersmith. Rapid multiplication of rectangular matrices. *SIAM J. Comput.*, 11(3):467–471, 1982.
- [CP19] Shiteng Chen and Periklis A Papakonstantinou. Depth reduction for composites. *SIAM Journal on Computing* 48(2):668–686, 2019.
- [CR04] Amit Chakrabarti and Oded Regev. An optimal randomised cell probe lower bound for approximate nearest neighbour searching. In FOCS, pages 473–482, 2004.
- [CS15] Ruiwen Chen and Rahul Santhanam. Improved algorithms for sparse MAX-SAT and MAX- k -CSP. In SAT, pages 33–45, 2015.
- [CSS16] Ruiwen Chen, Rahul Santhanam, and Srikanth Srinivasan. Average-case lower bounds and satisfiability algorithms for small threshold circuits. In CCC, pages 1:1–1:35, 2016.
- [CSV84] Ashok K. Chandra, Larry Stockmeyer, and Uzi Vishkin. Constant depth reducibility. *SIAM J. Computing*, 13(2):423–439, 1984.
- [CU03] Henry Cohn and Christopher Umans. A group-theoretic approach to fast matrix multiplication. In FOCS, pages 438–449, 2003.
- [CU13] Henry Cohn and Christopher Umans. Fast matrix multiplication using coherent configurations. In SODA, pages 1074–1086, 2013.
- [CVZ18] Matthias Christandl, Péter Vrana, and Jeroen Zuiddam. Universal points in the asymptotic spectrum of tensors. In STOC, pages 289–296, 2018.
- [CVZ19] Matthias Christandl, Péter Vrana, and Jeroen Zuiddam. Barriers for fast matrix multiplication from irreversibility. In CCC, pages 26:1–26:17, 2019.

- [CW82] Don Coppersmith and Shmuel Winograd. On the asymptotic complexity of matrix multiplication. *SIAM J. Comput.*, 11(3):472–492, 1982.
- [CW90] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of symbolic computation* 9(3):251–280, 1990.
- [CW16] Timothy M. Chan and Ryan Williams. Deterministic APSP, orthogonal vectors, and more: Quickly derandomizing Razborov Smolensky. In *SODA*, pages 1246–1255, 2016.
- [CW19a] Lijie Chen and Ruosong Wang. Classical algorithms from quantum and arthur-merlin communication protocols. *10th Innovations in Theoretical Computer Science* pages 23:1–23:20, 2019.
- [CW19b] Lijie Chen and Ryan Williams. Stronger connections between circuit analysis and circuit lower bounds, via PCPs of proximity. In *CCC*, pages 19:1–19:43, 2019.
- [DE17] Zeev Dvir and Benjamin Edelman. Matrix rigidity and the croot-levpach lemma. *arXiv preprint arXiv:1708.01646*, 2017.
- [Des07] Amit Jayant Deshpande. Sampling-based algorithms for dimension reduction. PhD thesis, Massachusetts Institute of Technology, 2007.
- [DGW19] Zeev Dvir, Alexander Golovnev, and Omri Weinstein. Static data structure lower bounds imply rigidity. In *STOC*, pages 967–978, 2019.
- [DL19] Zeev Dvir and Allen Liu. Fourier and Circulant Matrices Are Not Rigid. In *CCC*, pages 17:1–17:23, 2019.
- [DS13] A.M. Davie and A. J. Stothers. Improved bound for complexity of matrix multiplication. *Proceedings of the Royal Society of Edinburgh, Section: A Mathematics* 143:351–369, 4 2013.
- [Dub10] Moshe Dubiner. Bucketing coding and information theory for the statistical high-dimensional nearest-neighbor problem. *IEEE Transactions on Information Theory*, 56(8):4166–4179, 2010.
- [Dvi16] Zeev Dvir. On the non-rigidity of generating matrices of good codes. Writeup by Oded Goldreich. Available at <http://www.wisdomweizmann.ac.il/~oded/MC/209.html>, October 30 2016.
- [DW06a] Evgeny Dantsin and Alexander Wolpert. MAX-SAT for formulas with constant clause density can be solved faster than $O(2^n)$ time. In *SAT*, pages 266–276, 2006.
- [dW06b] Ronald de Wolf. Lower bounds on matrix rigidity via a quantum argument. In *ICALP*, volume 4051, pages 62–71, 2006.

- [EG17] Jordan S Ellenberg and Dion Gijswijt. On large subsets of \mathbb{F}_q^n with no three-term arithmetic progression. *Annals of Mathematics* 185(1):339–343, 2017.
- [EGOW18] Klim Efremenko, Ankit Garg, Rafael Oliveira, and Avi Wigderson. Barriers for rank methods in arithmetic complexity. In *ITCS*, pages 1:1–1:19, 2018.
- [FF93] Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM J. Comput.*, 22(5):994–1005, 1993.
- [FKL⁺01] 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 *FSTTCS*, pages 171–182, 2001.
- [For02] Jürgen Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *J. Comput. System. Sci.* 65(4):612–625, 2002.
- [Fri93] Joel Friedman. A note on matrix rigidity. *Combinatorica*, 13(2):235–239, 1993.
- [GHK⁺12] Anna Gál, Kristoffer Arnsfelt Hansen, Michal Koucký, Pavel Pudlák, and Emanuele Viola. Tight bounds on computing error-correcting codes by bounded-depth circuits with arbitrary gates. In *STOC*, pages 479–494, 2012.
- [GL01] Ben Gum and Richard J Lipton. Cheaper by the dozen: Batched algorithms. In *SDM*, pages 1–11, 2001.
- [GPW16] Mika Göös, Toniann Pitassi, and Thomas Watson. Zero-information protocols and unambiguity in arthur-merlin communication. *Algorithmica*, 76(3):684–719, 2016.
- [GPW18] Mika Göös, Toniann Pitassi, and Thomas Watson. The landscape of communication complexity classes. *Computational Complexity* 27(2):245–304, 2018.
- [GQ19] Joshua A Grochow and Youming Qiao. Isomorphism problems for tensors, groups, and cubic forms: completeness and reduction. *arXiv preprint arXiv:1907.00309* 2019.
- [Gri] D. Yu. Grigor'ev. Unpublished work. Cited in [KR98].
- [GST03] Dan Gutfreund, Ronen Shaltiel, and Amnon Ta-Shma. Uniform hardness versus randomness tradeoffs for arthur-merlin games. *Computational Complexity*, 12(3-4):85–130, 2003.

- [GT16] Oded Goldreich and Avishay Tal. Matrix rigidity of random toeplitz matrices. In STOC, pages 91 104, 2016.
- [HIM12] Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory of Computing* 8(1):321 350, 2012.
- [Hir03] Mika Hirvensalo. Studies on Boolean Functions Related to Quantum Computing PhD thesis, University of Turku, 2003.
- [HMP⁺ 93] András Hajnal, Wolfgang Maass, Pavel Pudlák, Mario Szegedy, and György Turán. Threshold circuits of bounded depth. *J. Comput. Syst. Sci.*, 46(2):129 154, 1993.
- [HNO08] Nicholas J. A. Harvey, Jelani Nelson, and Krzysztof Onak. Sketching and streaming entropy via approximation theory. In FOCS, pages 489 498, 2008.
- [Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13 30, 1963.
- [HY09] Pavel Hrubeč and Amir Yehudayoff. Arithmetic complexity in algebraic extensions, 2009.
- [IKW02] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. In search of an easy witness: exponential time vs. probabilistic polynomial time. *J. Comput. Syst. Sci.*, 65(4):672 694, 2002.
- [IM98] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In STOC, pages 604 613, 1998.
- [Ind04] Piotr Indyk. Nearest neighbors in high-dimensional spaces. In *Handbook of Discrete and Computational Geometry, Second Edition*, pages 877 892. Chapman and Hall, 2nd edition, 2004.
- [IPS13] Russell Impagliazzo, Ramamohan Paturi, and Stefan Schneider. A satisfiability algorithm for sparse depth two threshold circuits. In FOCS, pages 479 488, 2013.
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512 530, 2001.
- [JMV15] Hamid Jahanjou, Eric Miles, and Emanuele Viola. Local reductions. In ICALP, pages 749 760, 2015.

- [Juk01] Stasys Jukna. *Extremal Combinatorics, With Applications in Computer Science* EATCS Series. Springer, 2001.
- [KK19] Matti Karppa and Petteri Kaski. Probabilistic tensors and opportunistic boolean matrix multiplication. In *SODA*, pages 496–515, 2019.
- [KKK16] Matti Karppa, Petteri Kaski, and Jukka Kohonen. A faster subquadratic algorithm for finding outlier correlations. In *SODA*, pages 1288–1305, 2016.
- [KKKÓC16] Matti Karppa, Petteri Kaski, Jukka Kohonen, and Pádraig Ó Catháin. Explicit correlation amplifiers for finding outlier correlations in deterministic subquadratic time. In *ESA*, pages 52:1–52:17, 2016.
- [Kle97] Jon M Kleinberg. Two algorithms for nearest-neighbor search in high dimensions. In *STOC*, pages 599–608, 1997.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity* Cambridge University Press, 1997.
- [KOR00] Eyal Kushilevitz, Rafail Ostrovsky, and Yuval Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. *SIAM Journal on Computing* 30(2):457–474, 2000.
- [KOS04] 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.
- [KR98] B. S. Kashin and A. A. Razborov. Improved lower bounds on the rigidity of Hadamard matrices. *Matematicheskie Zametki* 63(4):535–540, 1998. (in Russian).
- [KS01] Adam R. Klivans and Rocco Servedio. Learning DNF in time $2^{O(n^{1/3})}$. In *STOC*, pages 258–265, 2001.
- [KS10] Adam R Klivans and Alexander A Sherstov. Lower bounds for agnostic learning via approximate rank. *Computational Complexity* 19(4):581–604, 2010.
- [KSS18] Robert Kleinberg, Will Sawin, and David Speyer. The growth rate of tri-colored sum-free sets. *Discrete Analysis* 12, 2018.
- [KV19] Mrinal Kumar and Ben Lee Volk. Lower bounds for matrix factorization. arXiv preprint arXiv:1904.01182, 2019.
- [KvM02] Adam R. Klivans and Dieter van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM J. Comput.*, 31(5):1501–1526, 2002.

- [KW16] Daniel M. Kane and Ryan Williams. Super-linear gate and super-quadratic wire lower bounds for depth-two and depth-three threshold circuits. In STOC, pages 633 643, 2016.
- [Lan17] Joseph M Landsberg. Geometry and complexity theory, volume 169. Cambridge University Press, 2017.
- [LG14] François Le Gall. Powers of tensors and fast matrix multiplication. In ISSAC, pages 296 303, 2014.
- [LGU17] François Le Gall and Florent Urrutia. Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor. In SODA, pages 1029 1046, 2017.
- [LN90] Nathan Linial and Noam Nisan. Approximate inclusion-exclusion. *Combinatorica*, 10(4):349 365, 1990.
- [Lok00] Satyanarayana V. Lokam. On the rigidity of vandermonde matrices. *Theoretical Computer Science* 237(1-2):477 483, 2000.
- [Lok01] Satyanarayana V. Lokam. Spectral methods for matrix rigidity with applications to size-depth tradeoffs and communication complexity. *Journal of Computer and System Sciences* 63(3):449 473, 2001.
- [Lok06] Satyanarayana V. Lokam. Quadratic lower bounds on matrix rigidity. In TAMC, volume 3959, pages 295 307. Springer, 2006.
- [Lok14] Satyanarayana V. Lokam. Exercises on matrix rigidity. Simons Institute for Theory of Computing. Available at <https://simons.berkeley.edu/sites/default/files/docs/1738/exercises.pdf>, 2014.
- [Lov11] Shachar Lovett. Computing polynomials with few multiplications. *Theory of Computing* 7(1):185 188, 2011.
- [LPT⁺17] Daniel Lokshtanov, Ramamohan Paturi, Suguru Tamaki, Ryan Williams, and Huacheng Yu. Beating brute force for systems of polynomial equations over finite fields. In SODA, pages 2190 2202, 2017.
- [LS09] Nathan Linial and Adi Shraibman. Learning complexity vs communication complexity. *Combinatorics, Probability & Computing* 18(1-2):227 245, 2009.
- [LTV03] JM Landsberg, J. Taylor, and N.K. Vishnoi. The geometry of matrix rigidity. Available at <https://smartech.gatech.edu/handle/1853/6514>, 2003.
- [Lup56] Oleg B Lupanov. On rectifier and switching-and-rectifier schemes. *Dokl. Akad. Nauk SSSR* 111(6):1171 1174, 1956.

- [Mat91] Jirí Matoušek. Computing dominances in E^n . *Inf. Process. Lett.*, 38(5):277–278, 1991.
- [Mat08] Jirí Matoušek. On variants of the Johnson Lindenstrauss lemma. *Random Struct. Algorithms*, 33(2):142–156, 2008.
- [Mid05] Gatis Midrijanis. Three lines proof of the lower bound for the matrix rigidity. *arXiv preprint arXiv:cs/0506081*, 2005.
- [MKZ09] Kerui Min, Ming-Yang Kao, and Hong Zhu. The closest pair problem under the hamming metric. In *Computing and Combinatorics* pages 205–214. Springer, 2009.
- [MP69] Marvin L Minsky and Seymour A Papert. *Perceptrons: An Introduction to Computational Geometry* MIT press Boston, MA:, 1969.
- [MPS16] Daniel Moeller, Ramamohan Paturi, and Stefan Schneider. Sub-quadratic algorithms for succinct stable matching. In *CSR*, pages 294–308, 2016.
- [MS82] Kurt Mehlhorn and Erik M Schmidt. Las vegas is better than determinism in vlsi and distributed computing. In *STOC*, pages 330–337, 1982.
- [MT98] Alexis Maciel and Denis Thérien. Threshold circuits of small majority-depth. *Information and Computation*, 146(1):55–83, 1998.
- [MT99] Alexis Maciel and Denis Thérien. Efficient threshold circuits for power series. *Information and Computation*, 152(1):62–73, 1999.
- [MV05] Peter Bro Miltersen and N. V. Vinodchandran. Derandomizing arthur-merlin games using hitting sets. *Computational Complexity* 14(3):256–279, 2005.
- [MW18] Cody Murray and Ryan Williams. Circuit lower bounds for nondeterministic quasi-polytime: an easy witness lemma for NP and NQP. In *STOC*, pages 890–901, 2018.
- [Nis] Noam Nisan. Unpublished work. Cited in [\[KR98\]](#).
- [NS94] Noam Nisan and Mario Szegedy. On the degree of boolean functions as real polynomials. *Computational Complexity* 4(4):301–313, 1994.
- [NS17] Eric Naslund and Will Sawin. Upper bounds for sun over-free sets. In *Forum of Mathematics, Sigma* volume 5, 2017.
- [NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *Comput. Syst. Sci*, 49(2):149–167, 1994.

- [OS10] Ryan O’Donnell and Rocco A. Servedio. New degree bounds for polynomial threshold functions. *Combinatorica*, 30(3):327–358, 2010.
- [OS17] Igor Carboni Oliveira and Rahul Santhanam. Pseudodeterministic constructions in subexponential time. In *STOC*, pages 665–677, 2017.
- [Pan78] V. Y. Pan. Strassen’s algorithm is not optimal. In *FOCS*, pages 166–176, 1978.
- [Pan06] Rina Panigrahy. Entropy based nearest neighbor search in high dimensions. In *SODA*, pages 1186–1195, 2006.
- [Par19] Orr Paradise. Smooth and strong PCPs. *ECCC preprint TR19-023*, 2019.
- [Pat92] Ramamohan Paturi. On the degree of polynomials that approximate symmetric boolean functions. In *STOC*, pages 468–474, 1992.
- [Pat08] Mihai Patrascu. *Lower bound techniques for data structures*. PhD thesis, Massachusetts Institute of Technology, 2008.
- [PRR95] Ramamohan Paturi, Sanguthevar Rajasekaran, and John Reif. The light bulb problem. *Information and Computation*, 117(2):187–192, 1995.
- [PS88] P. Pudlak and P. Savicky. Private communication. Cited in [Raz89], 1988.
- [Pud94] Pavel Pudlak. Large communication in constant depth circuits. *Combinatorica*, 14(2):203–216, 1994.
- [Ras16] Cyrus Rashtchian. Bounded matrix rigidity and John’s theorem. *ECCC preprint TR16-093*, 2016.
- [Raz87] A. A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.
- [Raz89] A. A. Razborov. On rigid matrices (in Russian). Manuscript can be found at <http://people.cs.uchicago.edu/~razborov/files/rigid.pdf>, 1989.
- [RS10] Alexander A. Razborov and Alexander A. Sherstov. The sign-rank of AC^0 . *SIAM J. Comput.*, 39(5):1833–1855, 2010.
- [RTS00] Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24, 2000.
- [Rub18] Aviad Rubinfeld. Hardness of approximate nearest neighbor search. In *STOC*, 2018.

- [Saw18] Will Sawin. Bounds for matchings in nonabelian groups. *The Electronic Journal of Combinatorics*, 25(4):4–23, 2018.
- [Sch81] A. Schönhage. Partial and total matrix multiplication. *SIAM J. Comput.*, 10(3):434–455, 1981.
- [Sch05] Rainer Schuler. An algorithm for the satisfiability problem of formulas in conjunctive normal form. *J. Algorithms*, 54(1):40–44, 2005.
- [She08] Alexander A Sherstov. Approximate inclusion-exclusion for arbitrary symmetric functions. In *CCC*, pages 112–123, 2008.
- [She13] Alexander A. Sherstov. Making polynomials robust to noise. *Theory of Computing*, 9:593–615, 2013.
- [She14] Alexander A. Sherstov. Breaking the Minsky–Papert barrier for constant-depth circuits. In *STOC*, pages 223–232, 2014.
- [Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *STOC*, pages 77–82, 1987.
- [Spi96] Daniel A. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Trans. Information Theory*, 42(6):1723–1731, 1996.
- [Sri13] Srikanth Srinivasan. On improved degree lower bounds for polynomial approximation. In *FSTTCS*, pages 201–212, 2013.
- [SS42] Raphaël Salem and Donald C Spencer. On sets of integers which contain no three terms in arithmetical progression. *Proceedings of the National Academy of Sciences*, 28(12):561–563, 1942.
- [SS96] Victor Shoup and Roman Smolensky. Lower bounds for polynomial evaluation and interpolation problems. *Computational Complexity*, 6(4):301–311, 1996.
- [SSS95] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff–Hoeffding bounds for applications with limited independence. *SIAM J. Discrete Mathematics*, 8(2):223–250, 1995.
- [SSS97] M.A. Shokrollahi, D.A. Spielman, and V. Stemann. A remark on matrix rigidity. *Information Processing Letters*, 64(6):283–285, 1997.
- [SST15] Takayuki Sakai, Kazuhisa Seto, and Suguru Tamaki. Solving sparse instances of Max SAT via width reduction and greedy restriction. *Theory Comput. Syst.*, 57(2):426–443, 2015.
- [SSTT15a] Takayuki Sakai, Kazuhisa Seto, Suguru Tamaki, and Junichi Teruyama. Improved exact algorithms for mildly sparse instances of Max SAT. In *IPEC*, pages 90–101, 2015.

- [SSTT15b] Takayuki Sakai, Kazuhisa Seto, Suguru Tamaki, and Junichi Teruyama. A satisfiability algorithm for depth-2 circuits with a symmetric gate at the top and AND gates at the bottom. *ECCC preprint TR15-136*, 2015.
- [Str69] Volker Strassen. Gaussian elimination is not optimal. *Numerische mathematik*, 13(4):354–356, 1969.
- [Str73] Volker Strassen. Vermeidung von divisionen. *Journal für die reine und angewandte Mathematik*, 264:184–202, 1973.
- [Str86] Volker Strassen. The asymptotic spectrum of tensors and the exponent of matrix multiplication. In *FOCS*, pages 49–54, 1986.
- [Str87] Volker Strassen. Relative bilinear complexity and matrix multiplication. *J. reine angew. Math. (Crelles Journal)*, 375–376:406–443, 1987.
- [Str91] Volker Strassen. Degeneration and complexity of bilinear maps: some asymptotic spectra. *Crelles J. Reine Angew. Math*, 413:127–180, 1991.
- [SV12] Rocco A. Servedio and Emanuele Viola. On a special case of rigidity. *ECCC preprint TR12-144*, 2012.
- [SV13] Sushant Sachdeva and Nisheeth K Vishnoi. Faster algorithms via approximation theory. *Theoretical Computer Science*, 9(2):125–210, 2013.
- [Sze75] Gabor Szegő. *Orthogonal Polynomials*. American Mathematical Society, 1975.
- [Tam16] Suguru Tamaki. A satisfiability algorithm for depth two circuits with a sub-quadratic number of symmetric and threshold gates. *ECCC preprint TR16-100*, 2016.
- [Tao16] Terence Tao. A symmetric formulation of the croot-lev-pach-ellenberg-gijswijt capset bound. <https://terrytao.wordpress.com/2016/05/18/a-symmetric-formulation-of-the-croot-lev-pach-ellenberg-gijswijt-capset-bound/>, 2016.
- [Tar93] Jun Tarui. Probabilistic polynomials, AC0 functions and the polynomial-time hierarchy. *Theor. Comput. Sci.*, 113(1):167–183, 1993.
- [TS16] Terence Tao and Will Sawin. Notes on the “slice rank” of tensors. <https://terrytao.wordpress.com/2016/08/24/notes-on-the-slice-rank-of-tensors/>, 2016.
- [Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *MFCS*, pages 162–176, 1977.
- [Val88] Leslie G Valiant. Functionality in neural nets. In *AAAI*, 1988.

- [Val15] Gregory Valiant. Finding correlations in subquadratic time, with applications to learning parities and the closest pair problem. *J. ACM*, 62(2):13, 2015.
- [Vol99] Heribert Vollmer. *Introduction to circuit complexity: a uniform approach*. Springer, 1999.
- [Wil05] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.
- [Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmith-winograd. In *STOC*, pages 887–898, 2012.
- [Wil13] Ryan Williams. Improving exhaustive search implies superpolynomial lower bounds. *SIAM J. Comput.*, 42(3):1218–1244, 2013.
- [Wil14a] Ryan Williams. Faster all-pairs shortest paths via circuit complexity. In *STOC*, pages 664–673, 2014.
- [Wil14b] Ryan Williams. New algorithms and lower bounds for circuits with linear threshold gates. In *STOC*, pages 194–202, 2014.
- [Wil14c] Ryan Williams. Nonuniform ACC circuit lower bounds. *J. ACM*, 61(1):2, 2014.
- [Wil14d] Ryan Williams. The polynomial method in circuit complexity applied to algorithm design (invited talk). In *FSTTCS*, pages 47–60, 2014.
- [Wil18] Ryan Williams. Counting solutions to polynomial systems via reductions. In *SOSA*, pages 6:1–6:15, 2018.
- [Win71] Shmuel Winograd. On multiplication of 2×2 matrices. *Linear algebra and its applications*, 4(4):381–388, 1971.
- [Wun12] Henning Wunderlich. On a theorem of razborov. *Computational Complexity*, 21(3):431–477, 2012.
- [WY14] Ryan Williams and Huacheng Yu. Finding orthogonal vectors in discrete structures. In *SODA*, pages 1867–1877, 2014.
- [Yao83] Andrew C. Yao. Lower bounds by probabilistic arguments. In *FOCS*, pages 420–428, 1983.
- [Yao90] Andrew C. Yao. On ACC and threshold circuits. In *FOCS*, pages 619–627, 1990.
- [Yat37] F. Yates. The design and analysis of factorial experiments. *Technical Communication No. 35, Commonwealth Bureau of Soil Science, Harpenden, UK*, 1937.

- [Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends in Theoretical Computer Science*, 6(3):139–255, 2012.
- [Zák83] Stanislav Zák. A Turing machine time hierarchy. *Theor. Comput. Sci.*, 26:327–333, 1983.