On the Approximability of Injective Tensor Norm

Vijay Bhattiprolu

CMU-CS-19-115

June 28, 2019

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Venkatesan Guruswami, Chair
Anupam Gupta
David Woodruff
Madhur Tulsiani, TTIC

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.*

This research was sponsored by the National Science Foundation under award CCF-1526092, and by the National Science Foundation under award 1422045. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

Keywords: Injective Tensor Norm, Operator Norm, Hypercontractive Norms, Sum of Squares Hierarchy, Convex Programming, Continuous Optimization, Optimization over the Sphere, Approximation Algorithms, Hardness of Approximation

Abstract

The theory of approximation algorithms has had great success with combinatorial optimization, where it is known that for a variety of problems, algorithms based on semidefinite programming are optimal under the unique games conjecture. In contrast, the approximability of most continuous optimization problems remains unresolved.

In this thesis we aim to extend the theory of approximation algorithms to a wide class of continuous optimization problems captured by the injective tensor norm framework. Given an order-d tensor T, and symmetric convex sets $C_1, \ldots C_d$, the injective tensor norm of T is defined as

$$\sup_{x^i \in C_i} \langle T, x^1 \otimes \cdots \otimes x^d \rangle,$$

Injective tensor norm has manifestations across several branches of computer science, optimization and analysis. To list some examples, it has connections to maximum singular value, max-cut, Grothendieck's inequality, non-commutative Grothendieck inequality, quantum information theory, k-XOR, refuting random constraint satisfaction problems, tensor PCA, densest-k-subgraph, and small set expansion. So a general theory of its approximability promises to be of broad scope and applicability.

We study various important special cases of the problem (through the lens of convex optimization and the sum of squares (SoS) hierarchy) and obtain the following results:

- We obtain the first NP-hardness of approximation results for hypercontractive norms. Specifically, we prove inapproximability results for computing the $p \to q$ operator norm (which is a special case of injective norm involving two convex sets) when $p \leq q$ and $2 \notin [p, q]$. Towards the goal of obtaining strong inapproximability results for $2 \to q$ norm when $q > 2$, we give random label cover (for which polynomial level SoS gaps are available) based hardness results for mixed norms, i.e., $2 \to \ell_q(\ell_{q'})$ for some $2 < q, q' < \infty$.

- We obtain improved approximation algorithms for computing the $p \to q$ operator norm when $p \geq 2 \geq q$.

- We introduce the technique of weak decoupling inequalities and use it to analyze the integrality gap of the SoS hierarchy for the maxima of various classes of polynomials over the sphere, namely arbitrary polynomials, polynomials with non-negative coefficients and sparse polynomials. We believe this technique is broadly applicable and could find use beyond optimization over the sphere.

  We also study how well higher levels of SoS approximate the maximum of a random polynomial over the sphere ([RRS16] concurrently obtained a similar result).

Acknowledgements

First and foremost I would like to thank my parents and sister for their undying support and patience. I owe a great deal to both of my grandfathers for encouraging me on my path and teaching me in my early years.

I am most fortunate to have been advised by Venkat Guruswami. As all of his students would confirm, he adapts to the needs of his students and I can't think of anything I wish he did differently. He has been an excellent collaborator: he had troves of questions relevant to my background and interests and would constantly make prescient comments about my research directions. Any confidence I have in being able to survive outside the nest, I owe to Venkat's style of advisory.

I spent two wonderful summers (and more) at TTI chicago hosted by Madhur Tulsiani. Madhur has always been incredibly generous with his time and knowledge, and has been virtually a second advisor to me. Thanks Madhur!

I greatly enjoyed my collaboration and friendship with Euiwoong and Mrinal. I've learned so much from both of them and I look forward to our collaborations for years to come.

I am very grateful to Assaf Naor, Prasad Raghavendra and Pravesh Kothari for their insightful and influential comments on my research directions. Special thanks to Prasad for hosting me at Berkeley for a few months.

I'd like to thank my committee (Venkat, Madhur, David Woodruff and Anupam Gupta) for taking the time and giving me valuable comments. My gratitude to Anupam, Danny and David for helping with Ph.D. program requirements.

Ryan and Anupam have given several talks that I found inspiring and they are both amongst my go-to references on pedagogical issues. Special thanks to Anupam for going out of his way many times to tell me how I can improve. Thanks are due to Boris, Tom, Po and Ravi for teaching fun courses. I am grateful to Gary for sparking my interest in spectral theory.

My five years at CMU have been the most enjoyable in my lifetime so far. A big part of this is due to the wonderful group of peers I had. While it would be impossible to list them all, some memories that stand out are: ploughing through coursework with David Wajc and Arash; working at coffee shops with Nic, Naama and Nika; working nights at GHC with Yu, Daniel and Lei; ping pong with David Wajc, John, Ainesh, Anish, Goran, Angela, Pedro, Sahil, Sai and David Kurokawa; climbing with Colin, Roie, Aria, Greg, Anson, Andrei, Ellis, and Woneui; classical music concerts with John, Yu and Ellen; board games with Sol, Ameya and many others; fun conversations with Amir, Ilqar, Laxman, Guru, Jakub, Alex, Sidhanth and David Witmer.

Last but not least I would like to thank Deb Cavlovich without whom I imagine SCS would cease functioning.

Contents

Part I

Preamble

Chapter 1

Introduction

In this thesis we will be concerned with the computational approximability of the injective tensor norm which given an order-d tensor T, and finite dimensional symmetric convex sets (i.e., $x \in C \iff -x \in C$) $C_1, \dots C_d$ is defined as

$$\|T\|_{C_1,\dots C_d} := \sup_{x^i \in C_i} \langle T, x^1 \otimes \cdots \otimes x^d \rangle .$$

To be precise, we are interested in the existence (and inexistence) of approximation algorithms with runtime polynomial in the dimension of each convex set C_i.

The injective tensor norm is very expressive and has a multitude of manifestations across branches of computer science, optimization, mathematics and physics. Many of the questions of interest surrounding these special cases are related to its approximability and inapproximability. To demonstrate this as well as to familiarize the reader with the object, we list some special cases below

Maximum Singular Value. $C_1 = \mathrm{Ball}(\ell_2^m),\ C_2 = \mathrm{Ball}(\ell_2^n)$.
This case corresponds to the maximum singular value of an $m \times n$ matrix which is well known to be exactly computable.

Max Column Norm. $C_1 = \mathrm{Ball}(\ell_p^m),\ C_2 = \mathrm{Ball}(\ell_1^n)$.
This corresponds to the maximum ℓ_{p^*} norm of a column of an $m \times n$ matrix – again exactly computable.

Grothendieck Inequality. $C_1 = \mathrm{Ball}(\ell_\infty^m),\ C_2 = \mathrm{Ball}(\ell_\infty^n)$.
Grothendieck's famous inequality implies that the natural SDP relaxation for maximizing $y^T A\, x$ where $y \in \ell_\infty^m,\ x \in \ell_\infty^n$ yields a constant factor approximation to this problem. This inequality has had a major impact on Banach space theory, computer science and quantum mechanics. See [Pis12, Pis86], [KN11], [LS$^+$09] for surveys on its applications to Banach space theory, combinatorial optimization and communication complexity respectively. Determining the precise value of Grothendieck's constant remains an outstanding open problem.

Max-Cut. $C_1 = \mathrm{Ball}(\ell_\infty^n),\ C_2 = \mathrm{Ball}(\ell_\infty^n),\ T$ **Laplacian.**

Maximizing the bilinear form of the Laplacian matrix of a graph on n vertices over $\mathrm{Ball}(\ell_\infty)$ can be shown to be equivalent to the well studied combinatorial optimization problem of finding the maximum sized cut in a graph. Goemans and Williamson's [GW95] .878 . . . approximation algorithm for this problem popularized the random hyperplane rounding algorithm and has since transformed the field of approximation algorithms.

Hypercontractive Norms. $C_1 = \mathrm{Ball}(\ell_q^m),\ C_2 = \mathrm{Ball}(\ell_p^n),\ p \le q^*$**.**

This corresponds to the operator norm of a linear operator mapping ℓ_p^n to $\ell_{q^*}^m$. When $p \le q^*$ this is referred to as a *Hypercontractive norm* and is well studied in various fields. It has connections to log-Sobolev inequalities [Gro14], certifying bounds on small-set expansion [BBH+12] and soundness proofs in Hardness of Approximation. While hypercontractive norms are unlikely to be computationally tractable (even $O(1)$ approximations), establishing NP-Hardness of $O(1)$-approximating hypercontractive norms (and related promise versions) would have important implications in quantum information theory and may also shed light on some important questions in hardness of approximation.

Non-Commutative Grothendieck Inequality. $C_1 = \mathrm{Ball}(\mathbb{S}_\infty^{m_1 \times n_1}),\ C_2 = \mathrm{Ball}(\mathbb{S}_\infty^{m_2 \times n_2})$**.**

Naor, Regev and Vidick [NRV13] made algorithmic Haagerup's [Haa85] sharp version of the non-commutative Grothendieck inequality (first established by Pisier [Pis78]) and used this to give approximation algorithms for robust principal component analysis and a generalization of the orthogonal procrustes problem. Regev and Vidick [RV15] showed how it can be used to bound the power of entanglement in quantum XOR games.

Homogeneous Polynomial Optimization $C_1 = \cdots = C_d$**.**

It can be shown for symmetric tensors T that the injective norm is within a $2^{O(d)}$ factor of $\max_{x \in C_1} |\langle T, x^{\otimes d} \rangle|$. Therefore injective tensor norm is closely related to the problem of maximizing (the magnitude) of a homogeneous degree-d polynomial over a convex set. the expected suprema of a random homogeneous polynomial over convex sets like the sphere and hypercube have been extensively studied in the statistical physics community [AC+17].

Optimization over the hypercube. $C_1 = \ldots C_d = \mathrm{Ball}(\ell_\infty^n)$**.**

Here we note a connection to the fundamental constraint satisfaction problem XOR. For any instance of d-XOR with m constraints there is a homogeneous multilinear polynomial p with m non-zero monomials and degree-d such that the number of constraints satisfied by an assignment $x \in \{\pm 1\}^n$ is precisely $m/2 + p(x)$. Therefore maximizing $|p(x)|$ over $\{\pm 1\}^n$ is precisely $\max\{\mathrm{SAT}(x), \mathrm{UNSAT}(x)\} - m/2$ where $\mathrm{SAT}(x)$ (resp. $\mathrm{UNSAT}(x)$) denotes the number of satisfied (resp. unsatisfied) d-XOR constraints. Thus injective norm can give an upper bound on the satisfiability of a

d-XOR instance and indeed many refutation algorithms (predominantly for random constraint satisfaction problems) exploit this connection.

Optimization over the sphere. $C_1 = \ldots C_d = \mathrm{Ball}(\ell_2^n)$.
For $d \geq 3$, this is a generalization of the spectral norm of a matrix. A certain promise variant of the problem is closely related to the quantum separability problem and consequently its approximability has connections to long standing open problems in quantum information theory. Optimization over the sphere also has connections to hypercontractivity and small-set expansion via $2 \to q$ norms, as well as to tensor principal component analysis and tensor decomposition [BKS15, GM15, MR14, HSS15]. The best approximation factor known for this case is polynomial in n. It is also very interesting from the perspective of inapproximability as it appears related to fundamental barriers in the theory of hardness of approximation and perhaps to inapproximability results of constraint satisfaction problems with very high density (a density at which random instances fail to be hard).

Optimization over ℓ_1. $C_1 = \ldots C_d = \mathrm{Ball}(\ell_1^n)$.
ℓ_1-optimization is closely related to optimization over the simplex and admits a PTAS for fixed d. Approximation algorithms for simplex optimization have been studied extensively in the optimization community [DK08] and have applications to portfolio optimization, game theory and population dynamics.

So in addition to being a mathematically intriguing pursuit, a characterization of the approximability of injective tensor norm promises to be of broad scope and applicability. It is then natural to ask the following questions:

Question.

1. *How does the approximability depend on the geometry of $C_1, \ldots, C_d$?*
2. *Can we determine the form of the best approximation factor (achieved by algorithms with polynomial runtime) as a function of $C_1, \ldots, C_d$?*
3. *What do the optimal approximation algorithms look like?*
4. *What does the approximation/runtime tradeoff look like?*

It is humbling how far this goal is from being achieved and there are yet many hurdles to cross before we can hope for a complete answer to these questions. There is a good deal of evidence in the combinatorial optimization community that convex programming relaxations and the sum of squares hierarchy are closely related to answering questions 3 and 4 respectively. We give a brief summary of convex programming and the sum of squares hierarchy in the next two sections.

1.1 Convex Programming Relaxations

An important paradigm from optimization theory is that a convex function f can be efficiently minimized over a compact convex set K given access to an oracle for f and a membership oracle oracle for K (under some additional technical conditions – see [GLS12] for a precise version of this statement).

A popular approach to approximating the optimum of NP-Hard problems is to relax the domain being optimized over to a convex domain; thus making the optimization problem tractable. As an example, given a $0/1$ integer program, one possible relaxation is to allow values to now be real numbers in the interval $[0, 1]$ or perhaps even vectors (as is the case in semidefinite programming (SDP) relaxations). Surprisingly, for an overwhelming majority of combinatorial optimization problems this method has produced the best known polynomial time approximation algorithms. In fact a beautiful result of Raghavendra [Rag08] establishes that assuming the unique games conjecture, a certain semidefinite programming relaxation is the optimal polynomial time approximation algorithm for a very wide class of problems known as constraint satisfaction problems. Similar theorems are known for Grothendieck's inequality [RS09] and strict constraint satisfaction problems [KMTV11] (which include many covering-packing problems like vertex cover).

This phenomenon suggests that similar statements might hold for continuous optimization problems and more general convex programming relaxations than SDPs. However, the only such result (i.e. matching approximation and hardness factors) we are aware of is [GRSW16] who established this for ℓ_p-subspace approximation and the problem of computing $\sup_{\|x\|_p \leq 1} x^T A\, x$ for $p \geq 2$. It would be very interesting to establish such convex programming optimality statements for injective tensor norms.

1.2 Sum of Squares Hierarchy

Relaxation hierarchies are procedures to obtain a hierarchy of convex relaxations. The convex relaxation obtained at each new level is stronger than that of the previous level at the cost of being larger in size. In a typical hierarchy, the q-th level relaxation has size $n^{O(q)}$. The first such hierarchy was given by Sherali and Adams [SA90] followed by Lovasz and Schrijver [LS91], both based on linear programming. The sum of squares (SoS) hierarchy is the strongest known convex programming hierarchy and there is considerable evidence in support of it achieving the right runtime vs. approximation trade-off for constraint satisfaction problems. Since CSPs are closely related to polynomial optimization over the hypercube it is reasonable to wonder if SoS might be the right hierarchy of approximation algorithms for polynomial optimization over other convex sets – for instance the sphere. Indeed, the SoS hierarchy of relaxations is defined precisely for the set of polynomial optimization problems and so is one natural candidate for studying polynomial optimization over convex sets that can be represented by polynomial constraints. The SoS hierarchy also captures (upto a logarithm in the exponent of the runtime) all known algorithmic results related to the HSEP problem from quantum information theory which can be viewed as a problem of obtaining additive approximations for degree-4 polynomial optimization over the sphere. The SoS hierarchy has also inspired new results in tensor decomposition/PCA for random tensors which are closely related to polynomial optimization over the sphere.

For these reasons and more, we will study the SoS hierarchy of relaxations for injective tensor norm (in the cases where it is well defined) as a means to predict the runtime approximation tradeoff and as evidence of intractability in cases where hardness of approximation results are difficult to obtain.

1.3 Brief Summary of Contributions

We next give a brief summary of the contributions of this thesis. A reader interested in perusing the document may skip this section and proceed to Chapter 2 and Chapter 10 where all results, related work and the relevant background are covered in detail.

- In Chapter 4 we obtain the first NP-hardness of approximation results for hypercontractive norms. Specifically, we prove inapproximability results for computing the $p \to q$ operator norm when $p \leq q$ and $2 \notin [p, q]$.

- In Chapter 6 Towards the goal of obtaining strong inapproximability results for $2 \to q$ norm when $q > 2$, we give random label cover (for which polynomial level SoS gaps are available) based hardness results for mixed norms, i.e., $2 \to \ell_q(\ell_{q'})$ for some $2 < q, q' < \infty$.

- In Chapter 5 we obtain improved approximation algorithms for computing the $p \to q$ operator norm when $p \geq 2 \geq q$.

- In Chapter 8 we introduce the technique of weak decoupling inequalities and use it to analyze the integrality gap of the SoS hierarchy for the maxima of various classes of polynomials over the sphere, namely arbitrary polynomials (improves on a result of Doherty and Wehner [DW12] for $q \ll n$), polynomials with non-negative coefficients and sparse polynomials.

- In Chapter 8 we also prove in the context of optimization over the sphere that "robust" integrality gaps for lower levels of a certain hierarchy of convex programs can be lifted to give higher level integrality gaps. This hierarchy is closely related to the SoS hierarchy but is possibly weaker. We hope that this method can find applications in other settings and perhaps even be shown to work in the context of the SoS hierarchy.

- In Chapter 9 we show an upper bound on the integrality gap of q levels of SoS on polynomials with random coefficients[1]. An interesting consequence of our result is that random/spiked-random instances cannot provide super-polylog level SoS gaps for the quantum Best Separable State problem.

[1] [RRS16] concurrently obtained slightly weaker bounds. However their bounds apply for the more general model of random polynomials with a sparsity parameter.

1.4 Chapter Credits

Chapter 4, Chapter 5, and Chapter 8 are based on joint works [BGG+18a, BGG+18b, BGG+17] respectively, with Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee and Madhur Tulsiani. Chapter 6 is based on unpublished joint work with the same authors.

Chapter 9 is based on the joint work [BGL16] with Venkatesan Guruswami and Euiwoong Lee.

1.5 Organization

We first provide a detailed exposition of our results in Chapter 2.

The rest of the document is then divided into two parts. The first part consists of results for the degree-2 case. It begins with the necessary normed space preliminaries (Chapter 3) which is followed by Chapter 4, Chapter 5, and Chapter 6 containing our results for operator norms.

The second part consists of results for degree-3 and beyond over the sphere. It begins with an SoS preliminaries chapter (Chapter 7) which is followed by our results for optimization over the sphere in Chapters 8 and 9.

Finally, in Chapter 10 we discuss the approximability landscape in full generality and also conclude with future directions and open problems.

Chapter 2

Detailed Results

In this chapter we specialize our discussion to the case of ℓ_p norms – an instructive case which itself involves multiple non-trivial results and also forms the mould for our conjectures in the more general case. We discuss the case of general norms in Chapter 10.

2.1 $\ell_p \to \ell_q$ Operator Norms

Consider the problem of finding the $\ell_p{\to}\ell_q$ norm of a given matrix $A \in \mathbb{R}^{m\times n}$ which is defined as

$$\|A\|_{p\to q} \;:=\; \max_{x\in\mathbb{R}^n\setminus\{0\}} \frac{\|Ax\|_q}{\|x\|_p}\,.$$

The quantity $\|A\|_{p\to q}$ is a natural generalization of the well-studied spectral norm, which corresponds to the case $p = q = 2$. For general p and q, this quantity computes the maximum distortion (stretch) of the operator A from the normed space ℓ_p^n to ℓ_q^m.

The case when $p = \infty$ and $q = 1$ relates to the well known Grothendieck inequality [KN12, Pis12], where the goal is to maximize $\langle y, Ax\rangle$ subject to $\|x\|_\infty, \|y\|_\infty \le 1$. In fact, via simple duality arguments, the general problem computing $\|A\|_{p\to q}$ can be seen to be equivalent to the following bilinear maximization problem (and to $\|A^T\|_{q^*\to p^*}$)

$$\|A\|_{p\to q} \;=\; \max_{\substack{\|x\|_p\le 1\\ \|y\|_{q^*}\le 1}} \langle y, Ax\rangle \;=\; \|A^T\|_{q^*\to p^*}\,,$$

where p^*, q^* denote the dual norms of p and q, satisfying $1/p + 1/p^* = 1/q + 1/q^* = 1$.

In Chapter 4 and Chapter 5 we study in detail, the algorithmic and complexity aspects of $\ell_p{\to}\ell_q$ norm. While this may seem either esoteric or narrow in scope, it turns out that resolving the $\ell_p{\to}\ell_q$ norm case is likely to have broad implications for the goal of characterizing the norms that admit constant factor approximations. A celebrated result of Maurey and Pisier states that every infinite dimensional Banach space X contains $(1 + \varepsilon)$-isomorphs of $\ell_{p_X}^k$ and $\ell_{q_X}^k$ where p_X (resp. q_X) is the modulus of Type (resp. Cotype) of X. Indeed combining finitary quantitative analogues of this result with the hardness of certain $\ell_p \to \ell_q$ norm norms derived in Chapter 4 yields inappproximability results for a broad class of norm sequences over $\mathbb{R}^n$. In addition to this, the ℓ_p case is connected to well studied problems in other areas. We next describe these connections, prior work, and our results in this context.

2.1.1 Hypercontractive norms, Small-Set Expansion, and Hardness.

$p \to q$ operator norms when $p < q$ are collectively referred to as *hypercontractive* norms, and have a special significance to the analysis of random walks, expansion and related problems in hardness of approximation [Bis11, BBH+12]. The problem of computing $\|A\|_{2\to 4}$ is also known to be equivalent to determining the maximum acceptance probability of a quantum protocol with multiple unentangled provers, and is related to several problems in quantum information theory [HM13, BH15].

Bounds on hypercontractive norms of operators are also used to prove expansion of small sets in graphs. Indeed, if f is the indicator function of set S of measure δ in a graph with adjacency matrix A, then we have that for any $p \le q$,

$$\Phi(S) \;=\; 1 - \frac{\langle f, Af \rangle}{\|f\|_2^2} \;\ge\; 1 - \frac{\|f\|_{q^*} \cdot \|Af\|_q}{\delta} \;\ge\; 1 - \|A\|_{p\to q} \cdot \delta^{1/p - 1/q}\,.$$

It was proved by Barak et al. [BBH+12] that the above connection to small-set expansion can in fact be made two-sided for a special case of the $2\to q$ norm. They proved that to resolve the promise version of the small-set expansion (SSE) problem, it suffices to distinguish the cases $\|A\|_{2\to q} \le c \cdot \sigma_{\min}$ and $\|A\|_{2\to q} \ge C \cdot \sigma_{\min}$, where $\sigma_{\min}$ is the least non-zero singular value of A and $C > c > 1$ are appropriately chosen constants based on the parameters of the SSE problem. Thus, the approximability of $2\to q$ norm is closely related to the small-set expansion problem. In particular, proving NP-hardness of approximating $2\to q$ norm is (necessarily) an intermediate goal towards proving the Small-Set Expansion Hypothesis of Raghavendra and Steurer [RS10].

However, relatively few algorithmic and hardness results are known for approximating hypercontractive norms. A result of Steinberg's [Ste05] gives an upper bound of $O(\max\{m, n\}^{25/128})$ on the approximation factor, for all p, q. For the case of $2\to q$ norm (for any $q > 2$), Barak et al. [BBH+12] give an approximation algorithm for the promise version of the problem described above, running in time $\exp\left(\tilde{O}(n^{2/q})\right)$. They also provide an additive approximation algorithm for $2\to 4$ norm (where the error depends on the $2\to 2$ norm and $2\to\infty$ norm of A), which was extended to the $2\to q$ norm by Harrow and Montanaro [HM13]. Barak et al. also prove NP-hardness of approximating $\|A\|_{2\to 4}$ within a factor of $1 + \tilde{O}(1/n^{o(1)})$, and hardness of approximating better than $\exp O((\log n)^{1/2-\varepsilon})$ in quasi-polynomial time, assuming the Exponential Time Hypothesis (ETH). This reduction was also used by Harrow, Natarajan and Wu [HNW16] to prove that $\tilde{O}(\log n)$ levels of the Sum-of-Squares SDP hierarchy cannot approximate $\|A\|_{2\to 4}$ within any constant factor.

It is natural to ask if the bottleneck in proving (constant factor) hardness of approximation for $2{\to}q$ norm arises from the fact from the nature of the domain (the ℓ_2 ball) or from hypercontractive nature of the objective. As discussed in Section 9.5, *all* hypercontractive norms present a barrier for gadget reductions, since if a "true" solution x is meant to encode the assignment to a (say) label cover problem with consistency checked via local gadgets, then (for $q > p$), a "cheating solution" may make the value of $\|Ax\|_q$ very large by using a sparse x which does not carry any meaningful information about the underlying label cover problem.

We show that (somewhat surprisingly) it is indeed possible to overcome the barrier for gadget reductions for hypercontractive norms, when $2 < p < q$ (and by duality, for any $p < q < 2$). This gives the first NP-hardness of approximation result for hypercontractive norms (under randomized reductions). Assuming ETH, this also rules out a constant factor approximation algorithm that runs in 2^{n^δ} for some $\delta := \delta(p,q)$.

Theorem. *For any p,q such that $1 < p \leq q < 2$ or $2 < p \leq q < \infty$ and $\varepsilon > 0$, there is no polynomial time algorithm that approximates the $p{\to}q$ norm of an $n \times n$ matrix within a factor $2^{\log^{1-\varepsilon} n}$ unless $\mathsf{NP} \subseteq \mathsf{BPTIME}\left(2^{(\log n)^{O(1)}}\right)$. When q is an even integer, the same inapproximability result holds unless $\mathsf{NP} \subseteq \mathsf{DTIME}\left(2^{(\log n)^{O(1)}}\right)$*

We also note that the operator A arising in our reduction satisfies $\sigma_{\min}(A) \approx 1$ (and is in fact a product of a carefully chosen projection and a scaled random Gaussian matrix). For such an A, we prove the hardness of distinguishing $\|A\|_{p\to q} \leq c$ and $\|A\|_{p\to q} \geq C$, for constants $C > c > 1$. For the corresponding problem in the case of $2{\to}q$ norm, Barak et al. [BBH+12] gave a subexponential algorithm running in time $\exp\left(O(n^{2/q})\right)$ (which works for every $C > c > 1$). On the other hand, since the running time of our reduction is $n^{O(q)}$, we get that assuming ETH, no algorithm can distinguish the above cases for $p{\to}q$ norm in time $\exp\left(n^{o(1/q)}\right)$, for any $p \leq q$ when $2 \notin [p,q]$.

While the above results give some possible reductions for working with hypercontractive norms, it remains an interesting problem to understand the role of the domain as a barrier to proving hardness results for the $2{\to}q$ norm problems. In fact, no hardness results are available even for the more general problem of polynomial optimization over the ℓ_2 ball. We view the above theorem as providing some evidence that while hypercontractive norms have been studied as a single class so far, the case when $2 \in [p,q]$ may be qualitatively different (with respect to techniques) from the case when $2 \notin [p,q]$. This is indeed known to be true in the *non-hypercontractive case* with $p \geq q$. In fact, our results are obtained via new hardness results for the case $p \geq q$, which we describe in a later subsection.

Towards $2 \to q$ Hardness. Strong inapproximability (SoS gaps or NP-hardness) results for the hypercontractive $2 \to q$ case remain elusive. Towards this goal, we consider the class of $2 \to X$ operator norms for exactly 2-convex norms X (see Section 3.6 for a definition). This class of contains all hypercontractive $2 \to q$ norms and moreover every operator norm in this class faces the same gadget reduction barrier discussed earlier. In

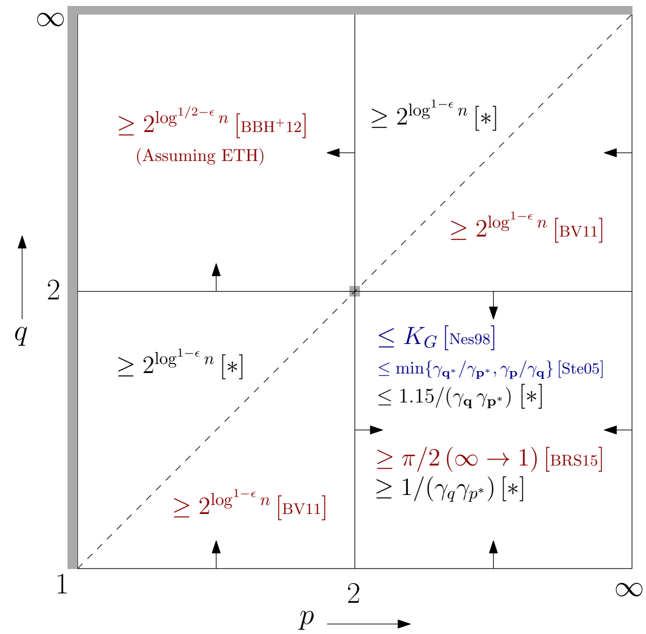

Figure 2.1: Upper and lower bounds for approximating $\|A\|_{p\to q}$. Arrows indicate the region to which a boundary belongs and thicker shaded regions represent exact algorithms. Our results are indicated by [∗]. We omit UGC-based hardness results in the figure.

Chapter 6 we show that this barrier can be overcome for certain exactly 2-convex norms (specifically mixed ℓ_p norms, i.e., $\ell_q(\ell_{q'})$ for $q, q' > 2$) and give a reduction from random label cover (for which polynomial level SoS gaps are available).

2.1.2 The non-hypercontractive case.

Several results are known in the case when $p \geq q$, and we summarize known results for matrix norms in Figure 2.1, for both the hypercontractive and non-hypercontractive cases. While the case of $p = q = 2$ corresponds to the spectral norm, the problem is also easy when $q = \infty$ (or equivalently $p = 1$) since this corresponds to selecting the row of A with the maximum ℓ_{p^*} norm. Note that in general, Figure 2.1 is symmetric about the principal diagonal. Also note that if $\|A\|_{p\to q}$ is a hypercontractive norm ($p < q$) then so is the equivalent $\|A^T\|_{q^*\to p^*}$ (the hypercontractive and non-hypercontractive case are separated by the non-principal diagonal).

As is apparent from the figure, the problem of approximating $\|A\|_{p\to q}$ for $p \geq q$ admits good approximations when $2 \in [q, p]$, and is hard otherwise. For the case when $2 \notin [q, p]$, an upper bound of $O(\max\{m, n\}^{25/128})$ on the approximation ratio was proved by Steinberg [Ste05]. Bhaskara and Vijayaraghavan [BV11] showed NP-hardness of approximation within any constant factor, and hardness of approximation within an $O\left(2^{(\log n)^{1-\varepsilon}}\right)$ factor for arbitrary $\varepsilon > 0$ assuming $\mathsf{NP} \not\subseteq \mathsf{DTIME}\left(2^{(\log n)^{O(1)}}\right)$.

Determining the right constants in these approximations when $2 \in [q, p]$ has been of considerable interest in the analysis and optimization community. For the case of $\infty \to 1$ norm, Grothendieck's theorem [Gro53] shows that the integrality gap of a semidefinite programming (SDP) relaxation is bounded by a constant, and the (unknown) optimal value is now called the Grothendieck constant K_G. Krivine [Kri77] proved an upper bound of $\pi/(2\ln(1+\sqrt{2})) = 1.782\ldots$ on K_G, and it was later shown by Braverman et al. that K_G is strictly smaller than this bound. The best known lower bound on K_G is about 1.676, due to (an unpublished manuscript of) Reeds [Ree91] (see also [KO09] for a proof).

An upper bound of K_G on the approximation factor also follows from the work of Nesterov [Nes98] for any $p \geq 2 \geq q$. A later work of Steinberg [Ste05] also gave an upper bound of $\min\{\gamma_p/\gamma_q, \gamma_{q^*}/\gamma_{p^*}\}$, where γ_p denotes p^{th} norm of a standard normal random variable (i.e., the p-th root of the p-th Gaussian moment). Note that Steinberg's bound is less than K_G for some values of (p,q), in particular for all values of the form $(2,q)$ with $q \leq 2$ (and equivalently $(p,2)$ for $p \geq 2$), where it equals $1/\gamma_q$ (and $1/\gamma_{p^*}$ for $(p,2)$).

On the hardness side, Briët, Regev and Saket [BRS15] showed NP-hardness of $\pi/2$ for the $\infty \to 1$ norm, strengthening a hardness result of Khot and Naor based on the Unique Games Conjecture (UGC) [KN09] (for a special case of the Grothendieck problem when the matrix A is positive semidefinite). Assuming UGC, a hardness result matching Reeds' lower bound was proved by Khot and O'Donnell [KO09], and hardness of approximating within K_G was proved by Raghavendra and Steurer [RS09].

For a related problem known as the L_p-Grothendieck problem, where the goal is to maximize $\langle x, Ax \rangle$ for $\|x\|_p \leq 1$, results by Steinberg [Ste05] and Kindler, Schechtman and Naor [KNS10] give an upper bound of γ_p^2, and a matching lower bound was proved assuming UGC by [KNS10], which was strengthened to NP-hardness by Guruswami et al. [GRSW16]. However, note that this problem is quadratic and not necessarily bilinear, and is in general much harder than the Grothendieck problems considered here. In particular, the case of $p = \infty$ only admits an $\Theta(\log n)$ approximation instead of K_G for the bilinear version [AMMN06, ABH$^+$05].

The Search For Optimal Constants and Optimal Rounding Algorithms. Determining the right approximation ratio for these problems often leads to the development of rounding algorithms that apply much more broadly. For the Grothendieck problem, the goal is to find $y \in \mathbb{R}^m$ and $x \in \mathbb{R}^n$ with $\|y\|_\infty, \|x\|_\infty \leq 1$, and one considers the following semidefinite relaxation:

$$\begin{aligned}
&\text{maximize} && \sum_{i,j} A_{i,j} \cdot \langle u^i, v^j \rangle \quad \text{s.t.} && \\
&\text{subject to} && \|u^i\|_2 \leq 1, \|v^j\|_2 \leq 1 && \forall i \in [m], j \in [n] \\
& && u^i, v^j \in \mathbb{R}^{m+n} && \forall i \in [m], j \in [n]
\end{aligned}$$

By the bilinear nature of the problem above, it is clear that the optimal x, y can be taken to have entries in $\{-1, 1\}$. A bound on the approximation ratio[1] of the above program is then obtained by designing a good "rounding" algorithm which maps the vectors u^i, v^j to values in $\{-1, 1\}$. Krivine's analysis [Kri77] corresponds to a rounding algorithm which considers a random vector $\mathbf{g} \sim \mathcal{N}(0, I_{m+n})$ and rounds to x, y defined as

$$y_i \;:=\; \operatorname{sgn}\left(\left\langle \varphi(u^i), \mathbf{g} \right\rangle\right) \quad \text{and} \quad x_j \;:=\; \operatorname{sgn}\left(\left\langle \psi(v^j), \mathbf{g} \right\rangle\right),$$

for some appropriately chosen transformations φ and ψ. This gives the following upper bound on the approximation ratio of the above relaxation, and hence on the value of the Grothendieck constant K_G:

$$K_G \;\leq\; \frac{1}{\sinh^{-1}(1)} \cdot \frac{\pi}{2} \;=\; \frac{1}{\ln(1+\sqrt{2})} \cdot \frac{\pi}{2}.$$

Braverman et al. [BMMN13] show that the above bound can be strictly improved (by a very small amount) using a two dimensional analogue of the above algorithm, where the value y_i is taken to be a function of the two dimensional projection $(\langle \varphi(u^i), \mathbf{g}_1 \rangle, \langle \varphi(u^i), \mathbf{g}_2 \rangle)$ for independent Gaussian vectors $\mathbf{g}_1, \mathbf{g}_2 \in \mathbb{R}^{m+n}$ (and similarly for x). Naor and Regev [NR14] show that such schemes are optimal in the sense that it is possible to achieve an approximation ratio arbitrarily close to the true (but unknown) value of K_G by using k-dimensional projections for a large (constant) k. A similar existential result was also proved by Raghavendra and Steurer [RS09] who proved that the there exists a (slightly different) rounding algorithm which can achieve the (unknown) approximation ratio K_G.

For the case of arbitrary $p \geq 2 \geq q$, Nesterov [Nes98] considered the convex program in Figure 2.2, denoted as $\mathsf{CP}(A)$, generalizing the one above. Note that since $q^* \geq 2$

$$\begin{aligned}
\text{maximize} \quad & \sum_{i,j} A_{i,j} \cdot \langle u^i, v^j \rangle \;=\; \left\langle A, UV^T \right\rangle \\
\text{subject to} \quad & \sum_{i \in [m]} \|u^i\|_2^{q^*} \;\leq\; 1 \\
& \sum_{j \in [n]} \|v^j\|_2^{p} \;\leq\; 1 \\
& u^i, v^j \in \mathbb{R}^{m+n}
\end{aligned}$$

u^i (resp. v^j) is the i-th (resp. j-th) row of U (resp. V)

Figure 2.2: The relaxation $\mathsf{CP}(A)$ for approximating $p \to q$ norm of a matrix $A \in \mathbb{R}^{m \times n}$.

and $p \geq 2$, the above program is convex in the entries of the Gram matrix of the vectors $\{u^i\}_{i \in [m]} \cup \{v^j\}_{j \in [n]}$. Although the stated bound in [Nes98] is slightly weaker (as it is proved for a larger class of problems), the approximation ratio of the above relaxation can

[1] Since we will be dealing with problems where the optimal solution may not be integral, we will use the term "approximation ratio" instead of "integrality gap".

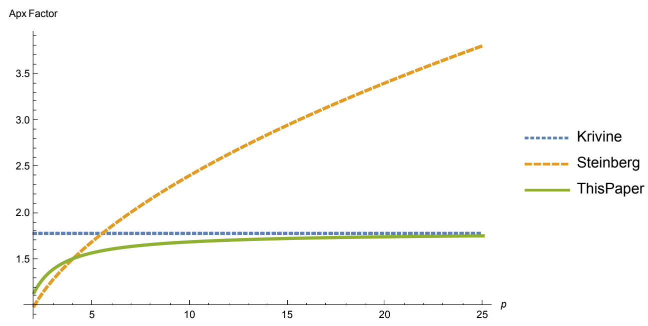

Figure 2.3: A comparison of the bounds for approximating $p \to p^*$ norm obtained from Krivine's rounding for K_G, Steinberg's analysis, and our bound. While our analysis yields an improved bound for $4 \le p \le 66$, we believe that the rounding algorithm achieves an improved bound for all $2 < p < \infty$.

be shown to be bounded by K_G. By using the Krivine rounding scheme of considering the sign of a random Gaussian projection (aka random hyperplane rounding) one can show that Krivine's upper bound on K_G still applies to the above problem.

Motivated by applications to robust optimization, Steinberg [Ste05] obtained an upper bound of $\min\{\gamma_p/\gamma_q, \gamma_{q^*}/\gamma_{p^*}\}$ on the approximation factor. Note that while Steinberg's bound is better (approaches 1) as p and q approach 2, it is unbounded when $p, q^* \to \infty$ (as in the Grothendieck problem).

Based on the inapproximability result of factor $1/(\gamma_{p^*} \cdot \gamma_q)$ obtained in this work, it is natural to ask if this is the "right form" of the approximation ratio. Indeed, this ratio is $\pi/2$ when $p^* = q = 1$, which is the ratio obtained by Krivine's rounding scheme, up to a factor of $\ln(1+\sqrt{2})$. We extend Krivine's result to all $p \ge 2 \ge q$ as below.

Theorem. *There exists a fixed constant $\varepsilon_0 \le 0.00863$ such that for all $p \ge 2 \ge q$, the approximation ratio of the convex relaxation $\mathsf{CP}(A)$ is upper bounded by*

$$\frac{1+\varepsilon_0}{\sinh^{-1}(1)} \cdot \frac{1}{\gamma_{p^*} \cdot \gamma_q} = \frac{1+\varepsilon_0}{\ln(1+\sqrt{2})} \cdot \frac{1}{\gamma_{p^*} \cdot \gamma_q}.$$

Perhaps more interestingly, the above theorem is proved via a generalization of hyperplane rounding, which we believe may be of independent interest. Indeed, for a given collection of vectors $w^1, \ldots, w^m$ considered as rows of a matrix W, Gaussian hyperplane rounding corresponds to taking the "rounded" solution y to be the

$$y := \underset{\|y'\|_\infty \le 1}{\operatorname{argmax}} \langle y', W\mathbf{g} \rangle = \left(\operatorname{sgn}\left(\left\langle w^i, \mathbf{g} \right\rangle\right)\right)_{i \in [m]}.$$

We consider the natural generalization to (say) ℓ_r norms, given by

$$y := \underset{\|y'\|_r \le 1}{\operatorname{argmax}} \langle y', W\mathbf{g} \rangle = \left(\frac{\operatorname{sgn}\left(\langle w^i, \mathbf{g}\rangle\right) \cdot \left|\langle w^i, \mathbf{g}\rangle\right|^{r^*-1}}{\|W\mathbf{g}\|_{r^*}^{r^*-1}}\right)_{i \in [m]}.$$

We refer to y as the "Hölder dual" of $W\mathbf{g}$, since the above rounding can be obtained by viewing $W\mathbf{g}$ as lying in the dual (ℓ_{r^*}) ball, and finding the y for which Hölder's inequality is tight. Indeed, in the above language, Nesterov's rounding corresponds to considering the ℓ_∞ ball (hyperplane rounding). While Steinberg used a somewhat different relaxation, the rounding there can be obtained by viewing $W\mathbf{g}$ as lying in the primal (ℓ_r) ball instead of the dual one. In case of hyperplane rounding, the analysis is motivated by the identity that for two unit vectors u and v, we have

$$\mathop{\mathbb{E}}_{\mathbf{g}}\left[\operatorname{sgn}(\langle \mathbf{g}, u\rangle) \cdot \operatorname{sgn}(\langle \mathbf{g}, v\rangle)\right] \;=\; \frac{2}{\pi} \cdot \sin^{-1}(\langle u, v\rangle)\,.$$

We prove the appropriate extension of this identity to ℓ_r balls (and analyze the functions arising there) which may also be of interest for other optimization problems over ℓ_r balls.

Hardness. We extend the hardness results of [BRS15] for the $\infty \to 1$ and $2 \to 1$ norms of a matrix to any $p \geq 2 \geq q$. The hardness factors obtained match the performance of known algorithms (due to Steinberg [Ste05]) for the cases of $2 \to q$ and $p \to 2$, and moreover almost match the algorithmic results in the more general case of $p \geq 2 \geq q$.

Theorem. *For any p, q such that $p \geq 2 \geq q$ and $\varepsilon > 0$, it is NP-hard to approximate the $p{\to}q$ norm within a factor $1/(\gamma_{p^*}\gamma_q) - \varepsilon$.*

2.2 Polynomial Optimization over the Sphere

In Chapter 8 and Chapter 9 we study the problem of optimizing homogeneous polynomials over the unit sphere. Formally, given an n-variate degree-d homogeneous polynomial f, the goal is to compute $\|f\|_2 := \sup_{\|x\|_2=1} |f(x)|$. For $d \geq 3$, it defines a natural higher-order analogue of the eigenvalue problem for matrices. The problem also provides an important testing ground for the development of new spectral and semidefinite programming (SDP) techniques, and techniques developed in the context of this problem have had applications to various other constrained settings [HLZ10, Lau09, Las09]. Besides being a natural and fundamental problem in its own right, it has connections to widely studied questions in many other areas, like the small set expansion hypothesis [BBH+12, BKS14], tensor low-rank decomposition and tensor PCA [BKS15, GM15, MR14, HSS15], refutation of random constraint satisfaction problems [RRS16] and planted clique [BV09].

Optimization over $\mathbb{S}^{n-1}$ has been given much attention in the optimization community, where for a fixed number of variables n and degree d of the polynomial, it is known that the estimates produced by q levels a certain hierarchy of SDPs (Sum of Squares) get arbitrarily close to the true optimal solution as q increases (see [Las09] for various applications). We refer the reader [dKL19, DW12, Fay04, dKLS14] for more information on convergence results. These algorithms run in time $n^{O(q)}$, which is polynomial for constant q. Unfortunately, known convergence results often give sub-optimal bounds when q is sub-linear in n.

In computer science, much attention has been given to the sub-exponential runtime regime (i.e. $q \ll n$) since many of the target applications such as SSE, QMA and refuting random CSPs are of considerable interest in this regime. In addition to the polytime $n^{d/2-1}$-approximation for general polynomials [HLZ10, So11], approximation guarantees have been proved for several special cases including $2 \to q$ norms [BBH+12], polynomials with non-negative coefficients [BKS14], and some polynomials that arise in quantum information theory [BKS17, BH13]. An outstanding open problem in quantum information theory is settling the complexity of the best separable state problem (which can be viewed as seeking a $(1+\varepsilon)$-approximation for maximizing a certain class of polynomials over the sphere), and the sum of squares (SoS) hierarchy captures all known algorithms for this problem upto a logarithm in the exponent [BKS14, BKS17]. The best known upper bound being $\sqrt{n}$ levels due to Barak, Kothari and Steurer [BKS17] and the best lower bound being $\log n$ levels due to Harrow, Natarajan and Wu [HNW16]. Hence there is considerable interest in understanding the approximation/runtime tradeoff (especially in the sub-exponential regime).

In this thesis, we develop general techniques to design and analyze algorithms for polynomial optimization over the sphere. The sphere constraint is one of the simplest constraints for polynomial optimization and thus is a good testbed for techniques. Indeed, we believe these techniques will also be useful in understanding polynomial optimization for other constrained settings.

In addition to giving an analysis the problem for arbitrary polynomials, these techniques can also be adapted to take advantage of the structure of the input polynomial, yielding better approximations for several special cases such as polynomials with non-negative coefficients, and sparse polynomials. Previous polynomial time algorithms for polynomial optimization work by reducing the problem to diameter estimation in convex bodies [So11] and seem unable to utilize structural information about the (class of) input polynomials. Development of a method which can use such information was stated as an open problem by Khot and Naor [KN08] (in the context of ℓ_∞ optimization).

Our approximation guarantees are with respect to the optimum at each level of the SoS hierarchy. Such SDPs are the most natural tools to bound the optima of polynomial optimization problems, and our results shed light on the efficacy of higher levels of the SoS hierarchy to deliver better approximations to the optimum.

- In Chapter 8 we introduce a technique we refer to as "weak decoupling inequalities" and use it to upper-bound the integrality gap of q levels of the SoS hierarchy for various classes of polynomials over the sphere, namely arbitrary polynomials (when $q \ll n$, our result yields better bounds than those implied by convergence results [DW12, dKL19][2]), polynomials with non-negative coefficients and sparse polynomials:

  Aribitrary. SoS_q gets an $(O(n)/q)^{d/2-1}$ approximation.

  Non-negative Coefficients. SoS_q gets an $(O(n)/q)^{d/4-1/2}$ approximation.

  Sparse. SoS_q gets a $\sqrt{m/q}$ approximation where m is the sparsity.

[2] [dKL19] appeared after the publication of this work

We believe that these techniques are applicable to much broader settings.

- In Chapter 8 we also prove for a certain hierarchy of relaxations for optimization over the sphere that "robust" integrality gaps for lower levels of the hierarchy can be lifted to integrality gaps for higher levels. This hierarchy is closely related to the SoS hierarchy but is possibly weaker (in fact a majority of the works studying SoS and optimization over the sphere can be seen as using only this weaker hierarchy). We give an example application of this result by using it to show polynomial level integrality gaps (for the aforementioned weaker hierarchy of relaxations) for optimizing non-negative coefficient polynomials.[3] We hope that this method can find applications in other settings and perhaps even be shown to work in the context of the SoS hierarchy.

- In Chapter 9, we show an upper bound on the integrality gap of q levels of SoS on polynomials with random coefficients[4]. Specifically we show that SoS certifies an upper bound that is an $(O(n)/q)^{d/4-1/2}$ approximation to the true value. An interesting consequence of our result is that random/spiked-random instances cannot provide Best Separable State gap instances for more than polylog levels of SoS.

[3] Integrality gaps for polynomial levels of SoS are already known in the case of arbitrary polynomials due to a result of Hopkins et al. [HKP+17]. More precisely, they show polynomial level integrality gaps for polynomials with i.i.d. ± 1 random coefficients.

[4] [RRS16] concurrently obtained slightly weaker bounds. However their bounds apply for the more general model of sparse random polynomials thereby finding applications to refutation of random CSPs.

Part II

Degree-2 (Operator Norms)

Chapter 3

Preliminaries (Normed Spaces)

3.1 Vectors

Unless otherwise specified, vectors are assumed to be finite dimensional with real valued coordinates. For a vector $x \in \mathbb{R}^n$, we denote its i-th coordinate by x_i (Chapter 5 is the only exception to this wherein it is more convenient to think of vectors as functions and so the notation $x(i)$ is used).

We denote sequences of vectors with superscripts, e.g. $x^1, x^2, \cdots \in \mathbb{R}^n$.

For $x, y \in \mathbb{R}^n$ we let $\langle x\, , y\rangle$ denote the inner product under the counting measure, i.e.,

$$\langle x\, , y\rangle := \sum_{i\in[n]} x_i \cdot y_i$$

For a scalar function $f : \mathbb{R} \to \mathbb{R}$ and $x \in \mathbb{R}^n$ we denote by $f[x] \in \mathbb{R}^n$ the vector obtained by applying f to x entry-wise.

For a vector u, we use D_u to denote the diagonal matrix with the entries of u forming the diagonal, and for a matrix M we use $\operatorname{diag}(M)$ to denote the vector of diagonal entries.

3.2 Norms

A function $\|\cdot\|_X$ from $\mathbb{R}^n$ to the non-negative reals is called a norm if it satisfies

Subadditivity. $\|x + y\|_X \leq \|x\|_X + \|y\|_X$.

Absolute Homogeneity. $\|c \cdot x\|_X = |c| \cdot \|x\|_X$ for any $c \in \mathbb{R}$.

Positive Definiteness. $\|x\|_X = 0$ implies $x = 0$.

We say a convex body K is symmetric if $K = -K$ (i.e., $x \in K \Rightarrow -x \in K$). There is a well known correspondence between norms and symmetric convex bodies. The map from norms to symmetric convex bodies is referred to as the unit ball of the norm and is defined as

$$\operatorname{Ball}(X) := \{x \mid \|x\|_X \leq 1\}.$$

The inverse map is known as the Minkowski functional and is defined as

$$\|x\| := \inf_{t>0}\{x/t \in K\}.$$

The dual space of $\mathbb{R}^n$ is $\mathbb{R}^n$ itself. For a norm X and $\xi \in \mathbb{R}^n$, the dual norm is defined as

$$\|\xi\|_{X^*} := \sup_{x\in \mathrm{Ball}(X)} \langle \xi , x\rangle$$

We will often consider families of normed spaces of increasing dimension, and we will denote this by $(X_n)_{n\in\mathbb{N}}$ where X_n is assumed to be a norm over $\mathbb{R}^n$.

3.3 ℓ_p Norms

For $p \geq 1$ and a vector x, we denote the counting ℓ_p-norm as $\|x\|_p := (\sum_i |x_i|^p)^{1/p}$ (when $p = \infty$ it is defined as $\max_{i\in[n]} |x_i|$).
For any $p \in [1, \infty]$, the dual norm of ℓ_p is ℓ_{p^*}, where p^* is defined as satisfying the equality: $\frac{1}{p} + \frac{1}{p^*} = 1$. Formally,

Fact 3.3.1. *For any $p \in [1, \infty]$, $\|\xi\|_{p^*} = \sup_{x\in \mathrm{Ball}(\ell_p)} \langle \xi , x\rangle$.*

For $p \geq 1$, we define the p-th Gaussian norm of a standard gaussian g as

$$\gamma_p := \Big(\mathop{\mathbb{E}}_{g\sim\mathcal{N}(0,1)}\big[|g|^p\big]\Big)^{1/p}.$$

3.4 Operator Norm

For a linear operator A mapping a normed space X over $\mathbb{R}^n$ to normed space Y over $\mathbb{R}^m$, the operator norm is defined as the maximum amount that A stretches an X-unit vector where stretch is measured according to Y i.e.,

$$\|A\|_{X\to Y} := \max_{x\in\mathbb{R}^n\setminus\{0\}} \|Ax\|_Y/\|x\|_X$$

We say two normed spaces X, Y are *isomorphic* (resp. *isometric*) to one another if there exists an invertible linear operator $A : X \to Y$ such that A and A^{-1} have bounded operator norm (resp. operator norm 1).[1] In the case where $X = \ell_p$, $Y = \ell_q$, we'll use the shorthand $\|A\|_{p\to q}$ to denote the operator norm.
We next record the equivalence of operator norms with bilinear form maximization.

[1] Since any two norms over a finite dimensional space are isomorphic, this notion is only interesting in infinite dimensional settings. The notion of isometry however remains interesting in the finite dimensional setting.

Fact 3.4.1. *For any linear operator $A : X \to Y$,*

$$\|A\|_{X\to Y} = \sup_{\|y\|_{Y^*}\le 1} \sup_{\|x\|_X\le 1} \langle y\,, Ax\rangle = \|A^T\|_{Y^*\to X^*}.$$

Proof. Using the fact that $\langle y\,, Ax\rangle = \langle x\,, A^T y\rangle$,

$$\begin{aligned}\|A\|_{X\to Y} &= \sup_{\|x\|_X\le 1} \|Ax\|_Y = \sup_{\|x\|_X\le 1} \sup_{\|y\|_{Y^*}\le 1} \langle y\,, Ax\rangle = \sup_{\|x\|_X\le 1} \sup_{\|y\|_{Y^*}\le 1} \langle x\,, A^T y\rangle \\ &= \sup_{\|y\|_{Y^*}\le 1} \|A^T y\|_{X^*} = \|A^T\|_{Y^*\to X^*}\,.\end{aligned}$$

■

Operator norms are submultiplicative in the following sense:

Fact 3.4.2. *For any norms X, Y, Z, and linear operators $C : X \to Y,\ B : Y \to Z$,*

$$\|BC\|_{X\to Z} \;=\; \sup_x \frac{\|BCx\|_Z}{\|x\|_X} \;\le\; \sup_x \frac{\|B\|_{Y\to Z}\|Cx\|_Y}{\|x\|_X} \;=\; \|C\|_{X\to Y}\|B\|_{Y\to Z}.$$

3.5 Type and Cotype

The notions of Type and Cotype are powerful classification tools from Banach space theory.

Definition 3.5.1. *The Type-2 constant of a Banach space X, denoted by $T_2(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x^i\}$ in X,*

$$\mathbb{E}\left[\|\sum_i \varepsilon_i \cdot x^i\|\right] \le C \cdot \sqrt{\sum_i \|x^i\|^2}$$

where ε_i is an independent Rademacher random variable. We say X is of Type-2 if $T_2(X) < \infty$.[2]

Definition 3.5.2. *The Cotype-2 constant of a Banach space X, denoted by $C_2(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x^i\}$ in X,*

$$\mathbb{E}\left[\|\sum_i \varepsilon_i \cdot x^i\|\right] \ge \frac{1}{C} \cdot \sqrt{\sum_i \|x^i\|^2}$$

where ε_i is an independent Rademacher random variable. We say X is of Cotype-2 if $C_2(X) < \infty$.

Remark 3.5.3.

[2] $T_2(X) < \infty$ is yet another property that is uninteresting in the finite dimensional setting since every finite dimensional norm is of Type-2 by John's theorem. However for a sequence of norms $(X_n)_{n\in\mathbb{N}}$, the dependence of $T_2(X_n)$ on the dimension n is an interesting and useful property to track and statements derived from Type-2 properties of infinite dimensional spaces can often be adapted to give quantitative finite dimensional versions when considering a sequence of norms of growing dimension.

- *It is known that* $C_2(X^*) \leq T_2(X)$.

- *It is known that for* $p \geq 2$, *we have* $T_2(\ell_p^n) = \gamma_p$ *(while* $C_2(\ell_p^n) \to \infty$ *as* $n \to \infty$*) and for* $q \leq 2$, $C_2(\ell_q^n) = \max\{2^{1/q-1/2}, 1/\gamma_q\}$ *(while* $T_2(\ell_q^n) \to \infty$ *as* $n \to \infty$*).*

We say X is Type-2 (resp. Cotype-2) if $T_2(X) < \infty$ (resp. $C_2(X) < \infty$). $T_2(X)$ and $C_2(X)$ can be regarded as measures of the "closeness" of X to a Hilbert space. Some notable manifestations of this correspondence are:

- $T_2(X) = C_2(X) = 1$ if and only if X is isometric to a Hilbert space.

- Kwapień [Kwa72a]: X is of Type-2 and Cotype-2 if and only if it is isomorphic to a Hilbert space.

- Figiel, Lindenstrauss and Milman [FLM77]: If X is a Banach space of Cotype-2, then any n-dimensional subspace of X has an $m = \Omega(n)$-dimensional subspace with Banach-Mazur distance at most 2 from ℓ_2^m.

More generally one defines

Definition 3.5.4 (Type/Cotype)**.** *For* $1 \leq p \leq 2$, *the Type-p constant of a normed space* X, *denoted by* $T_p(X)$, *is the smallest constant* C *such that for every finite sequence of vectors* $\{x^i\}$ *in* X,

$$\mathbb{E}\left[\|\sum_i \varepsilon_i \cdot x^i\|\right] \leq C \cdot \left(\sum_i \|x^i\|^p\right)^{1/p}$$

where ε_i *is an independent Rademacher random variable.* X *is said to have Type-p if* $T_p(X) < \infty$.

For $2 \leq q \leq \infty$, *the Cotype-q constant of a normed space* X, *denoted by* $C_q(X)$, *is the smallest constant* C *such that for every finite sequence of vectors* $\{x^i\}$ *in* X,

$$\mathbb{E}\left[\|\sum_i \varepsilon_i \cdot x^i\|\right] \geq \frac{1}{C} \cdot \left(\sum_i \|x^i\|^q\right)^{1/q}.$$

X *is said to have Cotype-q if* $C_q(X) < \infty$.

Any normed space X trivially is Type-1 and Cotype-∞. It is easily checked that Type-p implies Type-p' for any $p' \leq p$ and Cotype-q implies Cotype-q' for any $q \geq q'$. Let $p_X := \sup\{p \mid T_p(X) < \infty\}$ and $q_X := \inf\{q \mid C_q(X) < \infty\}$. p_X (resp. q_X) is referred to as the modulus of Type (resp. Cotype).
Another example of the power of these notions in classifying Banach spaces is the celebrated MP+K theorem:

Theorem 3.5.5 (Maurey and Pisier + Krivine)**.** *Any infinite dimensional Banach space* X *contains for any* $\varepsilon > 0$, $(1+\varepsilon)$*-isomorphs of* ℓ_{p_X} *and* ℓ_{q_X} *of arbitrarily large dimension.*

3.6 p-convexity and q-concavity

The notions of p-convexity and q-concavity are well defined for a wide class of normed spaces known as Banach lattices. In this document we only define these notions for finite dimensional norms that are 1-unconditional in the elementary basis (i.e., those norms X for which flipping the sign of an entry of x does not change the norm. We shall refer to such norms as *sign-invariant norms*). Most of the statements we make in this context can be readily extended to the case of norms admitting some 1-unconditional basis, but we choose to fix the elementary basis in the interest of clarity. With respect to the goals of this document, we believe most of the key insights are already manifest in the elementary basis case.

Definition 3.6.1 (p-convexity/q-concavity)**.** *Let X be a sign-invariant norm over $\mathbb{R}^n$. Then for $1 \leq p \leq \infty$ the p-convexity constant of X, denoted by $M^{(p)}(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x^i\}$ in X,*

$$\left\| \left[\sum_i |[x^i]|^p \right]^{1/p} \right\| \leq C \cdot \left(\sum_i \|x^i\|^p \right)^{1/p}$$

X is said to be p-convex if $M^{(p)}(X) < \infty$. We will say X is exactly p-convex if $M^{(p)}(X) = 1$.

For $1 \leq q \leq \infty$, the q-concavity constant of X, denoted by $M_{(q)}(X)$, is the smallest constant C such that for every finite sequence of vectors $\{x^i\}$ in X,

$$\left\| \left[\sum_i |[x^i]|^q \right]^{1/q} \right\| \geq \frac{1}{C} \cdot \left(\sum_i \|x^i\|^q \right)^{1/q}.$$

X is said to be q-concave if $M_{(q)}(X) < \infty$.
We will say X is exactly q-concave if $M_{(q)}(X) = 1$.

Every sign-invariant norm is exactly 1-convex and ∞-concave.

For a sign-invariant norm X over $\mathbb{R}^n$, and any $0 < p < \infty$ let $X^{(p)}$ denote the function $\||[x]|^p\|_X^{1/p}$. $X^{(p)}$ is referred to as the p-convexification of X. It is easily verified that $M^{(p)}(X^{(p)}) = M^{(1)}(X)$ and further that $X^{(p)}$ is an exactly p-convex sign-invariant norm if and only if X is a sign-invariant norm (and therefore exactly 1-convex).

3.7 Convex Relaxation for Operator Norm

In this section we will see that there is a natural convex relaxation for a wide class of operator norms. It is instructive to first consider the pertinent relaxation for Grothendieck's inequality. Recall the bilinear formulation of the problem wherein given an $m \times n$ matrix

A, the goal is to maximize $y^T A\, x$ over $\|y\|_\infty, \|x\|_\infty \le 1$. One then considers the following semidefinite programming relaxation:

$$\begin{array}{lll} \text{maximize} & \sum_{i,j} A_{i,j} \cdot \langle u^i, v^j \rangle \quad \text{s.t.} & \\ \text{subject to} & \|u^i\|_2 \le 1, \|v^j\|_2 \le 1 & \forall i \in [m], j \in [n] \\ & u^i, v^j \in \mathbb{R}^{m+n} & \forall i \in [m], j \in [n] \end{array}$$

which is equivalent to

$$\begin{array}{ll} \textbf{maximize} & \frac{1}{2} \cdot \left\langle \begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}, \begin{bmatrix} \mathbb{Y} & \mathbb{W} \\ \mathbb{W}^T & \mathbb{X} \end{bmatrix} \right\rangle \quad \text{s.t.} \\ & \mathbb{X}_{i,i} \le 1, \quad \mathbb{Y}_{j,j} \le 1 \\ & \begin{bmatrix} \mathbb{Y} & \mathbb{W} \\ \mathbb{W}^T & \mathbb{X} \end{bmatrix} \succeq 0, \quad \mathbb{Y} \in \mathbb{S}^{m\times m},\ \mathbb{X} \in \mathbb{S}^{n\times n},\ \mathbb{W} \in \mathbb{R}^{m\times n} \end{array}$$

where $\mathbb{S}^{m\times m}$ is the set of $m \times m$ symmetric positive semidefinite matrices in $\mathbb{R}^{m\times m}$.

Nesterov [Nes98, NWY00][3] and independently Naor and Schechtman[4] observed that if X and Y^* are exactly 2-convex, then there is a natural computable convex relaxation for the bilinear formulation of $X \to Y$ operator norm. Recall the goal is to maximize $y^T A\, x$ over $\|y\|_{Y^*}, \|x\|_X \le 1$. The relaxation which we will call $\mathsf{CP}(A)$ is as follows:

$$\begin{array}{ll} \textbf{maximize} & \frac{1}{2} \cdot \left\langle \begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}, \begin{bmatrix} \mathbb{Y} & \mathbb{W} \\ \mathbb{W}^T & \mathbb{X} \end{bmatrix} \right\rangle \quad \text{s.t.} \\ & \mathrm{diag}(\mathbb{X}) \in \mathrm{Ball}(X^{(1/2)}), \quad \mathrm{diag}(\mathbb{Y}) \in \mathrm{Ball}(Y^{*(1/2)}) \\ & \begin{bmatrix} \mathbb{Y} & \mathbb{W} \\ \mathbb{W}^T & \mathbb{X} \end{bmatrix} \succeq 0, \quad \mathbb{Y} \in \mathbb{S}^{m\times m},\ \mathbb{X} \in \mathbb{S}^{n\times n},\ \mathbb{W} \in \mathbb{R}^{m\times n} \end{array}$$

For a vector s, let D_s denote the diagonal matrix with s as diagonal entries. Let $\overline{X} := (X^{(1/2)})^*$, $\overline{Y} := (Y^{*(1/2)})^*$. We can then define the dual program $\mathrm{DP}(A)$ as follows:

$$\begin{array}{l} \textbf{minimize} \quad (\|s\|_{\overline{Y}} + \|t\|_{\overline{X}})/2 \quad \text{s.t.} \\ \begin{bmatrix} D_s & -A \\ -A^T & D_t \end{bmatrix} \succeq 0, \quad s \in \mathbb{R}^m,\ t \in \mathbb{R}^n . \end{array}$$

Strong duality is satisfied, i.e. $\mathrm{DP}(A) = \mathsf{CP}(A)$, and a proof can be found in [NWY00] (see Lemma 13.2.2 and Theorem 13.2.3).

---

[3] Nesterov uses the language of quadratic programming and appears not to have noticed the connections to Banach space theory. In fact, it appears that Nesterov even gave yet another proof of an $O(1)$ upper bound on Grothendieck's constant.

[4] personal communication

3.8 Factorization of Linear Operators

Let X, Y, E be Banach spaces and let $A : X \to Y$ be a continuous linear operator. We say that A *factorizes* through E if there exist continuous operators $C : X \to E$ and $B : E \to Y$ such that $A = BC$. Factorization theory has been a major topic of study in functional analysis, going as far back as Grothendieck's famous "Resume" [Gro53]. It has many striking applications, like the isomorphic characterization of Hilbert spaces and L_p spaces due to Kwapień [Kwa72a, Kwa72b], connections to type and cotype through the work of Kwapień [Kwa72a], Rosenthal [Ros73], Maurey [Mau74] and Pisier [Pis80], connections to Sidon sets through the work of Pisier [Pis86], characterization of weakly compact operators due to Davis et al. [DFJP74], connections to the theory of p-summing operators through the work of Grothendieck [Gro53], Pietsch [Pie67] and Lindenstrauss and Pelczynski [LP68].

Let $\Phi(A)$ denote

$$\Phi(A) := \inf_{H} \inf_{BC=A} \frac{\|C\|_{X\to H} \cdot \|B\|_{H\to Y}}{\|A\|_{X\to Y}}$$

where the infimum runs over all Hilbert spaces H. We say A factorizes through a Hilbert space if $\Phi(A) < \infty$. Further, let

$$\Phi(X, Y) := \sup_{A} \Phi(A)$$

where the supremum runs over continuous operators $A : X \to Y$. As a quick example of the power of factorization theorems, observe that if $\mathrm{I} : X \to X$ is the identity operator on a Banach space X and $\Phi(\mathrm{I}) < \infty$, then X is isomorphic to a Hilbert space and moreover the distortion (Banach-Mazur distance) is at most $\Phi(\mathrm{I})$ (i.e., there exists an invertible operator $T : X \to H$ for some Hilbert space H such that $\|T\|_{X\to H} \cdot \|T^{-1}\|_{H\to X} \leq \Phi(\mathrm{I})$). In fact (as observed by Maurey), Kwapień gave an isomorphic characterization of Hilbert spaces by proving a factorization theorem. Maurey observed that a more general factorization result underlies Kwapień's work:

Theorem 3.8.1 (Kwapień-Maurey)**.** *Let X be a Banach space of Type-2 and Y be a Banach space of Cotype-2. Then any operator $T : X \to Y$ factorizes through a Hilbert space. Moreover $\Phi(X, Y) \leq T_2(X) C_2(Y)$.*

Surprisingly Grothendieck's work which predates the work of Kwapień and Maurey, established that $\Phi(\ell_\infty^n, \ell_1^m) \leq K_G$ for all $m, n \in \mathbb{N}$, which is not implied by the above theorem since $T_2(\ell_\infty^n) \to \infty$ as $n \to \infty$. Pisier [Pis80] unified the above results for the case of approximable operators by proving the following:

Theorem 3.8.2 (Pisier)**.** *Let X, Y be Banach spaces such that X^*, Y are of Cotype-2. Then any approximable operator $T : X \to Y$ factorizes through a Hilbert space. Moreover $\Phi(T) \leq (2\, C_2(X^*) C_2(Y))^{3/2}$.*

In the Chapter 5 we show that for any $p^*, q \in [1, 2]$, any $m, n \in \mathbb{N}$

$$\Phi(\ell_p^n, \ell_q^m) \leq \frac{1 + \varepsilon_0}{\sinh^{-1}(1)} \cdot C_2(\ell_{p^*}^n) \cdot C_2(\ell_q^m)$$

which improves upon Pisier's bound and for certain ranges of (p, q), improves upon K_G as well as the bound of Kwapień-Maurey.

Chapter 4

Hardness results for $p \to q$ norm

In this chapter we prove NP-hardness results for approximating hypercontractive norms (i.e., $p \to q$ norm when $p \leq q$). We show

Theorem 4.0.1. *For any p, q such that $1 < p \leq q < 2$ or $2 < p \leq q < \infty$ and a constant $c > 1$, unless $\mathsf{NP} \in \mathsf{BPP}$, no polytime algorithm approximates $p \to q$ norm within a factor of c. The reduction runs in time $n^{B_p q}$ for $2 < p < q$, where $B_p = \mathrm{poly}(1/(1 - \gamma_{p^*}))$.*

We show that the above hardness can be strengthened to any constant factor via a simple tensoring argument. In fact, this also shows that it is hard to approximate $\|A\|_{p \to q}$ within almost polynomial factors unless NP is in randomized quasi-polynomial time. This is the content of the following theorem.

Theorem 4.0.2. *For any p, q such that $1 < p \leq q < 2$ or $2 < p \leq q < \infty$ and $\varepsilon > 0$, there is no polynomial time algorithm that approximates the $p \to q$ norm of an $n \times n$ matrix within a factor $2^{\log^{1-\varepsilon} n}$ unless $\mathsf{NP} \subseteq \mathsf{BPTIME}\left(2^{(\log n)^{O(1)}}\right)$. When q is an even integer, the same inapproximability result holds unless $\mathsf{NP} \subseteq \mathsf{DTIME}\left(2^{(\log n)^{O(1)}}\right)$*

En route to the above result, we also prove new results for the case when $p \geq q$ with $2 \in [q, p]$:

Theorem 4.0.3. *For any p, q such that $p \geq 2 \geq q$ and $\varepsilon > 0$, it is NP-hard to approximate the $p \to q$ norm within a factor $1/(\gamma_{p^*} \gamma_q) - \varepsilon$.*

where γ_r denotes the r^{th} norm of a standard normal random variable, and $p^* := p/(p-1)$ is the dual norm of p.

Both Theorem 4.0.1 and Theorem 4.0.3 are consequences of a more technical theorem, which proves hardness of approximating $\|A\|_{2 \to r}$ for $r < 2$ (and hence $\|A\|_{r^* \to 2}$ for $r^* > 2$) while providing additional structure in the matrix A produced by the reduction. We also show our methods can be used to provide a simple proof (albeit via randomized reductions) of the $2^{\Omega((\log n)^{1-\varepsilon})}$ hardness for the non-hypercontractive case when $2 \notin [q, p]$, which was proved by [BV11].

See Fig. 2.1 for a pictorial summary of the hardness and algorithmic results in various regimes.

4.1 Proof Overview

The hardness of proving hardness for hypercontractive norms. Reductions for various geometric problems use a "smooth" version of the Label Cover problem, composed with long-code functions for the labels of the variables. In various reductions, including the ones by Guruswami et al. [GRSW16] and Briët et al. [BRS15] (which we closely follow) the solution vector x to the geometric problem consists of the Fourier coefficients of the various long-code functions, with a "block" x_v for each vertex of the label-cover instance. The relevant geometric operation (transformation by the matrix A in our case) consists of projecting to a space which enforces the consistency constraints derived from the label-cover problem, on the Fourier coefficients of the encodings.

However, this strategy presents with two problems when designing reductions for hypercontractive norms. Firstly, while projections maintain the ℓ_2 norm of encodings corresponding to consistent labelings and reduce that of inconsistent ones, their behaviour is harder to analyze for ℓ_p norms for $p \neq 2$. Secondly, the *global* objective of maximizing $\|Ax\|_q$ is required to enforce different behavior within the blocks x_v, than in the full vector x. The block vectors x_v in the solution corresponding to a satisfying assignment of label cover are intended to be highly sparse, since they correspond to "dictator functions" which have only one non-zero Fourier coefficient. This can be enforced in a test using the fact that for a vector $x_v \in \mathbb{R}^t$, $\|x_v\|_q$ is a convex function of $\|x_v\|_p$ when $p \leq q$, and is maximized for vectors with all the mass concentrated in a single coordinate. However, a global objective function which tries to maximize $\sum_v \|x_v\|_q^q$, also achieves a high value from global vectors x which concentrate all the mass on coordinates corresponding to few vertices of the label cover instance, and do not carry any meaningful information about assignments to the underlying label cover problem.

Since we can only check for a global objective which is the ℓ_q norm of some vector involving coordinates from blocks across the entire instance, it is not clear how to enforce local Fourier concentration (dictator functions for individual long codes) and global well-distribution (meaningful information regarding assignments of most vertices) using the same objective function. While the projector A also enforces a linear relation between the block vectors x_u and x_v for all edges (u, v) in the label cover instance, using this to ensure well-distribution across blocks seems to require a very high density of constraints in the label cover instance, and no hardness results are available in this regime.

Our reduction. We show that when $2 \notin [p, q]$, it is possible to bypass the above issues using hardness of $\|A\|_{2\to r}$ as an intermediate (for $r < 2$). Note that since $\|z\|_r$ is a *concave* function of $\|z\|_2$ in this case, the test favors vectors in which the mass is well-distributed and thus solves the second issue. For this, we use local tests based on the Berry-Esséen theorem (as in [GRSW16] and [BRS15]). Also, since the starting point now is the ℓ_2 norm, the effect of projections is easier to analyze. This reduction is discussed in Section 4.3.

By duality, we can interpret the above as a hardness result for $\|A\|_{p\to 2}$ when $p > 2$ (using $r = p^*$). We then convert this to a hardness result for $p{\to}q$ norm in the hypercontractive case by composing A with an "approximate isometry" B from $\ell_2 \to \ell_q$ (i.e., $\forall y \; \|By\|_q \approx \|y\|_2$) since we can replace $\|Ax\|_2$ with $\|BAx\|_q$. Milman's version of the

Dvoretzky theorem [Ver17] implies random operators to a sufficiently high dimensional ($n^{O(q)}$) space satisfy this property, which then yields constant factor hardness results for the $p \to q$ norm.

We also show that the hardness for hypercontractive norms can be amplified via tensoring. This was known previously for the $2 \to 4$ norm using an argument based on parallel repetition for QMA [HM13], and for the case of $p = q$ [BV11]. We give a simple argument based on convexity, which proves this for all $p \le q$, but appears to have gone unnoticed previously. The amplification is then used to prove hardness of approximation within almost polynomial factors.

Non-hypercontractive norms. We also use the hardness of $\|A\|_{2 \to r}$ to obtain hardness for the non-hypercontractive case of $\|A\|_{p \to q}$ with $q < 2 < p$, by using an operator that "factorizes" through ℓ_2. In particular, we obtain hardness results for $\|A\|_{p \to 2}$ and $\|A\|_{2 \to q}$ (of factors $1/\gamma_{p^*}$ and $1/\gamma_q$ respectively) using the reduction in Section 4.3. We then combine these hardness results using additional properties of the operator A obtained in the reduction, to obtain a hardness of factor $(1/\gamma_{p^*}) \cdot (1/\gamma_q)$ for the $p \to q$ norm for $p > 2 > q$. The composition, as well as the hardness results for hypercontractive norms, are presented in Section 4.4.

We also obtain a simple proof of the $2^{\Omega((\log n)^{1-\varepsilon})}$ hardness for the non-hypercontractive case when $2 \notin [q, p]$ (already proved by Bhaskara and Vijayaraghavan [BV11]) via an approximate isometry argument as used in the hypercontractive case. In the hypercontractive case, we started from a constant factor hardness of the $p \to 2$ norm and the same factor for $p \to q$ norm using the fact that for a random Gaussian matrix B of appropriate dimensions, we have $\|Bx\|_q \approx \|x\|_2$ for all x. We then amplify the hardness via tensoring. In the non-hypercontractive case, we start with a hardness for $p \to p$ norm (obtained via the above isometry), which we *first* amplify via tensoring. We then apply another approximate isometry result due to Schechtman [Sch87], which gives a samplable distribution $\mathcal{D}$ over random matrices B such that with high probability over B, we have $\|Bx\|_q \approx \|x\|_p$ for all x.

We illustrate in this chapter how combining hardness for $p \to 2$ norm, with geometric principles like duality, tensoring, composition and embedding yields strong results in both the hypercontractive and non-hypercontractive regimes.

4.2 Preliminaries and Notation

For a vector $x \in \mathbb{R}^n$, exclusively this chapter we will use $x(i)$ to denote its i-th coordinate since in certain situations it will be convenient to think of vectors as functions. For $p \in [1, \infty)$, we define $\|\cdot\|_{\ell_p}$ to denote the counting p-norm and $\|\cdot\|_{L_p}$ to denote the expectation p-norm; i.e., for a vector $x \in \mathbb{R}^n$,

$$\|x\|_{\ell_p} := \left(\sum_{i \in [n]} |x(i)|^p \right)^{1/p} \quad \text{and} \quad \|x\|_{L_p} := \mathop{\mathbb{E}}_{i \sim [n]} \left[|x(i)|^p \right]^{1/p} = \left(\frac{1}{n} \cdot \sum_{i \in [n]} |x(i)|^p \right)^{1/p}.$$

While $\|\cdot\|_{\ell_p}$ and $\|x\|_p$ both denote the counting p-norm in this document, we will use $\|\cdot\|_{\ell_p}$ in this chapter to highlight the distinction between counting and expectation norms. Clearly $\|x\|_{\ell_p} = \|x\|_{L_p} \cdot n^{1/p}$. For $p = \infty$, we define $\|x\|_{\ell_\infty} = \|x\|_{L_\infty} := \max_{i \in [n]} |x(i)|$. We will use p^* to denote the 'dual' of p, i.e. $p^* = p/(p-1)$. Unless stated otherwise, we usually work with $\|\cdot\|_{\ell_p}$. We also define inner product $\langle x, y \rangle$ to denote the inner product under the counting measure unless stated otherwise; i.e., for two vectors $x, y \in \mathbb{R}^n$, $\langle x, y \rangle := \sum_{i \in [n]} x(i) y(i)$.

4.2.1 Fourier Analysis

We introduce some basic facts about Fourier analysis of Boolean functions. Let $R \in \mathbb{N}$ be a positive integer, and consider a function $f : \{\pm 1\}^R \to \mathbb{R}$. For any subset $S \subseteq [R]$ let $\chi_S := \prod_{i \in S} x_i$. Then we can represent f as

$$f(x_1, \ldots, x_R) = \sum_{S \subseteq [R]} \widehat{f}(S) \cdot \chi_S(x_1, \ldots x_R), \tag{4.1}$$

where

$$\widehat{f}(S) = \mathbb{E}_{x \in \{\pm 1\}^R}[f(x) \cdot \chi_S(x)] \text{ for all } S \subseteq [R]. \tag{4.2}$$

The *Fourier transform* refers to a linear operator F that maps f to $\widehat{f}$ as defined as (4.2). We interpret $\widehat{f}$ as a 2^R-dimensional vector whose coordinates are indexed by $S \subseteq [R]$. Endow the expectation norm and the expectation norm to f and $\widehat{f}$ respectively; i.e.,

$$\|f\|_{L_p} := \left(\mathop{\mathbb{E}}_{x \in \{\pm 1\}^R} [|f(x)|^p] \right)^{1/p} \quad \text{and} \quad \|\widehat{f}\|_{\ell_p} := \left(\sum_{S \subseteq [R]} |\widehat{f}(S)|^p \right)^{1/p}.$$

as well as the corresponding inner products $\langle f, g \rangle$ and $\langle \widehat{f}, \widehat{g} \rangle$ consistent with their 2-norms. We also define the *inverse Fourier transform* F^T to be a linear operator that maps a given $\widehat{f} : 2^R \to \mathbb{R}$ to $f : \{\pm 1\}^R \to \mathbb{R}$ defined as in (4.1). We state the following well-known facts from Fourier analysis.

Observation 4.2.1 (Parseval's Theorem). *For any $f : \{\pm 1\}^R \to \mathbb{R}$, $\|f\|_{L_2} = \|Ff\|_{\ell_2}$.*

Observation 4.2.2. *F and F^T form an adjoint pair; i.e., for any $f : \{\pm 1\}^R \to \mathbb{R}$ and $\widehat{g} : 2^R \to \mathbb{R}$,*

$$\langle \widehat{g}, Ff \rangle = \langle F^T \widehat{g}, f \rangle.$$

Observation 4.2.3. *$F^T F$ is the identity operator.*

In Section 4.3, we also consider a *partial* Fourier transform F_P that maps a given function $f : \{\pm 1\}^R \to \mathbb{R}$ to a vector $\widehat{f} : [R] \to \mathbb{R}$ defined as $\widehat{f}(i) = \mathbb{E}_{x \in \{\pm 1\}^R}[f(x) \cdot x_i]$ for all $i \in [R]$. It is the original Fourier transform where $\widehat{f}$ is further projected to R coordinates corresponding to linear coefficients. The partial inverse Fourier transform F_P^T is a transformation that maps a vector $\widehat{f} : [R] \to \mathbb{R}$ to a function $f : \{\pm 1\}^R \to \mathbb{R}$ as in (4.1) restricted to $S = \{i\}$ for some $i \in [R]$. These partial transforms satisfy similar observations as above: (1) $\|f\|_{L_2} \geq \|F_P f\|_{\ell_2}$, (2) $\|F_P^T \widehat{f}\|_{L_2} = \|\widehat{f}\|_{\ell_2}$, (3) F_P and F_P^T form an adjoint pair, and (4) $(F_P^T F_P) f = f$ if and only if f is a linear function.

4.2.2 Smooth Label Cover

An instance of Label Cover is given by a quadruple $\mathcal{L} = (G, [R], [L], \Sigma)$ that consists of a regular connected graph $G = (V, E)$, a label set $[R]$ for some positive integer n, and a collection $\Sigma = ((\pi_{e,v}, \pi_{e,w}) : e = (v, w) \in E)$ of pairs of maps both from $[R]$ to $[L]$ associated with the endpoints of the edges in E. Given a *labeling* $\ell : V \to [R]$, we say that an edge $e = (v, w) \in E$ is *satisfied* if $\pi_{e,v}(\ell(v)) = \pi_{e,w}(\ell(w))$. Let $\mathsf{OPT}(\mathcal{L})$ be the maximum fraction of satisfied edges by any labeling.

The following hardness result for Label Cover, given in [GRSW16], is a slight variant of the original construction due to [Kho02]. The theorem also describes the various structural properties, including smoothness, that are identified by the hard instances.

Theorem 4.2.4. *For any $\xi > 0$ and $J \in \mathbb{N}$, there exist positive integers $R = R(\xi, J), L = L(\xi, J)$ and $D = D(\xi)$, and a Label Cover instance $(G, [R], [L], \Sigma)$ as above such that*

- *(Hardness): It is NP-hard to distinguish between the following two cases:*
  - *(Completeness):* $\mathsf{OPT}(\mathcal{L}) = 1$.
  - *(Soundness):* $\mathsf{OPT}(\mathcal{L}) \leq \xi$.
- *(Structural Properties):*
  - *(J-Smoothness): For every vertex $v \in V$ and distinct $i, j \in [R]$, we have*
$$\mathop{\mathbb{P}}_{e: v\in e}[\pi_{e,v}(i) = \pi_{e,v}(j)] \leq 1/J.$$
  - *(D-to-1): For every vertex $v \in V$, edge $e \in E$ incident on v, and $i \in [L]$, we have $|\pi_{e,v}^{-1}(i)| \leq D$; that is at most D elements in $[R]$ are mapped to the same element in $[L]$.*
  - *(Weak Expansion): For any $\delta > 0$ and vertex set $V' \subseteq V$ such that $|V'| = \delta \cdot |V|$, the number of edges among the vertices in $|V'|$ is at least $(\delta^2/2)|E|$.*

4.3 Hardness of $2 \to r$ norm with $r < 2$

This section proves the following theorem that serves as a starting point of our hardness results. The theorem is stated for the expectation norm for consistency with the current literature, but the same statement holds for the counting norm, since if A is an $n \times n$ matrix, $\|A\|_{\ell_2 \to \ell_r} = n^{1/r - 1/2} \cdot \|A\|_{L_2 \to L_r}$. Note that the matrix A used in the reduction below does not depend on r.

Theorem 4.3.1. *For any $\varepsilon > 0$, there is a polynomial time reduction that takes a 3-CNF formula φ and produces a symmetric matrix $A \in \mathbb{R}^{n\times n}$ with $n = |\varphi|^{\mathrm{poly}(1/\varepsilon)}$ such that*

- *(Completeness) If φ is satisfiable, there exists $x \in \mathbb{R}^n$ with $|x(i)| = 1$ for all $i \in [n]$ and $Ax = x$. In particular, $\|A\|_{L_2 \to L_r} \geq 1$ for all $1 \leq r \leq \infty$.*

- *(Soundness)* $\|A\|_{L_2\to L_r} \le \gamma_r + \varepsilon^{2-r}$ *for all* $1 \le r < 2$.

We adapt the proof by Briët, Regev and Saket for the hardness of $2 \to 1$ and $\infty \to 1$ norms to prove the above theorem. A small difference is that, unlike their construction which starts with a Fourier encoding of the long-code functions, we start with an evaluation table (to ensure that the resulting matrices are symmetric). We also analyze their dictatorship tests for the case of fractional r.

4.3.1 Reduction and Completeness

Let $\mathcal{L} = (G, [R], [L], \Sigma)$ be an instance of Label Cover with $G = (V, E)$. In the rest of this section, $n = |V|$ and our reduction will construct a self-adjoint linear operator $\mathbf{A} : \mathbb{R}^N \to \mathbb{R}^N$ with $N = |V| \cdot 2^R$, which yields a symmetric $N \times N$ matrix representing $\mathbf{A}$ in the standard basis. This section concerns the following four Hilbert spaces based on the standard Fourier analysis composed with $\mathcal{L}$.

1. Evaluation space $\mathbb{R}^{2^R}$. Each function in this space is denoted by $f : \{\pm 1\}^R \to \mathbb{R}$. The inner product is defined as $\langle f, g\rangle := \mathbb{E}_{x\in\{\pm1\}^R}[f(x)g(x)]$, which induces $\|f\|_2 := \|f\|_{L_2}$. We also define $\|f\|_{L_p} := \mathbb{E}_x[|f(x)|^p]^{1/p}$ in this space.

2. Fourier space $\mathbb{R}^R$. Each function in this space is denoted by $\widehat{f} : [R] \to \mathbb{R}$. The inner product is defined as $\langle \widehat{f}, \widehat{g}\rangle := \sum_{i\in[R]} \widehat{f}(i)\widehat{g}(i)$, which induces $\|\widehat{f}\|_2 := \|\widehat{f}\|_{\ell_2}$.

3. Combined evaluation space $\mathbb{R}^{V\times 2^R}$. Each function in this space is denoted by $\mathbf{f} : V \times \{\pm1\}^R \to \mathbb{R}$. The inner product is defined as $\langle \mathbf{f}, \mathbf{g}\rangle := \mathbb{E}_{v\in V}[\mathbb{E}_{x\in\{\pm1\}^R}[\mathbf{f}(v,x)\mathbf{g}(v,x)]]$, which induces $\|\mathbf{f}\|_{L_2} := \|\mathbf{f}\|_{L_2}$. We also define $\|\mathbf{f}\|_p := \mathbb{E}_{v,x}[|\mathbf{f}(v,x)|^p]^{1/p}$ in this space.

4. Combined Fourier space $\mathbb{R}^{V\times R}$. Each function in this space is denoted by $\widehat{\mathbf{f}} : V \times [R] \to \mathbb{R}$. The inner product is defined as $\langle \widehat{\mathbf{f}}, \widehat{\mathbf{g}}\rangle := \mathbb{E}_{v\in V}[\sum_{i\in[R]} \widehat{\mathbf{f}}(v,i)\widehat{\mathbf{g}}(v,i)]$, which induces $\|\widehat{\mathbf{f}}\|_2$, which is neither a counting nor an expectation norm.

Note that $\mathbf{f} \in \mathbb{R}^{V\times 2^R}$ and a vertex $v \in V$ induces $f_v \in \mathbb{R}^{2^R}$ defined by $f_v(x) := \mathbf{f}(v, x)$, and similarly $\widehat{\mathbf{f}} \in \mathbb{R}^{V\times R}$ and a vertex $v \in V$ induces $\widehat{f_v} \in \mathbb{R}^R$ defined by $\widehat{f_v}(x) := \widehat{\mathbf{f}}(v, x)$. As defined in Section 4.2.1, we use the standard following (partial) *Fourier transform* F that maps $f \in \mathbb{R}^{2^R}$ to $\widehat{f} \in \mathbb{R}^R$ as follows. [1]

$$\widehat{f}(i) = (Ff)(i) := \mathop{\mathbb{E}}_{x\in\{\pm1\}^R} [x_i f(x)]\,. \tag{4.3}$$

The (partial) *inverse Fourier transform* F^T that maps $\widehat{f} \in \mathbb{R}^R$ to $f \in \mathbb{R}^{2^R}$ is defined by

$$f(x) = (F^T\widehat{f})(x) := \sum_{i\in[R]} x_i \widehat{f}(i). \tag{4.4}$$

[1] We use only *linear Fourier coefficients* in this work. F was defined as F_P in Section 4.2.1.

This Fourier transform can be naturally extended to combined spaces by defining $\mathbf{F} : \mathbf{f} \mapsto \widehat{\mathbf{f}}$ as $f_v \mapsto \widehat{f_v}$ for all $v \in V$. Then $\mathbf{F}^T$ maps $\widehat{\mathbf{f}}$ to $\mathbf{f}$ as $\widehat{f_v} \mapsto f_v$ for all $v \in V$.

Finally, let $\widehat{\mathbf{P}} : \mathbb{R}^{V \times R} \to \mathbb{R}^{V \times R}$ be the orthogonal projector to the following subspace of the combined Fourier space:

$$\widehat{\mathbf{L}} := \left\{ \widehat{\mathbf{f}} \in \mathbb{R}^{V \times R} : \sum_{j \in \pi_{e,u}^{-1}(i)} \widehat{f_u}(i) = \sum_{j \in \pi_{e,v}^{-1}(i)} \widehat{f_v}(j) \text{ for all } (u,v) \in E \text{ and } i \in [L] \right\}. \tag{4.5}$$

Our transformation $\mathbf{A} : \mathbb{R}^{V \times 2^R} \to \mathbb{R}^{V \times 2^R}$ is defined by

$$\mathbf{A} := (\mathbf{F}^T)\widehat{\mathbf{P}}\mathbf{F}. \tag{4.6}$$

In other words, given $\mathbf{f}$, we apply the Fourier transform for each $v \in V$, project the combined Fourier coefficients to $\widehat{\mathbf{L}}$ that checks the Label Cover consistency, and apply the inverse Fourier transform. Since $\widehat{\mathbf{P}}$ is a projector, $\mathbf{A}$ is self-adjoint by design.

We also note that a similar reduction that produces $(\mathbf{F}^T)\widehat{\mathbf{P}}$ was used in Guruswami et al. [GRSW16] and Briët et al. [BRS15] for subspace approximation and Grothendieck-type problems, and indeed this reduction suffices for Theorem 4.3.1 except the self-adjointness and additional properties in the completeness case.

Completeness. We prove the following lemma for the completeness case. A simple intuition is that if $\mathcal{L}$ admits a good labeling, we can construct a $\mathbf{f}$ such that each f_v is a linear function and $\widehat{\mathbf{f}}$ is already in the subspace $\widehat{\mathbf{L}}$. Therefore, each of Fourier transform, projection to $\widehat{\mathbf{L}}$, and inverse Fourier transform does not really change $\mathbf{f}$.

Lemma 4.3.2 (Completeness). *Let $\ell : V \to [R]$ be a labeling that satisfies every edge of $\mathcal{L}$. There exists a function $\mathbf{f} \in \mathbb{R}^{V \times 2^R}$ such that $\mathbf{f}(v,x)$ is either $+1$ or -1 for all $v \in V, x \in \{\pm 1\}^R$ and $\mathbf{A}\mathbf{f} = \mathbf{f}$.*

Proof. Let $\mathbf{f}(v,x) := x_{\ell(v)}$ for every $v \in V, x \in \{\pm 1\}^R$. Consider $\widehat{\mathbf{f}} = \mathbf{F}\mathbf{f}$. For each vertex $v \in V$, $\widehat{\mathbf{f}}(v,i) = \widehat{f_v}(i) = 1$ if $i = \ell(v)$ and 0 otherwise. Since ℓ satisfies every edge of $\mathcal{L}$, $\widehat{\mathbf{f}} \in \widehat{\mathbf{L}}$ and $\widehat{\mathbf{P}}\widehat{\mathbf{f}} = \widehat{\mathbf{f}}$. Finally, since each f_v is a linear function, the partial inverse Fourier transform F^T satisfies $(F^T)\widehat{f_v} = f_v$, which implies that $(\mathbf{F}^T)\widehat{\mathbf{f}} = \mathbf{f}$. Therefore, $\mathbf{A}\mathbf{f} = (\mathbf{F}^T\widehat{\mathbf{P}}\mathbf{F})\mathbf{f} = \mathbf{f}$. ∎

4.3.2 Soundness

We prove the following soundness lemma. This finishes the proof of Theorem 4.3.1 since Theorem 4.2.4 guarantees NP-hardness of Label Cover for arbitrarily small $\xi > 0$ and arbitrarily large $J \in \mathbb{N}$.

Lemma 4.3.3 (Soundness). *For every $\varepsilon > 0$, there exist $\xi > 0$ (that determines $D = D(\xi)$ as in Theorem 4.2.4) and $J \in \mathbb{N}$ such that if $\mathsf{OPT}(\mathcal{L}) \leq \xi$, $\mathcal{L}$ is D-to-1, and $\mathcal{L}$ is J-smooth, $\|\mathbf{A}\|_{L_2 \to L_r} \leq \gamma_r + 4\varepsilon^{2-r}$ for every $1 \leq r < 2$.*

Proof. Let $\mathbf{f} \in \mathbb{R}^{V \times 2^R}$ be an arbitrary vector such that $\|\mathbf{f}\|_{L_2} = 1$. Let $\widehat{\mathbf{f}} = \mathbf{F}\mathbf{f}$, $\widehat{\mathbf{g}} = \widehat{\mathbf{L}}\widehat{\mathbf{f}}$, and $\mathbf{g} = \mathbf{F}^T \widehat{\mathbf{g}}$ so that $\mathbf{g} = (\mathbf{F}^T \widehat{\mathbf{L}} \mathbf{F})\mathbf{f} = \mathbf{A}\mathbf{f}$. By Parseval's theorem, $\|\widehat{f}_v\|_{\ell_2} \leq \|f_v\|_{L_2}$ for all $v \in V$ and $\|\widehat{\mathbf{f}}\|_2 \leq \|\mathbf{f}\|_{L_2} \leq 1$. Since $\widehat{\mathbf{L}}$ is an orthogonal projection, $\|\widehat{\mathbf{g}}\|_2 \leq \|\widehat{\mathbf{f}}\|_2 \leq 1$. Fix $1 \leq r < 2$ and suppose

$$\|\mathbf{g}\|_{L_r}^r = \mathop{\mathbb{E}}_{v \in V}\left[\|g_v\|_{L_r}^r\right] \geq \gamma_r^r + 4\varepsilon^{2-r}. \tag{4.7}$$

Use Lemma 4.5.2 to obtain $\delta = \delta(\varepsilon)$ such that $\|g_v\|_{L_p}^p > (\gamma_p^p + \varepsilon)\|\widehat{g}_v\|_{\ell_2}^p$ implies $\|\widehat{g}\|_{\ell_4} > \delta\|\widehat{g}\|_{\ell_2}$ for all $1 \leq p < 2$ (so that δ does not depend on r), and consider

$$V_0 := \{v \in V : \|\widehat{g}_v\|_{\ell_4} > \delta\varepsilon \text{ and } \|\widehat{g}_v\|_{\ell_2} \leq 1/\varepsilon\}. \tag{4.8}$$

We prove the following lemma that lower bounds the size of V_0.

Lemma 4.3.4. *For $V_0 \subseteq V$ defined as in (4.8), we have $|V_0| \geq \varepsilon^2 |V|$.*

Proof. The proof closely follows the proof of Lemma 3.4 of [BRS15]. Define the sets

$$\begin{aligned} V_1 &= \{v \in V : \|\widehat{g}_v\|_{\ell_4} \leq \delta\varepsilon \text{ and } \|\widehat{g}_v\|_{\ell_2} < \varepsilon\}, \\ V_2 &= \{v \in V : \|\widehat{g}_v\|_{\ell_4} \leq \delta\varepsilon \text{ and } \|\widehat{g}_v\|_{\ell_2} \geq \varepsilon\}, \\ V_3 &= \{v \in V : \|\widehat{g}_v\|_{\ell_2} > 1/\varepsilon\}. \end{aligned}$$

From (4.7), we have

$$\sum_{v \in V_0} \|g_v\|_{L_r}^r + \sum_{v \in V_1} \|g_v\|_{L_r}^r + \sum_{v \in V_2} \|g_v\|_{L_r}^r + \sum_{v \in V_3} \|g_v\|_{L_r}^r \geq (\gamma_r^r + 4\varepsilon^{2-r})|V|. \tag{4.9}$$

We bound the four sums on the left side of (4.9) individually. Parseval's theorem and the fact that $r < 2$ implies $\|g_v\|_{L_r} \leq \|g_v\|_{L_2} = \|\widehat{g}_v\|_{\ell_2}$, and since $\|\widehat{g}_v\|_{\ell_2} \leq 1/\varepsilon$ for every $v \in V_0$, the first sum in (4.9) can be bounded by

$$\sum_{v \in V_0} \|g_v\|_{L_r}^r \leq |V_0|/\varepsilon^r. \tag{4.10}$$

Similarly, using the definition of V_1 the second sum in (4.9) is at most $\varepsilon^r |V|$. By Lemma 4.5.2, for each $v \in V_2$, we have $\|g_v\|_{L_r}^r \leq (\gamma_r^r + \varepsilon)\|\widehat{g}_v\|_{\ell_2}^r$. Therefore, the third sum in (4.9) is bounded as

$$\begin{aligned} \sum_{v \in V_2} \|g_v\|_{L_r}^r &\leq (\gamma_r^r + \varepsilon) \sum_{v \in V_2} \|\widehat{g}_v\|_{\ell_2}^r \\ &= (\gamma_r^r + \varepsilon)|V_2| \mathbb{E}_{v \in V_2}[\|\widehat{g}_v\|_{\ell_2}^r] \\ &\leq (\gamma_r^r + \varepsilon)|V_2| \mathbb{E}_{v \in V_2}[\|\widehat{g}_v\|_{\ell_2}^2]^{r/2} && \text{(By Jensen using } r < 2) \\ &= (\gamma_r^r + \varepsilon)|V_2| \left(\frac{\sum_{v \in V_2} \|\widehat{g}_v\|_{\ell_2}^2}{|V_2|}\right)^{r/2} \\ &\leq (\gamma_r^r + \varepsilon)|V_2|^{1-r/2}|V|^{r/2} && \Big(\sum_{v \in V_2} \|\widehat{g}_v\|_{\ell_2}^2 \leq \sum_{v \in V} \|\widehat{g}_v\|_{\ell_2}^2 \leq |V|\Big) \end{aligned}$$

$$\leq (\gamma_r^r + \varepsilon)|V|. \tag{4.11}$$

Finally, the fourth sum in (4.9) is bounded by

$$\begin{aligned}
\sum_{v \in V_3} \|g_v\|_{L_r}^r &\leq \sum_{v \in V_3} \|g_v\|_{L_2}^r && \text{(Since } r < 2\text{)} \\
&= \sum_{v \in V_3} \|\widehat{g}_v\|_{\ell_2}^r && \text{(By Parseval's theorem)} \\
&= \sum_{v \in V_3} \|\widehat{g}_v\|_{\ell_2}^{r-2} \|\widehat{g}_v\|_{\ell_2}^2 \\
&< \sum_{v \in V_3} \varepsilon^{2-r} \|\widehat{g}_v\|_{\ell_2}^2 && (\|\widehat{g}_v\|_{\ell_2} > 1/\varepsilon \text{ for } v \in V_3, \text{ and } r < 2) \\
&= \varepsilon^{2-r} \sum_{v \in V_3} \|\widehat{g}_v\|_{\ell_2}^2 \leq \varepsilon^{2-r}|V|. && (4.12)
\end{aligned}$$

Combining the above with (4.9) yields

$$\begin{aligned}
|V_0| &\geq \varepsilon^r \sum_{v \in V_0} \|g_v\|_{L_r}^r \\
&\geq \varepsilon^r \Big((\gamma_r^r + 4\varepsilon^{2-r})|V| - \varepsilon^r |V| - (\gamma_r^r + \varepsilon)|V| - \varepsilon^{2-r}|V| \Big) \\
&\geq \varepsilon^r \varepsilon^{2-r} |V| = \varepsilon^2 |V|, && (4.13)
\end{aligned}$$

where the last inequality uses the fact that $\varepsilon^{2-r} \geq \varepsilon \geq \varepsilon^r$. ∎

Therefore, $|V_0| \geq \varepsilon^2 |V|$ and every vertex of v satisfies $\|\widehat{g}_v\|_{\ell_4} > \delta\varepsilon$ and $\|\widehat{g}_v\|_{\ell_2} \leq 1/\varepsilon$. Using only these two facts together with $\widehat{\mathbf{g}} \in \widehat{\mathbf{L}}$, Briët et al. [BRS15] proved that if the smoothness parameter J is large enough given other parameters, $\mathcal{L}$ admits a labeling that satisfies a significant fraction of edges.

Lemma 4.3.5 (Lemma 3.6 of [BRS15])**.** *Let $\beta := \delta^2\varepsilon^3$. There exists an absolute constant $c' > 0$ such that if $\mathcal{L}$ is T-to-1 and $T/(c'\varepsilon^8\beta^4)$-smooth for some $T \in \mathbb{N}$, there is a labeling that satisfies at least $\varepsilon^8\beta^4/1024$ fraction of E.*

This finishes the proof of Lemma 4.3.3 by setting $\xi := \varepsilon^8\beta^4/1024$ and $J := D(\xi)/(c'\varepsilon^8\beta^4)$ with $D(\xi)$ defined in Theorem 4.2.4. Given a 3-SAT formula, φ, by the standard property of Smooth Label Cover, the size of the reduction is $|\varphi|^{O(J\log(1/\xi))} = |\varphi|^{\mathrm{poly}(1/\varepsilon)}$. ∎

4.4 Hardness of $p \to q$ norm

In this section, we prove the main results of the chapter. We prove Theorem 4.0.3 on hardness of approximating $p \to q$ norm when $p \geq 2 \geq q$, and Theorem 4.0.2 on hardness of approximating $p \to q$ norm when $2 < p < q$. By duality, the same hardness is implied for the case of $p < q < 2$.

Our result for $p \geq 2 \geq q$ in Section 4.4.1 follows from Theorem 4.3.1 using additional properties in the completeness case. For hypercontractive norms, we start by showing constant factor hardness via reduction from $p \to 2$ norm (see Section 4.4.2), and then amplify the hardness factor by using the fact that all hypercontractive norms productivize under Kronecker product, which we prove in Section 4.4.4.

4.4.1 Hardness for $p \geq 2 \geq q$

We use Theorem 4.3.1 to prove hardness of $p \to q$ norm for $p \geq 2 \geq q$, which proves Theorem 4.0.3.

Proof of Theorem 4.0.3: Fix p, q, and $\delta > 0$ such that $\infty \geq p \geq 2 \geq q$ and $p > q$. Our goal is to prove that $p \to q$ norm is NP-hard to approximate within a factor $1/(\gamma_{p^*}\gamma_q + \delta)$. For $2 \to q$ norm for $1 \leq q < 2$, Theorem 4.3.1 (with $\varepsilon \leftarrow \delta^{1/(2-q)}$) directly proves a hardness ratio of $1/(\gamma_q + \varepsilon^{2-q}) = 1/(\gamma_q + \delta)$. By duality, it also gives an $1/(\gamma_{p^*} + \delta)$ hardness for $p \to 2$ norm for $p > 2$.

For $p \to q$ norm for $p > 2 > q$, apply Theorem 4.3.1 with $\varepsilon = (\delta/3)^{\max(1/(2-p^*),1/(2-q))}$. It gives a polynomial time reduction that produces a symmetric matrix $A \in \mathbb{R}^{n \times n}$ given a 3-SAT formula φ. Our instance for $p \to q$ norm is $AA^T = A^2$.

- (Completeness) If φ is satisfiable, there exists $x \in \mathbb{R}^n$ such that $|x(i)| = 1$ for all $i \in [N]$ and $Ax = x$. Therefore, $A^2 x = x$ and $\|A^2\|_{L_p \to L_q} \geq 1$.
- (Soundness) If φ is not satisfiable,

$$\begin{aligned}\|A\|_{L_p \to L_2} &= \|A\|_{L_2 \to L_{p^*}} \leq \gamma_{p^*} + \varepsilon^{2-p^*} \leq \gamma_{p^*} + \delta/3, \text{ and} \\ \|A\|_{L_2 \to L_q} &\leq \gamma_q + \varepsilon^{2-q} \leq \gamma_q + \delta/3.\end{aligned}$$

  This implies that

$$\|A^2\|_{L_p \to L_q} \leq \|A\|_{L_p \to L_2} \|A\|_{L_2 \to L_q} \leq (\gamma_{p^*} + \delta/3)(\gamma_q + \delta/3) \leq \gamma_{p^*}\gamma_q + \delta\,.$$

This creates a gap of $1/(\gamma_{p^*}\gamma_q + \delta)$ between the completeness and the soundness case. The same gap holds for the counting norm since $\|A^2\|_{\ell_p \to \ell_q} = n^{1/q-1/p} \cdot \|A^2\|_{L_p \to L_q}$. ∎

4.4.2 Reduction from $p \to 2$ norm via Approximate Isometries

Let $A \in \mathbb{R}^{n \times n}$ be a hard instance of $p \to 2$ norm. For any $q \geq 1$, if a matrix $B \in \mathbb{R}^{m \times n}$ satisfies $\|Bx\|_{\ell_q} = (1 \pm o(1))\|x\|_{\ell_2}$ for all $x \in \mathbb{R}^n$, then $\|BA\|_{p \to q} = (1 \pm o(1))\|A\|_{p \to 2}$. Thus BA will serve as a hard instance for $p \to q$ norm if one can compute such a matrix B efficiently. In fact, a consequence of the Dvoretzky-Milman theorem is that a sufficiently tall random matrix B satisfies the aforementioned property with high probability. In other words, for $m = m(q, n)$ sufficiently large, a random linear operator from ℓ_2^n to ℓ_q^m is an approximate isometry.

To restate this from a geometric perspective, for $m(q,n)$ sufficiently larger than n, a random section of the unit ball in ℓ_q^m is approximately isometric to the unit ball in ℓ_2^n. In the interest of simplicity, we will instead state and use a corollary of the following matrix deviation inequality due to Schechtman (see [Sch06], Chapter 11 in [Ver17]).

Theorem 4.4.1 (Schechtman [Sch06])**.** *Let B be an $m \times n$ matrix with i.i.d. $\mathcal{N}(0,1)$ entries. Let $f : \mathbb{R}^m \to \mathbb{R}$ be a positive-homogeneous and subadditive function, and let b be such that $f(y) \le b\|y\|_{\ell_2}$ for all $y \in \mathbb{R}^m$. Then for any $T \subset \mathbb{R}^n$,*

$$\sup_{x\in T} |f(Bx) - \mathbb{E}\left[f(Bx)\right]| = O(b \cdot \gamma(T) + t \cdot \mathrm{rad}(T))$$

with probability at least $1 - e^{-t^2}$, where $\mathrm{rad}(T)$ is the radius of T, and $\gamma(T)$ is the Gaussian complexity of T defined as

$$\gamma(T) := \mathop{\mathbb{E}}_{g\sim\mathcal{N}(0,I_n)}\left[\sup_{t\in T} |\langle g, t\rangle|\right]$$

The above theorem is established by proving that the random process given by $X_x := f(Bx) - \mathbb{E}[f(Bx)]$ has sub-gaussian increments with respect to L_2 and subsequently appealing to Talagrand's Comparison tail bound.

We will apply this theorem with $f(\cdot) = \|\cdot\|_{\ell_q}$, $b = 1$ and T being the unit ball under $\|\cdot\|_{\ell_2}$. We first state a known estimate of $\mathbb{E}[f(Bx)] = \mathbb{E}[\|Bx\|_{\ell_q}]$ for any fixed x satisfying $\|x\|_{\ell_2} = 1$. Note that when $\|x\|_{\ell_2} = 1$, Bx has the same distribution as an m-dimensional random vector with i.i.d. $\mathcal{N}(0,1)$ coordinates.

Theorem 4.4.2 (Biau and Mason [BM15])**.** *Let $X \in \mathbb{R}^m$ be a random vector with i.i.d. $\mathcal{N}(0,1)$ coordinates. Then for any $q \ge 2$,*

$$\mathbb{E}\left[\|X\|_{\ell_q}\right] = m^{1/q} \cdot \gamma_q + O(m^{(1/q)-1}).$$

We are now equipped to see that a tall random Gaussian matrix is an approximate isometry (as a linear map from ℓ_2^n to ℓ_q^m) with high probability.

Corollary 4.4.3. *Let B be an $m \times n$ matrix with i.i.d. $\mathcal{N}(0,1)$ entries where $m = \omega(n^{q/2})$. Then with probability at least $1 - e^{-n}$, every vector $x \in \mathbb{R}^n$ satisfies,*

$$\|Bx\|_{\ell_q} = (1 \pm o(1)) \cdot m^{1/q} \cdot \gamma_q \cdot \|x\|_{\ell_2}.$$

Proof. We apply Theorem 4.4.1 with function f being the ℓ_q norm, $b = 1$, and $t = \sqrt{n}$. Further we set T to be the ℓ_2 unit sphere, which yields $\gamma(T) = \Theta(\sqrt{n})$ and $\mathrm{rad}(T) = 1$. Applying Theorem 4.4.2 yields that with probability at least $1 - e^{t^2} = 1 - e^{-n}$, for all x with $\|x\|_{\ell_2} = 1$, we have

$$\begin{aligned}\left|\|Bx\|_{\ell_q} - m^{1/q}\cdot\gamma_q\right| &\le \left|\|Bx\|_{\ell_q} - \mathbb{E}\left[\|X\|_{\ell_q}\right]\right| + \left|\mathbb{E}\left[\|X\|_{\ell_q}\right] - m^{1/q}\cdot\gamma_q\right| \\ &\le O(b\cdot\gamma(T) + t\cdot\mathrm{rad}(T) + m^{(1/q)-1}) \\ &\le O(\sqrt{n} + \sqrt{n} + m^{(1/q)-1}) \\ &\le o(m^{1/q}).\end{aligned}$$

■

We thus obtain the desired constant factor hardness:

Proposition 4.4.4. *For any $p > 2$, $2 \leq q < \infty$ and any $\varepsilon > 0$, there is no polynomial time algorithm that approximates $p \rightarrow q$ norm (and consequently $q^* \rightarrow p^*$ norm) within a factor of $1/\gamma_{p^*} - \varepsilon$ unless NP $\not\subseteq$ BPP.*

Proof. By Corollary 4.4.3, for every $n \times n$ matrix A and a random $m \times n$ matrix B with i.i.d. $\mathcal{N}(0,1)$ entries ($m = \omega(n^{q/2})$), with probability at least $1 - e^{-n}$, we have

$$\|BA\|_{\ell_p \to \ell_q} = (1 \pm o(1)) \cdot \gamma_q \cdot m^{1/q} \cdot \|A\|_{\ell_p \to \ell_2}.$$

Thus the reduction $A \rightarrow BA$ combined with $p \rightarrow 2$ norm hardness implied by Theorem 4.3.1, yields the claim. ■

The generality of the concentration of measure phenomenon underlying the proof of the Dvoretzky-Milman theorem allows us to generalize Proposition 4.4.4, to obtain constant factor hardness of maximizing various norms over the ℓ_p ball ($p > 2$). In this more general version, the strength of our hardness assumption is dependent on the Gaussian width of the dual of the norm being maximized. Its proof is identical to that of Proposition 4.4.4.

Theorem 4.4.5. *Consider any $p > 2, \varepsilon > 0$, and any family $(f_m)_{m \in \mathbb{N}}$ of positive-homogeneous and subadditive functions where $f_m : \mathbb{R}^m \to \mathbb{R}$. Let $(b_m)_{m \in \mathbb{N}}$ be such that $f_m(y) \leq b_m \cdot \|y\|_{\ell_2}$ for all y and let $N = N(n)$ be such that $\gamma_*(f_N) = \omega(b_N \cdot \sqrt{n})$, where*

$$\gamma^*(f_N) := \mathop{\mathbb{E}}_{g \sim \mathcal{N}(0, I_N)} [f_N(g)] .$$

Then unless NP $\not\subseteq$ BPTIME$(N(n))$, there is no polynomial time $(1/\gamma_{p^*} - \varepsilon)$-approximation algorithm for the problem of computing $\sup_{\|x\|_p = 1} f_m(Ax)$, given an $m \times n$ matrix A.

4.4.3 Derandomized Reduction

In this section, we show how to derandomize the reduction in Proposition 4.4.4 to obtain NP-hardness when $q \geq 2$ is an even integer and $p > 2$. Similarly to Section 4.4.2, given $A \in \mathbb{R}^{n \times n}$ as a hard instance of $p \rightarrow 2$ norm, our strategy is to construct a matrix $B \in \mathbb{R}^{m \times n}$ and output BA as a hard instance of $p \rightarrow q$ norm.

Instead of requiring B to satisfy $\|Bx\|_{\ell_q} = (1 \pm o(1))\|x\|_{\ell_2}$ for all $x \in \mathbb{R}^n$, we show that $\|Bx\|_{\ell_q} \leq (1 + o(1))\|x\|_{\ell_2}$ for all $x \in \mathbb{R}^n$ and $\|Bx\|_{\ell_q} \geq (1 - o(1))\|x\|_{\ell_2}$ when every coordinate of x has the same absolute value. Since Theorem 4.3.1 ensures that $\|A\|_{\ell_p \to \ell_2}$ is achieved by $x = Ax$ for such a well-spread x in the completeness case, BA serves as a hard instance for $p \rightarrow q$ norm.

We use the following construction of q-wise independent sets to construct such a B deterministically.

Theorem 4.4.6 (Alon, Babai, and Itai [ABI86]). *For any $k \in \mathbb{N}$, one can compute a set S of vectors in $\{\pm 1\}^n$ of size $O(n^{k/2})$, in time $n^{O(k)}$, such that the vector random variable Y obtained by sampling uniformly from S satisfies that for any $I \in \binom{[n]}{k}$, the marginal distribution $Y|_I$ is the uniform distribution over $\{\pm 1\}^k$.*

For a matrix B as above, a randomly chosen row behaves similarly to an n-dimensional Rademacher random vector with respect to $\|\cdot\|_{\ell_q}$.

Corollary 4.4.7. *Let $R \in \mathbb{R}^n$ be a vector random variable with i.i.d. Rademacher (± 1) coordinates. For any even integer $q \geq 2$, there is an $m \times n$ matrix B with $m = O(n^{q/2})$, computable in $n^{O(q)}$ time, such that for all $x \in \mathbb{R}^n$, we have*

$$\|Bx\|_{\ell_q} = m^{1/q} \cdot \mathop{\mathbb{E}}_{R} \left[\langle R, x\rangle^q\right]^{1/q}.$$

Proof. Let B be a matrix, the set of whose rows is precisely S. By Theorem 4.4.6,

$$\|Bx\|_{\ell_q}^q = \sum_{Y \in S} \langle Y, x\rangle^q = m \cdot \mathop{\mathbb{E}}_{R} \left[\langle R, x\rangle^q\right] . \qquad \blacksquare$$

We use the following two results that will bound $\|BA\|_{\ell_p \to \ell_q}$ for the completeness case and the soundness case respectively.

Theorem 4.4.8 (Stechkin [Ste61]). *Let $R \in \mathbb{R}^n$ be a vector random variable with i.i.d. Rademacher coordinates. Then for any $q \geq 2$ and any $x \in \mathbb{R}^n$ whose coordinates have the same absolute value,*

$$\mathbb{E}\left[\langle R, x\rangle\right]^{1/q} = (1 - o(1)) \cdot \gamma_q \|x\|_{\ell_2}.$$

Theorem 4.4.9 (Khintchine inequality [Haa81]). *Let $R \in \mathbb{R}^n$ be a vector random variable with i.i.d. Rademacher coordinates. Then for any $q \geq 2$ and any $x \in \mathbb{R}^n$,*

$$\mathbb{E}\left[\langle R, x\rangle^q\right]^{1/q} \leq \gamma_q \cdot \|x\|_{\ell_2}.$$

We finally prove the derandomimzed version of Proposition 4.4.4 for even $q \geq 2$.

Proposition 4.4.10. *For any $p > 2, \varepsilon > 0$, and any even integer $q \geq 2$, it is NP-hard to approximate $p \to q$ norm within a factor of $1/\gamma_{p^*} - \varepsilon$.*

Proof. Apply Theorem 4.3.1 with $r_1 \leftarrow p^*$ and $\varepsilon \leftarrow \varepsilon$. Given an instance φ of 3-SAT, Theorem 4.3.1 produces a symmetric matrix $A \in \mathbb{R}^{n \times n}$ in polynomial time as a hard instance of $p \to 2$ norm. Our instance for $p \to q$ norm is BA where B is the $m \times n$ matrix given by Corollary 4.4.7 with $m = O(n^{q/2})$.

- (Completeness) If φ is satisfiable, there exists a vector $x \in \{\pm \frac{1}{\sqrt{n}}\}^n$ such that $Ax = x$. So we have $\|BAx\|_{\ell_q} = \|Bx\|_{\ell_q} = (1 - o(1)) \cdot m^{1/q} \cdot \gamma_q$, where the last equality uses Corollary 4.4.7 and Theorem 4.4.8. Thus $\|BA\|_{\ell_p \to \ell_q} \geq (1 - o(1)) \cdot m^{1/q} \cdot \gamma_q$.

- (Soundness) If φ is not satisfiable, then for any x with $\|x\|_{\ell_p} = 1$,

$$\|BAx\|_{\ell_q} = m^{1/q} \cdot \mathop{\mathbb{E}}_{R} \left[\langle R, Ax\rangle^q\right]^{1/q} \leq m^{1/q} \cdot \gamma_q \cdot \|Ax\|_{\ell_2}$$
$$\leq m^{1/q} \cdot \gamma_q \cdot \|A\|_{\ell_p \to \ell_2} \leq m^{1/q} \cdot \gamma_q \cdot (\gamma_{p^*} - \varepsilon)$$

where the first inequality is a direct application of Theorem 4.4.9. $\blacksquare$

4.4.4 Hypercontractive Norms Productivize

We will next amplify our hardness results using the fact that hypercontractive norms productivize under the natural operation of Kronecker or tensor product. Bhaskara and Vijayraghavan [BV11] showed this for the special case of $p = q$ and the Harrow and Montanaro [HM13] showed this for 2→4 norm (via parallel repetition for QMA(2)). In this section we prove this claim whenever $p \leq q$.

Theorem 4.4.11. *Let A and B be $m_1 \times n_1$ and $m_2 \times n_2$ matrices respectively. Then for any $1 \leq p \leq q < \infty$, $\|A \otimes B\|_{\ell_p \to \ell_q} \leq \|A\|_{\ell_p \to \ell_q} \cdot \|B\|_{\ell_p \to \ell_q}$.*

Proof. We will begin with some notation. Let a_i, b_j respectively denote the i-th and j-th rows of A and B. Consider any $z \in \mathbb{R}^{[n_1] \times [n_2]}$ satisfying $\|z\|_{\ell_p} = 1$. For $k \in [n_1]$, let $z_k \in \mathbb{R}^{n_2}$ denote the vector given by $z_k(\ell) := z(k, \ell)$. For $j \in [m_2]$, let $\bar{z}_j \in \mathbb{R}^{n_1}$ denote the vector given by $\bar{z}_j(k) := \langle b_j, z_k \rangle$. Finally, for $k \in [n_1]$, let $\lambda_k := \|z_k\|_{\ell_p}^p$ and let $v_k \in \mathbb{R}^{m_2}$ be the vector given by $v_k(j) := |\bar{z}_j(k)|^p / \lambda_k$.

We begin by 'peeling off' A:

$$\begin{aligned} \|(A \otimes B)z\|_{\ell_q}^q &= \sum_{i,j} |\langle a_i \otimes b_j, z \rangle|^q = \sum_j \sum_i |\langle a_i, \bar{z}_j \rangle|^q \\ &= \sum_j \|A\bar{z}_j\|_{\ell_q}^q \\ &\leq \|A\|_{\ell_p \to \ell_q}^q \cdot \sum_j \|\bar{z}_j\|_{\ell_p}^q \\ &= \|A\|_{\ell_p \to \ell_q}^q \cdot \sum_j \left(\|\bar{z}_j\|_{\ell_p}^p \right)^{q/p} \end{aligned}$$

In the special case of $p = q$, the proof ends here since the expression is a sum of terms of the form $\|By\|_{\ell_p}^p$ and can thus be upper bounded term-wise by $\|B\|_{\ell_p \to \ell_p}^p \cdot \|z_k\|_{\ell_p}^p$ which sums to $\|B\|_{\ell_q \to \ell_p}^p$. To handle the case of $q > p$, we will use a convexity argument:

$$\begin{aligned} &\|A\|_{\ell_p \to \ell_q}^q \cdot \sum_j \left(\|\bar{z}_j\|_{\ell_p}^p \right)^{q/p} \\ = \ &\|A\|_{\ell_p \to \ell_q}^q \cdot \sum_j \left(\sum_k |\bar{z}_j(k)|^p \right)^{q/p} \\ = \ &\|A\|_{\ell_p \to \ell_q}^q \cdot \|\sum_k \lambda_k \cdot v_k\|_{\ell_{q/p}}^{q/p} && (|\bar{z}_j(k)|^p = \lambda_k v_k(j)) \\ \leq \ &\|A\|_{\ell_p \to \ell_q}^q \cdot \sum_k \lambda_k \cdot \|v_k\|_{\ell_{q/p}}^{q/p} && (\text{by convexity of } \|\cdot\|_{q/p}^{q/p} \text{ when } q \geq p) \\ \leq \ &\|A\|_{\ell_p \to \ell_q}^q \cdot \max_k \|v_k\|_{\ell_{q/p}}^{q/p} \end{aligned}$$

It remains to show that $\|v_k\|_{\ell_{q/p}}^{q/p}$ is precisely $\|Bz_k\|_{\ell_q}^q / \|z_k\|_{\ell_p}^q$.

$$\begin{aligned}
\|A\|_{\ell_p\to\ell_q}^q \cdot \max_k \|v_k\|_{\ell_{q/p}}^{q/p} &= \|A\|_{\ell_p\to\ell_q}^q \cdot \max_k \frac{1}{\|z_k\|_{\ell_p}^q} \cdot \sum_j |\bar{z}_j(k)|^q \\
&= \|A\|_{\ell_p\to\ell_q}^q \cdot \max_k \frac{1}{\|z_k\|_{\ell_p}^q} \cdot \sum_j |\langle b_j, z_k\rangle|^q \\
&= \|A\|_{\ell_p\to\ell_q}^q \cdot \max_k \frac{\|Bz_k\|_{\ell_q}^q}{\|z_k\|_{\ell_p}^q} \\
&\leq \|A\|_{\ell_p\to\ell_q}^q \cdot \|B\|_{\ell_p\to\ell_q}^q
\end{aligned}$$

Thus we have established $\|A \otimes B\|_{\ell_p\to\ell_q} \leq \|A\|_{\ell_p\to\ell_q} \cdot \|B\|_{\ell_p\to\ell_q}$. Lastly, the claim follows by observing that the statement is equivalent to the statement obtained by replacing the counting norms with expectation norms. ■

We finally establish super constant NP-Hardness of approximating $p\to q$ norm, proving Theorem 4.0.2.

Proof of Theorem 4.0.2: Fix $2 < p \leq q < \infty$. Proposition 4.4.4 states that there exists $c = c(p,q) > 1$ such that any polynomial time algorithm approximating the $p\to q$ norm of an $n \times n$-matrix A within a factor of c will imply $\mathsf{NP} \subseteq \mathsf{BPP}$. Using Theorem 4.4.11, for any integer $k \in \mathbb{N}$ and $N = n^k$, any polynomial time algorithm approximating the $p\to q$ norm of an $N \times N$-matrix $A^{\otimes k}$ within a factor of c^k implies that NP admits a randomized algorithm running in time $\mathrm{poly}(N) = n^{O(k)}$. Under $\mathsf{NP} \not\subseteq \mathsf{BPP}$, any constant factor approximation algorithm is ruled out by setting k to be a sufficiently large constant. For any $\varepsilon > 0$, setting $k = \log^{1/\varepsilon} n$ rules out an approximation factor of $c^k = 2^{O(\log^{1-\varepsilon} N)}$ unless $\mathsf{NP} \subseteq \mathsf{BPTIME}\left(2^{\log^{O(1)} n}\right)$.

By duality, the same statements hold for $1 < p \leq q < 2$. When $2 < p \leq q$ and q is an even integer, all reductions become deterministic due to Proposition 4.4.10. ■

4.4.5 A Simple Proof of Hardness for the Case $2 \notin [q, p]$

In this section, we show how to prove an almost-polynomial factor hardness for approximating $p\to q$ norm in the non-hypercontractive case when $2 > p \geq q$ (and the case $p \geq q > 2$ follows by duality). This result is already known from the work of Bhaskara and Vijayaraghavan [BV11]. We show how to obtain a more modular proof, composing our previous results with a simple embedding argument. However, while the reduction in [BV11] was deterministic, we will only give a randomized reduction below.

As in [BV11], we start with a strong hardness for the $p\to p$ norm, obtained in Theorem 4.0.2. While the reduction in [BV11] relied on special properties of the instance for $\ell_p\to\ell_p$ norm, we can simply use the following embedding result of Schechtman [Sch87] (phrased in a way convenient for our application).

Theorem 4.4.12 (Schechtman [Sch87], Theorem 5)**.** *Let $q < p < 2$ and $\varepsilon > 0$. Then, there exists a polynomial time samplable distribution $\mathcal{D}$ on random matrices in $\mathbb{R}^{m\times n}$ with $m = \Omega_\varepsilon(n^3)$, such that with probability $1 - o(1)$, we have for every $x \in R^n$, $\|Bx\|_{\ell_q} = (1 \pm \varepsilon) \cdot \|x\|_{\ell_p}$.*

In fact the distribution $\mathcal{D}$ is based on p-stable distributions. While the theorem in [Sch87] does not mention the high probability bound or samplability, it is easy to modify the proof to obtain there properties. We provide a proof sketch below for completeness. We note that Schechtman obtains a stronger bound of $O(n^{1+p/q})$ on the dimension m of the ℓ_q space, which requires a more sophisticated argument using "Lewis weights". However, we only state weaker $O(n^3)$ bound above, which suffices for our purposes and is easier to convert to a samplable distribution.

We first prove the following hardness result for approximating $p\to q$ norm in the reverse-hypercontractive case, using Theorem 4.4.12.

Theorem 4.4.13. *For any p, q such that $1 < q \le p < 2$ or $2 < q \le p < \infty$ and $\varepsilon > 0$, there is no polynomial time algorithm that approximates the $p\to q$ norm of an $n \times n$ matrix within a factor $2^{\log^{1-\varepsilon} n}$ unless $\mathsf{NP} \subseteq \mathsf{BPTIME}\left(2^{(\log n)^{O(1)}}\right)$.*

Proof. We consider the case $1 < q \le p < 2$ (the other case follows via duality). Theorem 4.0.2 gives a reduction from SAT on n variables, approximating the $p\to p$ norm of matrices $A \in \mathbb{R}^{N\times N}$ with $N = 2^{(\log n)^{O(1/\varepsilon)}}$, within a factor $2^{(\log N)^{1-\varepsilon}}$. Sampling a matrix B from the distribution $\mathcal{D}$ given by Theorem 4.4.12 (with dimension N) gives that it is also hard to approximate $\|BA\|_{p\to q} \approx \|A\|_{p\to p}$, within a factor $2^{(\log N)^{1-\varepsilon}}$. ■

We now give a sketch of the proof of Theorem 4.4.12 including the samplability condition. The key idea is to embed the space ℓ_p^n into the infinite-dimensional space L_q (for $0 \le q \le p < 2$) using p-stable random variables. The corresponding subspace of L_q can then be embedded into ℓ_q^m if the random variables (elements of L_q) constructed in the previous space are bounded in L_∞ norm. This is the content of the following claim.

Claim 4.4.14 (Schechtman [Sch87], Proposition 4)**.** *Let $\varepsilon > 0$ and Ω be an efficiently samplable probability space and let V be an n-dimensional subspace of $L_q(\Omega)$, such that*

$$M := \sup\left\{\|f\|_{L_\infty} \mid \|f\|_{L_q} \le 1, f \in V\right\} < \infty .$$

Then there exists a polynomial time samplable distribution $\mathcal{D}$ over linear operators $T : L_q(\Omega) \to \mathbb{R}^m$ for $m = C(\varepsilon, q) \cdot n \cdot M^q$ such that with probability $1 - o(1)$, we have that for every $f \in V$, $\|Tf\|_{\ell_q} = (1 \pm \varepsilon) \cdot \|f\|_{L_q}$.

Proof Sketch: The linear operator is simply defined by sampling $x_1, \ldots, x_m \sim \Omega$ independently, and taking

$$Tf := \frac{1}{m^{1/q}} \cdot (f(x_1), \ldots, f(x_m)) \qquad \forall f .$$

The proof then follows by concentration bounds for L_∞-bounded random variables, and a union bound over an epsilon net for the space V. □

The problem then reduces to constructing an embedding of ℓ_p^n into L_q, which is bounded in L_∞ norm. While a simple embedding can be constructed using p-stable distributions, Schechtman uses a clever reweighting argument to control the L_∞ norm. We show below that a simple truncation argument can also be used to obtain a somewhat crude bound on the L_∞ norm, which suffices for our purposes and yields an easily samplable distribution.

We collect below the relevant facts about p-stable random variables needed for our argument, which can be found in many well-known references, including [Ind06, AK06].

Fact 4.4.15. *For all $p \in (0,2)$, there exist (normalized) p-stable random variables Z satisfying the following properties:*

1. *For $Z_1, \ldots, Z_n$ iid copies of Z, and for all $a \in \mathbb{R}^n$, the random variable*

$$S \;:=\; \frac{a_1 \cdot Z_1 + \cdots + a_n \cdot Z_n}{\|a\|_{\ell_p}}\,,$$

   has distribution identical to Z.

2. *For all $q < p$, we have*

$$C_{p,q} \;:=\; \|Z\|_{L_q} \;=\; \left(\mathbb{E}\left[|Z|^q\right]\right)^{1/q} \;<\; \infty\,.$$

3. *There exists a constant C_p such that for all $t > 0$,*

$$\mathbb{P}\left[|Z| \geq t\right] \;<\; \frac{C_p}{t}\,.$$

4. *Z can be sampled by choosing $\theta \in_R [-\pi/2, \pi/2]$, $r \in_R [0,1]$, and taking*

$$Z \;=\; \frac{\sin(p\theta)}{(\cos(\theta))^{1/p}} \cdot \left(\frac{\cos((1-p)\cdot\theta)}{\ln(1/r)}\right)^{(1-p)/p}\,.$$

We now define an embedding of ℓ_p^n into L_q with bounded L_∞, using truncated p-stable random variables. Let $Z = (Z_1, \ldots, Z_n)$ be a vector of iid p-stable random variables as above, and let B be a parameter to be chosen later. We consider the random variables

$$\Delta(Z) \;:=\; \mathbb{1}_{\{\exists i \in [n]\; |Z_i| > B\}} \quad \text{and} \quad Y \;:=\; (1 - \Delta(Z)) \cdot Z \;=\; \mathbb{1}_{\{\forall i \in [n]\; |Z_i| \leq B\}} \cdot Z\,.$$

For all $a \in \mathbb{R}^n$, we define the (linear) embedding

$$\varphi(a) \;:=\; \frac{\langle a, Y\rangle}{C_{p,q}} \;=\; \frac{\langle a, Z\rangle}{C_{p,q}} - \Delta(Z) \cdot \frac{\langle a, Z\rangle}{C_{p,q}}\,.$$

By the properties of p-stable distributions, we know that $\|\langle a, Z\rangle / C_{p,q}\|_{L_q} = \|a\|_{\ell_p}$ for all $a \in \mathbb{R}^n$. By the following claim, we can choose B so that the second term only introduces a small error.

Claim 4.4.16. *For all $\varepsilon > 0$, there exists $B = O_{p,q,\varepsilon}(n^{1/p})$ such that for the embedding φ defined above*

$$\left| \|\varphi(a)\|_{L_q} - \|a\|_{\ell_p} \right| \;\leq\; \varepsilon \cdot \|a\|_{\ell_p} .$$

Proof. By triangle inequality, it suffices to bound $\|\Delta(Z) \cdot \langle a, Z\rangle\|_{L_q}$ by $\varepsilon \cdot C_{p,q} \cdot \|a\|_{\ell_p}$. Let $\delta > 0$ be such that $(1+\delta) \cdot q < p$. Using the fact that $\Delta(Z)$ is Boolean and Hölder's inequality, we observe that

$$\begin{aligned}
\|\Delta(Z) \cdot \langle a, Z\rangle\|_{L_q} &= \left(\mathbb{E}\left[|\langle a, Z\rangle|^q \cdot \Delta(Z)\right]\right)^{1/q} \\
&\leq \left(\mathbb{E}\left[|\langle a, Z\rangle|^{q(1+\delta)}\right]\right)^{1/(q(1+\delta))} \cdot \left(\mathbb{E}\left[\Delta(Z)\right]\right)^{\delta/(q(1+\delta))} \\
&= C_{p,(1+\delta)q} \cdot \|a\|_{\ell_p} \cdot \left(\mathbb{P}\left[\exists i \in [n] \;\; |Z_i| \geq B\right]\right)^{\delta/(q(1+\delta))} \\
&\leq C_{p,(1+\delta)q} \cdot \|a\|_{\ell_p} \cdot \left(n \cdot \frac{C_p}{B^p}\right)^{\delta/(q(1+\delta))}
\end{aligned}$$

Thus, choosing $B = O_{\varepsilon,p,q}(n^{1/p})$ such that

$$\frac{C_{p,(1+\delta)q}}{C_{p,q}} \cdot \left(n \cdot \frac{C_p}{B^p}\right)^{\delta/(q(1+\delta))} \;\leq\; \varepsilon$$

proves the claim. ∎

Using the value of B as above, we now observe a bound on $\|\varphi(a)\|_{L_\infty}$.

Claim 4.4.17. *Let $B = O_{\varepsilon,p,q}(n^{1/p})$ be chosen as above. Then, we have that*

$$M \;:=\; \sup\left\{ \|\langle a, Y\rangle\|_{L_\infty} \;\middle|\; \|\langle a, Y\rangle\|_{L_q} \leq 1 \right\} \;=\; O_{\varepsilon,p,q}(n) .$$

Proof. By the choice of B, we have that $\|\langle a, Y\rangle\|_{L_q} \geq (1-\varepsilon)\|a\|_{\ell_p}$. Thus, we can assume that $\|a\|_{\ell_p} \leq 2$. Hölder's inequality then gives for all such a,

$$\begin{aligned}
|\langle a, Y\rangle| &\leq \|a\|_{\ell_1} \cdot \|Y\|_{\ell_\infty} \\
&\leq n^{1-1/p} \cdot \|a\|_{\ell_p} \cdot B \\
&\leq 2 \cdot n^{1-1/p} \cdot B \;=\; O_{\varepsilon,p,q}(n) ,
\end{aligned}$$

which proves the claim. ∎

Using the above bound on M in Claim 4.4.14 gives a bound of $m = O_{\varepsilon,p,q}(n^{q+1}) = O_{\varepsilon,p,q}(n^3)$. Moreover, the distribution over embeddings is efficiently samplable, since it obtained by truncating p-stable random variables. This completes the proof of Theorem 4.4.12.

4.5 Dictatorship Test

First we prove an implication of Berry-Esséen estimate for fractional moments (similar to Lemma 3.3 of [GRSW16], see also [KNS10]).

Lemma 4.5.1. *There exist universal constants $c > 0$ and $\delta_0 > 0$ such that the following statement is true. If $X_1, \cdots, X_n$ are bounded independent random variables with $|X_i| \le 1$, $\mathbb{E}[X_i] = 0$ for $i \in [n]$, and $\sum_{i\in[n]} \mathbb{E}[X_i^2] = 1$, $\sum_{i\in[n]} \mathbb{E}[|X_i|^3] \le \delta$ for some $0 < \delta < \delta_0$, then for every $p \ge 1$:*

$$\left(\mathbb{E}\left[\left|\sum_{j=1}^{n} X_j\right|^p\right]\right)^{\frac{1}{p}} \le \gamma_p \cdot \left(1 + c\delta \left(\log\left(1/\delta\right)\right)^{\frac{p}{2}}\right).$$

Now we state and prove the main lemma of this section:

Lemma 4.5.2. *Let $f : \{\pm 1\}^R \to \mathbb{R}$ be a linear function for some positive integer $R \in \mathbb{N}$ and $\widehat{f} : [R] \to \mathbb{R}$ be its linear Fourier coefficients defined by*

$$\widehat{f}(i) := \mathop{\mathbb{E}}_{x\in\{\pm 1\}^R}\left[x_i f(x)\right].$$

For all $\varepsilon > 0$, there exists $\delta > 0$ such that if $\|f\|_{L_r} > (\gamma_r + \varepsilon)\|\widehat{f}\|_{\ell_2}$ then $\|\widehat{f}\|_{\ell_4} > \delta\|\widehat{f}\|_{\ell_2}$ for all $1 \le r < 2$.

Proof. We will prove this lemma by the method of contradiction. Let us assume $\|\widehat{f}\|_{\ell_4} \le \delta\|\widehat{f}\|_{\ell_2}$, for δ to be fixed later.

Let us define $y_i := \frac{\widehat{f}(i)}{\|\widehat{f}\|_{\ell_2}}$. Then, for all $x \in \{-1, 1\}^R$,

$$g(x) := \sum_{i\in[n]} x_i \cdot y_i = \frac{f(x)}{\|\widehat{f}\|_{\ell_2}}.$$

Let $Z_i = x_i \cdot y_i$ be the random variable when x_i is independently uniformly randomly chosen from $\{-1, 1\}$. Now

$$\sum_{i\in[n]} \mathbb{E}\left[Z_i^2\right] = \sum_{i\in[n]} \frac{\widehat{f}(i)^2}{\|\widehat{f}\|_{\ell_2}^2} = 1.$$

and

$$\sum_{i\in[n]} \mathbb{E}\left[|Z_i|^3\right] = \sum_{i\in[n]} \frac{\left|\widehat{f}(i)\right|^3}{\|\widehat{f}\|_{\ell_2}^3} = \sum_{i\in[n]} \frac{\left|\widehat{f}(i)\right|^2}{\|\widehat{f}\|_{\ell_2}^2} \cdot \frac{\left|\widehat{f}(i)\right|}{\|\widehat{f}\|_{\ell_2}} \le \frac{\|\widehat{f}\|_{\ell_4}^2}{\|\widehat{f}\|_{\ell_2}^2} \le \delta^2,$$

where the penultimate inequality follows from Cauchy-Schwarz ineqality.

Hence, by applying Lemma 4.5.1 on the random variables $Z_1, \cdots, Z_n$, we get:

$$\frac{\|f\|_{L_r}}{\|\widehat{f}\|_{\ell_2}} = \|g\|_{L_r} = \left(\mathop{\mathbb{E}}_{x\in\{-1,1\}^n}\left[|g(x)|^r\right]\right)^{\frac{1}{r}}$$
$$= \left(\mathop{\mathbb{E}}_{x\in\{-1,1\}^n}\left[\left|\sum_{i\in[n]} Z_i\right|^r\right]\right)^{\frac{1}{r}}$$
$$\leq \gamma_r\left(1 + c\delta^2\left(\log\frac{1}{\delta}\right)^r\right)$$

We choose $\delta > 0$ small enough (since $1 \leq r < 2$, setting $\delta < \frac{\sqrt{\varepsilon}}{\min(\delta_0, \sqrt{\gamma_2}\log\frac{c\gamma_2}{\varepsilon})} = \frac{\sqrt{\varepsilon}}{\min(\delta_0, \log\frac{c}{\varepsilon})}$ suffices) so that $\delta^2(\log\frac{1}{\delta})^r < \frac{\varepsilon}{c\gamma_r}$. For this choise of δ, we get: $\|f\|_{L_r} \leq (\gamma_r + \varepsilon)\|\widehat{f}\|_{\ell_2}$ – a contradiction. And hence the proof follows. ∎

Finally we prove Lemma 4.5.1:

Proof of Lemma 4.5.1: The proof is almost similar to that of Lemma 2.1 of [KNS10]. From Berry-Esséen theorem (see [vB72] for the constant), we get that:

$$\mathbb{P}\left[\left|\sum_{i=1}^n X_i\right| \geq u\right] \leq \mathbb{P}\left[|g| \geq u\right] + 2\sum_{i=1}^n \mathbb{E}\left[|X_i|^3\right] \leq \mathbb{P}\left[|g| \geq u\right] + 2\delta\,,$$

for every $u > 0$ and where $g \sim \mathcal{N}(0,1)$. By Hoeffding's lemma,

$$\mathbb{P}\left[\left|\sum_{i\in[n]} X_i\right| \geq t\right] < 2\mathrm{e}^{-2t^2}$$

for every $t > 0$. Combining the above observations, we get:

$$\mathbb{E}\left[\left|\sum_{i=1}^n X_i\right|^p\right] = \int_0^\infty pu^{p-1}\mathbb{P}\left[\left|\sum_{i=1}^n X_i\right| \geq u\right] du$$
$$\leq \int_0^a pu^{p-1}\mathbb{P}\left[|g| > u\right] du + 2\delta a^p + 2\int_a^\infty pu^{p-1}\mathrm{e}^{-2u^2} du$$
$$= \sqrt{\frac{2}{\pi}}\int_0^a u^p \mathrm{e}^{-u^2/2} du + 2\delta a^p + \frac{2p}{2^{\frac{p-1}{2}}}\int_{2a^2}^\infty z^{\frac{p+1}{2}-1}\mathrm{e}^{-z} dz$$
$$= \gamma_p^p - \sqrt{\frac{2}{\pi}}\int_a^\infty u^p \mathrm{e}^{-u^2/2} du + 2\delta a^p + \Gamma\left(\frac{p+1}{2}, 2a^2\right)\ ,$$

where $\Gamma(\cdot,\cdot)$ is the upper incomplete gamma function and a is a large constant determined later depending on δ and p. The second term is bounded as

$$\int_a^\infty u^p \mathrm{e}^{-u^2/2} du = a^{p-1}\mathrm{e}^{-a^2/2} + (p-1)\int_a^\infty u^{p-2}\mathrm{e}^{-u^2/2} du \leq a^{p-1}\mathrm{e}^{-a^2/2} + \frac{p-1}{a^2}\int_a^\infty u^p \mathrm{e}^{-u^2/2} du\,.$$

Hence $\int_a^\infty u^p \mathrm{e}^{-u^2/2} du \leq \frac{a^{p+1} e^{-a^2/2}}{1+a^2-p}$.

We know, $\Gamma(p+1/2, x) \to x^{\frac{p-1}{2}} \mathrm{e}^{-x}$ as $x \to \infty$. We choose $a = \gamma_p \sqrt{\log \frac{1}{\delta}}$. Hence there exists δ_0 so that for all small enough $\delta < \delta_0$, we have $\Gamma(p+1/2, 2a^2) \sim 2^{\frac{p-1}{2}} a^{p-1} \delta^{2\gamma_p^2} \ll \delta a^p$ where the last inequality follows from the fact that $2\gamma_p^2 > 1$ (as $p > 1$). Putting all this together, we get:

$$2\delta a^p + \Gamma\left(\frac{p+1}{2}, 2a^2\right) - \sqrt{\frac{2}{\pi}} \int_a^\infty u^p \mathrm{e}^{-u^2/2} du \ll 3\delta a^p - \sqrt{\frac{2}{\pi}} \frac{a^{p+1} e^{-a^2/2}}{1 + a^2 - p} \leq c\gamma_p^p \delta \left(\log \frac{1}{\delta}\right)^{p/2},$$

where c is an absolute constant independent of a and p. This completes the proof of the lemma. ∎

Chapter 5

Algorithmic results for $p \to q$ norm

In this chapter we generalize Krivine's Rounding algorithm for Grothendieck's constant to obtain improved approximation algorithms for $p \to q$ norm when $p \geq 2 \geq q$. Specifically we prove the following:

Theorem 5.0.1. *There exists a fixed constant $\varepsilon_0 \leq 0.00863$ such that for all $p \geq 2 \geq q$, the approximation ratio of the convex relaxation $\mathsf{CP}(A)$ is upper bounded by*

$$\frac{1+\varepsilon_0}{\sinh^{-1}(1)} \cdot \frac{1}{\gamma_{p^*} \cdot \gamma_q} \;=\; \frac{1+\varepsilon_0}{\ln(1+\sqrt{2})} \cdot \frac{1}{\gamma_{p^*} \cdot \gamma_q}\,.$$

See Fig. 2.1 for a pictorial summary of the hardness and algorithmic results in various regimes and see Fig. 2.3 for a plot comparing ours and known approximation algorithms in the case of $p \geq 2 \geq q$.

5.1 Proof overview

As discussed in Section 2.1.2, we consider Nesterov's convex relaxation and generalize the hyperplane rounding scheme using "Hölder duals" of the Gaussian projections, instead of taking the sign. Further as in Krivine's rounding scheme, we apply this rounding to a transformation of the convex relaxation's solution. The nature of this transformation depends on how the rounding procedure changes the correlation between two vectors. Let $u, v \in \mathbb{R}^N$ be two unit vectors with $\langle u, v \rangle = \rho$. Then, for $\mathbf{g} \sim \mathcal{N}(0, I_N)$, $\langle \mathbf{g}, u \rangle$ and $\langle \mathbf{g}, v \rangle$ are ρ-correlated Gaussian random variables. Hyperplane rounding then gives ± 1 valued random variables whose correlation is given by

$$\mathop{\mathbb{E}}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2} [\mathrm{sgn}(\mathbf{g}_1) \cdot \mathrm{sgn}(\mathbf{g}_2)] \;=\; \frac{2}{\pi} \cdot \sin^{-1}(\rho)\,.$$

The transformations φ and ψ (to be applied to the vectors u and v) in Krivine's scheme are then chosen depending on the Taylor series for the sin function, which is the inverse of function computed on the correlation. For the case of Hölder-dual rounding, we prove the following generalization of the above identity

$$\mathop{\mathbb{E}}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2} \left[\operatorname{sgn}(\mathbf{g}_1) \, |\mathbf{g}_1|^{q-1} \cdot \operatorname{sgn}(\mathbf{g}_2) \, |\mathbf{g}_2|^{p^*-1} \right] \; = \; \gamma_q^q \cdot \gamma_{p^*}^{p^*} \cdot \rho \cdot {}_2F_1\left(1 - \frac{q}{2}, 1 - \frac{p^*}{2} \, ; \, \frac{3}{2} \, ; \, \rho^2 \right) \, ,$$

where ${}_2F_1$ denotes a hypergeometric function with the specified parameters. The proof of the above identity combines simple tools from Hermite analysis with known integral representations from the theory of special functions, and may be useful in other applications of the rounding procedure.

Note that in the Grothendieck case, we have $\gamma_{p^*}^{p^*} = \gamma_q^q = \sqrt{2/\pi}$, and the remaining part is simply the $\sin^{-1}$ function. In the Krivine rounding scheme, the transformations φ and ψ are chosen to satisfy $(2/\pi) \cdot \sin^{-1}\left(\langle \varphi(u), \psi(v) \rangle\right) = c \cdot \langle u, v \rangle$, where the constant c then governs the approximation ratio. The transformations $\varphi(u)$ and $\psi(v)$ taken to be of the form $\varphi(u) = \oplus_{i=1}^{\infty} a_i \cdot u^{\otimes i}$ such that

$$\langle \varphi(u), \psi(v) \rangle \; = \; c' \cdot \sin\left(\langle u, v \rangle\right) \qquad \text{and} \qquad \|\varphi(u)\|_2 = \|\psi(v)\| = 1 \, .$$

If f represents (a normalized version of) the function of ρ occurring in the identity above (which is $\sin^{-1}$ for hyperplane rounding), then the approximation ratio is governed by the function h obtained by replacing every Taylor coefficient of f^{-1} by its absolute value. While f^{-1} is simply the sin function (and thus h is the sinh function) in the Grothendieck problem, no closed-form expressions are available for general p and q.

The task of understanding the approximation ratio thus reduces to the analytic task of understanding the *family* of the functions h obtained for different values of p and q. Concretely, the approximation ratio is given by the value $1/(h^{-1}(1) \cdot \gamma_q \, \gamma_{p^*})$. At a high level, we prove bounds on $h^{-1}(1)$ by establishing properties of the Taylor coefficients of the family of functions f^{-1}, i.e., the family given by

$$\left\{ f^{-1} \;\mid\; f(\rho) = \rho \cdot {}_2F_1\left(a_1, b_1 \, ; \, 3/2 \, ; \, \rho^2 \right) \; , \; a_1, b_1 \in [0, 1/2] \right\} \, .$$

While in the cases considered earlier, the functions h are easy to determine in terms of f^{-1} via succinct formulae [Kri77, Haa81, AN04] or can be truncated after the cubic term [NR14], neither of these are true for the family of functions we consider. Hypergeometric functions are a rich and expressive class of functions, capturing many of the special functions appearing in Mathematical Physics and various ensembles of orthogonal polynomials. Due to this expressive power, the set of inverses is not well understood. In particular, while the coefficients of f are monotone in p and q, this is not true for f^{-1}. Moreover, the rates of decay of the coefficients may range from inverse polynomial to super-exponential. We analyze the coefficients of f^{-1} using complex-analytic methods inspired by (but quite different from) the work of Haagerup [Haa81] on bounding the complex Grothendieck constant. The key technical challenge in our work is in *arguing systematically about a family of inverse hypergeometric functions* which we address by developing methods to estimate the values of a family of contour integrals.

While our methods only gives a bound of the form $h^{-1}(1) \geq \sinh^{-1}(1)/(1+\varepsilon_0)$, we believe this is an artifact of the analysis and the true bound should indeed be $h^{-1}(1) \geq \sinh^{-1}(1)$.

5.2 Relation to Factorization Theory

Let X, Y be Banach spaces, and let $A : X \to Y$ be a continuous linear operator. As before, the norm $\|A\|_{X\to Y}$ is defined as

$$\|A\|_{X\to Y} := \sup_{x\in X\setminus\{0\}} \frac{\|Ax\|_Y}{\|x\|_X}.$$

The operator A is said to be factorize through Hilbert space if the factorization constant of A defined as

$$\Phi(A) := \inf_H \inf_{BC=A} \frac{\|C\|_{X\to H}\cdot\|B\|_{H\to Y}}{\|A\|_{X\to Y}}$$

is bounded, where the infimum is taken over all Hilbert spaces H and all operators $B : H \to Y$ and $C : X \to H$. The factorization gap for spaces X and Y is then defined as $\Phi(X,Y) := \sup_A \Phi(A)$ where the supremum runs over all continuous operators $A : X \to Y$.

The theory of factorization of linear operators is a cornerstone of modern functional analysis and has also found many applications outside the field (see [Pis86, AK06] for more information). An application to theoretical computer science was found by Tropp [Tro09] who used the Grothendieck factorization [Gro53] to give an algorithmic version of a celebrated column subset selection result of Bourgain and Tzafriri [BT87].

As an almost immediate consequence of convex programming duality, our new algorithmic results also imply some improved factorization results for ℓ_p^n, ℓ_q^m (a similar observation was already made by Tropp [Tro09] in the special case of ℓ_∞^n, ℓ_1^m and for a slightly different relaxation). We first state some classical factorization results, for which we will use $T_2(X)$ and $C_2(X)$ to respectively denote the Type-2 and Cotype-2 constants of X. We refer the interested reader to Section 3.8 for a more detailed description of factorization theory as well as the relevant functional analysis preliminaries.

The Kwapień-Maurey [Kwa72a, Mau74] theorem states that for any pair of Banach spaces X and Y

$$\Phi(X,Y) \le T_2(X)\cdot C_2(Y).$$

However, Grothendieck's result [Gro53] shows that a much better bound is possible in a case where $T_2(X)$ is unbounded. In particular,

$$\Phi(\ell_\infty^n, \ell_1^m) \le K_G,$$

for all $m, n \in \mathbb{N}$. Pisier [Pis80] showed that if X or Y satisfies the approximation property (which is always satisfied by finite-dimensional spaces), then

$$\Phi(X,Y) \le (2\cdot C_2(X^*)\cdot C_2(Y))^{3/2}.$$

We show that the approximation ratio of Nesterov's relaxation is in fact an upper bound on the factorization gap for the spaces ℓ_p^n and ℓ_q^m. Combined with our upper bound on the integrality gap, we show an improved bound on the factorization constant, i.e., for any $p \geq 2 \geq q$ and $m, n \in \mathbb{N}$, we have that for $X = \ell_p^n$, $Y = \ell_q^m$

$$\Phi(X, Y) \leq \frac{1 + \varepsilon_0}{\sinh^{-1}(1)} \cdot (C_2(X^*) \cdot C_2(Y)) ,$$

where $\varepsilon_0 \leq 0.00863$ as before. This improves on Pisier's bound for all $p \geq 2 \geq q$, and for certain ranges of (p, q) it also improves upon K_G and the bound of Kwapień-Maurey.

5.3 Approximability and Factorizability

Let (X_n) and (Y_m) be sequences of Banach spaces such that X_n is over the vector space $\mathbb{R}^n$ and Y_m is over the vector space $\mathbb{R}^m$. We shall say a pair of sequences $((X_n), (Y_m))$ factorize if $\Phi(X_n, Y_m)$ is bounded by a constant independent of m and n. Similarly, we shall say a pair of families $((X_n), (Y_m))$ are computationally approximable if there exists a polynomial $R(m, n)$, such that for every $m, n \in \mathbb{N}$, there is an algorithm with runtime $R(m, n)$ approximating $\|A\|_{X_n \to Y_m}$ within a constant independent of m and n (given an oracle for computing the norms of vectors and a separation oracle for the unit balls of the norms). We consider the natural question of characterizing the families of norms that are approximable and their connection to factorizability and Cotype.

The pairs (p, q) for which (ℓ_p^n, ℓ_q^m) is known (resp. not known) to factorize, are precisely those pairs (p, q) which are known to be computationally approximable (resp. inapproximable assuming hardness conjectures like $\mathsf{P} \neq \mathsf{NP}$ and ETH). Moreover the Hilbertian case which trivially satisfies factorizability, is also known to be computationally approximable (with approximation factor 1).

It is tempting to ask whether the set of computationally approximable pairs includes the set of factorizable pairs or the pairs for which X_n^*, Y_m have bounded (independent of m, n) Cotype-2 constant. Further yet, is there a connection between the approximation factor and the factorization constant, or approximation factor and Cotype-2 constants (of X_n^* and Y_m)? Our work gives some modest additional evidence towards such conjectures. Such a result would give credibility to the appealing intuitive idea of the approximation factor being dependent on the "distance" to a Hilbert space.

5.4 Notation

For $p \geq 2 \geq q \geq 1$, we will use the following notation: $a := p^* - 1$ and $b := q - 1$. We note that $a, b \in [0, 1]$.

For an $m \times n$ matrix M (or vector, when $n = 1$). For a scalar function f, we define $f[M]$ to be the matrix with entries defined as $(f[M])_{i,j} = f(M_{i,j})$ for $i \in [m], j \in [n]$. For vectors $u, v \in \mathbb{R}^\ell$, we denote by $u \circ v \in \mathbb{R}^\ell$ the entry-wise/Hadamard product of u and v. We denote the concatenation of two vectors u and v by $u \oplus v$.

For a function $f(\tau) = \sum_{k \geq 0} f_k \cdot \tau^k$ defined as a power series, we denote by $\mathrm{abs}\,(f)$, the function $\mathrm{abs}\,(f)\,(\tau) := \sum_{k \geq 0} |f_k| \cdot \tau^k$.

5.5 Analyzing the Approximation Ratio via Rounding

We will show that $\mathsf{CP}(A)$ is a good approximation to $\|A\|_{p\to q}$ by using an appropriate generalization of Krivine's rounding procedure. Before stating the generalized procedure, we shall give a more detailed summary of Krivine's procedure.

5.5.1 Krivine's Rounding Procedure

Krivine's procedure centers around the classical random hyperplane rounding. In this context, we define the random hyperplane rounding procedure on an input pair of matrices $U \in \mathbb{R}^{m\times\ell}$, $V \in \mathbb{R}^{n\times\ell}$ as outputting the vectors $\mathrm{sgn}[U\mathbf{g}]$ and $\mathrm{sgn}[V\mathbf{g}]$ where $\mathbf{g} \in \mathbb{R}^{\ell}$ is a vector with i.i.d. standard Gaussian coordinates ($f[v]$ denotes entry-wise application of a scalar function f to a vector v. We use the same convention for matrices.). The so-called Grothendieck identity states that for vectors $u, v \in \mathbb{R}^{\ell}$,

$$\mathbb{E}\left[\mathrm{sgn}\langle \mathbf{g}, u\rangle \cdot \mathrm{sgn}\langle \mathbf{g}, v\rangle\right] = \frac{\sin^{-1}\langle \widehat{u}, \widehat{v}\rangle}{\pi/2}$$

where $\widehat{u}$ denotes $u/\|u\|_2$. This implies the following equality which we will call the hyperplane rounding identity:

$$\mathbb{E}\left[\mathrm{sgn}[U\mathbf{g}](\mathrm{sgn}[V\mathbf{g}])^T\right] = \frac{\sin^{-1}[\widehat{U}\widehat{V}^T]}{\pi/2}. \tag{5.1}$$

where for a matrix U, we use $\widehat{U}$ to denote the matrix obtained by replacing the rows of U by the corresponding unit (in ℓ_2 norm) vectors. Krivine's main observation is that for any matrices U, V, there exist matrices $\varphi(\widehat{U}), \psi(\widehat{V})$ with unit vectors as rows, such that

$$\varphi(\widehat{U})\,\psi(\widehat{V})^T = \sin[(\pi/2)\cdot c \cdot \widehat{U}\widehat{V}^T]$$

where $c = \sinh^{-1}(1)\cdot 2/\pi$. Taking $\widehat{U}, \widehat{V}$ to be the optimal solution to $\mathsf{CP}(A)$, it follows that

$$\|A\|_{\infty\to 1} \geq \left\langle A, \mathbb{E}\left[\mathrm{sgn}[\varphi(\widehat{U})\,\mathbf{g}]\;(\mathrm{sgn}[\psi(\widehat{V})\,\mathbf{g}])^T\right]\right\rangle = \langle A, c\cdot\widehat{U}\widehat{V}^T\rangle = c\cdot\mathsf{CP}(A).$$

The proof of Krivine's observation follows from simulating the Taylor series of a scalar function using inner products. We will now describe this more concretely.

Observation 5.5.1 (Krivine). *Let $f : [-1,1] \to \mathbb{R}$ be a scalar function satisfying $f(\rho) = \sum_{k\geq 1} f_k\,\rho^k$ for an absolutely convergent series (f_k). Let $\mathrm{abs}(f)(\rho) := \sum_{k\geq 1}|f_k|\,\rho^k$ and further for vectors $u, v \in \mathbb{R}^{\ell}$ of ℓ_2-length at most 1, let*

$$S_L(f,u) := (\mathrm{sgn}(f_1)\sqrt{f_1}\cdot u) \oplus (\mathrm{sgn}(f_2)\sqrt{f_2}\cdot u^{\otimes 2}) \oplus (\mathrm{sgn}(f_3)\sqrt{f_3}\cdot u^{\otimes 3}) \oplus \cdots$$
$$S_R(f,v) := (\sqrt{f_1}\cdot v) \oplus (\sqrt{f_2}\cdot v^{\otimes 2}) \oplus (\sqrt{f_3}\cdot v^{\otimes 3}) \oplus \cdots$$

Then for any $U \in \mathbb{R}^{m\times\ell}$, $V \in \mathbb{R}^{n\times\ell}$, $S_L(f, \sqrt{c_f}\cdot\widehat{U})$ and $S_R(f, \sqrt{c_f}\cdot\widehat{V})$ have ℓ_2-unit vectors as rows, and

$$S_L(f, \sqrt{c_f}\cdot\widehat{U})\; S_R(f, \sqrt{c_f}\cdot\widehat{V})^T = f\,[c_f\cdot\widehat{U}\widehat{V}^T]$$

where $S_L(f, W)$ for a matrix W, is applied to row-wise and $c_f := (\text{abs}\,(f)^{-1})(1)$.

Proof. Using the facts $\langle y^1\otimes y^2\,, y^3\otimes y^4\rangle = \langle y^1\,, y^3\rangle\cdot\langle y^2\,, y^4\rangle$ and $\langle y^1\oplus y^2\,, y^3\oplus y^4\rangle = \langle y^1\,, y^3\rangle + \langle y^2\,, y^4\rangle$, we have

- $\langle S_L(f,u)\,, S_R(f,v)\rangle = f(\langle u\,,v\rangle)$
- $\|S_L(f,u)\|_2 = \sqrt{\text{abs}\,(f)\,(\|u\|_2^2)}$
- $\|S_R(f,v)\|_2 = \sqrt{\text{abs}\,(f)\,(\|v\|_2^2)}$

The claim follows. ∎

Before stating our full rounding procedure, we first discuss a natural generalization of random hyperplane rounding, and much like in Krivine's case this will guide the final procedure.

5.5.2 Generalizing Random Hyperplane Rounding – Hölder Dual Rounding

Fix any convex bodies $B_1 \subset \mathbb{R}^m$ and $B_2 \subset \mathbb{R}^k$. Suppose that we would like a strategy that for given vectors $y \in \mathbb{R}^m$, $x \in \mathbb{R}^n$, outputs $\overline{y} \in B_1$, $\overline{x} \in B_2$ so that $y^T A\, x = \langle A\,, y\, x^T\rangle$ is close to $\langle A\,, \overline{y}\,\overline{x}^T\rangle$ for all A. A natural strategy is to take

$$(\overline{y}, \overline{x}) := \underset{(\widetilde{y},\widetilde{x})\in B_1\times B_2}{\text{argmax}} \left\langle \widetilde{y}\,\widetilde{x}^T\,, y\,x^T\right\rangle = \left(\underset{\widetilde{y}\in B_1}{\text{argmax}}\,\langle\widetilde{y}\,,y\rangle\;\;,\;\;\underset{\widetilde{x}\in B_2}{\text{argmax}}\,\langle\widetilde{x}\,,x\rangle\right)$$

In the special case where B is the unit ℓ_p ball, there is a closed form for an optimal solution to $\max_{\widetilde{x}\in B}\langle\widetilde{x}\,,x\rangle$, given by $\Psi_{p^*}(x)/\|x\|_{p^*}^{p^*-1}$, where $\Psi_{p^*}(x) := \text{sgn}[x]\circ|[x]|^{p^*-1}$. Note that for $p=\infty$, this strategy recovers the random hyperplane rounding procedure. We shall call this procedure, *Gaussian Hölder Dual Rounding* or *Hölder Dual Rounding* for short.

Just like earlier, we will first understand the effect of *Hölder Dual Rounding* on a solution pair U, V. For $\rho \in [-1,1]$, let $\mathbf{g}_1 \sim_\rho \mathbf{g}_2$ denote ρ-correlated standard Gaussians, i.e., $\mathbf{g}_1 = \rho\,\mathbf{g}_2 + \sqrt{1-\rho^2}\,\mathbf{g}_3$ where $(\mathbf{g}_2, \mathbf{g}_3) \sim \mathcal{N}(0, \mathrm{I}_2)$, and let

$$\widetilde{f}_{a,b}(\rho) := \mathop{\mathbb{E}}_{\mathbf{g}_1\sim_\rho\mathbf{g}_2}\left[\text{sgn}(\mathbf{g}_1)|\mathbf{g}_1|^b\,\text{sgn}(\mathbf{g}_2)|\mathbf{g}_1|^a\right]$$

We will work towards a better understanding of $\widetilde{f}_{a,b}(\cdot)$ in later sections. For now note that we have for vectors $u, v \in \mathbb{R}^\ell$,

$$\mathbb{E}\left[\text{sgn}\langle\mathbf{g}\,,u\rangle\,|\langle\mathbf{g}\,,u\rangle|^b\cdot\text{sgn}\langle\mathbf{g}\,,v\rangle\,|\langle\mathbf{g}\,,v\rangle|^a\right] = \|u\|_2^b\cdot\|v\|_2^a\cdot\widetilde{f}_{a,b}(\langle\widehat{u}\,,\widehat{v}\rangle)\,.$$

Thus given matrices U, V, we obtain the following generalization of the hyperplane rounding identity for *Hölder Dual Rounding* :

$$\mathbb{E}\left[\Psi_q([U\mathbf{g}])\,\Psi_{p^*}([V\mathbf{g}])^T\right] = D_{(\|u^i\|_2^b)_{i\in[m]}} \cdot \widetilde{f}_{a,b}([\widehat{U}\widehat{V}^T]) \cdot D_{(\|v^j\|_2^a)_{j\in[n]}}\,. \tag{5.2}$$

5.5.3 Generalized Krivine Transformation and the Full Rounding Procedure

We are finally ready to state the generalized version of Krivine's algorithm. At a high level the algorithm simply applies *Hölder Dual Rounding* to a transformed version of the optimal convex program solution pair U, V. Analogous to Krivine's algorithm, the transformation is a type of "inverse" of Eq. (5.2).

(Inversion 1) Let (U, V) be the optimal solution to $\mathsf{CP}(A)$, and let $(u^i)_{i\in[m]}$ and $(v^j)_{j\in[n]}$ respectively denote the rows of U and V.

(Inversion 2) Let $c_{a,b} := \left(\text{abs}\left(\widetilde{f}^{-1}_{a,b}\right)\right)^{-1}(1)$ and let

$$\varphi(U) := D_{(\|u^i\|_2^{1/b})_{i\in[m]}}\ S_L(\widetilde{f}^{-1}_{a,b}, \sqrt{c_{a,b}}\cdot\widehat{U})\,,$$
$$\psi(V) := D_{(\|v^j\|_2^{1/a})_{j\in[n]}}\ S_R(\widetilde{f}^{-1}_{a,b}, \sqrt{c_{a,b}}\cdot\widehat{V})\,.$$

(Hölder-Dual 1) Let $\mathbf{g} \sim \mathcal{N}(0, \mathrm{I})$ be an infinite dimensional i.i.d. Gaussian vector.

(Hölder-Dual 2) Return $y := \Psi_q(\varphi(U)\,\mathbf{g})/\|\varphi(U)\,\mathbf{g}\|_q^b$ and $x := \Psi_{p^*}(\psi(V)\,\mathbf{g})/\|\psi(V)\,\mathbf{g}\|_{p^*}^a$.

Remark 5.5.2. *Note that $\|\Psi_r(\overline{x})\|_{r^*} = \|\overline{x}\|_r^{r-1}$ and so the returned solution pair always lie on the unit ℓ_{q^*} and ℓ_p spheres respectively.*

Remark 5.5.3. *Like in [AN04] the procedure above can be made algorithmic by observing that there always exist $\varphi'(U) \in \mathbb{R}^{m\times(m+n)}$ and $\psi'(V) \in \mathbb{R}^{m\times(m+n)}$, whose rows have the exact same lengths and pairwise inner products as those of $\varphi(U)$ and $\psi(V)$ above. Moreover they can be computed without explicitly computing $\varphi(U)$ and $\psi(V)$ by obtaining the Gram decomposition of*

$$M := \begin{bmatrix} \text{abs}\left(\widetilde{f}^{-1}_{a,b}\right)[c_{a,b}\cdot\widehat{V}\widehat{V}^T] & \widetilde{f}^{-1}_{a,b}([c_{a,b}\cdot\widehat{U}\widehat{V}^T]) \\ \widetilde{f}^{-1}_{a,b}([c_{a,b}\cdot\widehat{V}\widehat{U}^T]) & \text{abs}\left(\widetilde{f}^{-1}_{a,b}\right)[c_{a,b}\cdot\widehat{V}\widehat{V}^T] \end{bmatrix},$$

and normalizing the rows of the decomposition according to the definition of $\varphi(\cdot)$ and $\psi(\cdot)$ above. The entries of M can be computed in polynomial time with exponentially (in m and n) good accuracy by implementing the Taylor series of $\widetilde{f}^{-1}_{a,b}$ upto $\mathrm{poly}(m,n)$ terms (Taylor series inversion can be done upto k terms in time $\mathrm{poly}(k)$).

Remark 5.5.4. *Note that the 2-norm of the i-th row (resp. j-th row) of $\varphi(U)$ (resp. $\psi(V)$) is $\|u^i\|_2^{1/b}$ (resp. $\|v^j\|_2^{1/a}$).*

We commence the analysis by defining some convenient normalized functions and we will also show that $c_{a,b}$ above is well-defined.

5.5.4 Auxiliary Functions

Let $f_{p,q}(\rho) := \widetilde{f}_{p,q}(\rho)/(\gamma_{p^*}^{p^*}\gamma_q^q)$, $\widetilde{h}_{a,b} := \text{abs}\left(\widetilde{f}_{a,b}^{-1}\right)$, and $h_{a,b} := \text{abs}\left(f_{a,b}^{-1}\right)$. Also note that $h_{a,b}^{-1}(\rho) = \widetilde{h}_{a,b}^{-1}(\rho)/(\gamma_{p^*}^{p^*}\gamma_q^q)$.

Well Definedness. By Lemma 5.7.7, $f_{a,b}^{-1}(\rho)$ and $h_{a,b}(\rho)$ are well defined for $\rho \in [-1,1]$. By (M1) in Corollary 5.6.19, $\widetilde{f}_1^{-1} = 1$ and hence $h_{a,b}(1) \geq 1$ and $h_{a,b}(-1) \leq -1$. Combining this with the fact that $h_{a,b}(\rho)$ is continuous and strictly increasing on $[-1,1]$, implies that $h_{a,b}^{-1}(x)$ is well defined on $[-1,1]$.

We can now proceed with the analysis.

5.5.5 $1/(h_{p,q}^{-1}(1)\cdot\gamma_{p^*}\gamma_q)$ Bound on Approximation Factor

For any vector random variable $\mathbf{X}$ in a universe Ω, and scalar valued functions $f_1 : \Omega \to \mathbb{R}$ and $f_2 : \Omega \to (0,\infty)$. Let $\lambda = \mathbb{E}[f_1(\mathbf{X})]/\mathbb{E}[f_2(\mathbf{X})]$. Now we have

$$\begin{aligned} &\max_{x\in\Omega} f_1(x) - \lambda\cdot f_2(x) \geq \mathbb{E}\left[f_1(\mathbf{X}) - \lambda\cdot f_2(\mathbf{X})\right] = 0 \\ \Rightarrow\quad &\max_{x\in\Omega} f_1(x)/f_2(x) \geq \lambda = \mathbb{E}\left[f_1(\mathbf{X})\right]/\mathbb{E}\left[f_2(\mathbf{X})\right]. \end{aligned}$$

Thus we have

$$\|A\|_{p\to q} \;\geq\; \frac{\mathbb{E}[\langle A\,,\Psi_q(\varphi(U)\,\mathbf{g})\;\Psi_{p^*}(\psi(V)\,\mathbf{g})^T\rangle]}{\mathbb{E}[\|\Psi_q(\varphi(U)\,\mathbf{g})\|_{q^*}\cdot\|\Psi_{p^*}(\psi(V)\,\mathbf{g})\|_p]} \;=\; \frac{\langle A\,,\mathbb{E}[\Psi_q(\varphi(U)\,\mathbf{g})\;\Psi_{p^*}(\psi(V)\,\mathbf{g})^T]\rangle}{\mathbb{E}[\|\Psi_q(\varphi(U)\,\mathbf{g})\|_{q^*}\cdot\|\Psi_{p^*}(\psi(V)\,\mathbf{g})\|_p]},$$

which allows us to consider the numerator and denominator separately. We begin by proving the equality that the above algorithm was designed to satisfy:

Lemma 5.5.5. $\mathbb{E}[\Psi_q(\varphi(U)\,\mathbf{g})\;\Psi_{p^*}(\psi(V)\,\mathbf{g})^T] \;=\; c_{a,b}\cdot(\widetilde{U}\widetilde{V}^T)$

Proof.

$$\begin{aligned} &\mathbb{E}\left[\Psi_q(\varphi(U)\,\mathbf{g})\;\Psi_{p^*}(\psi(V)\,\mathbf{g})^T\right] \\ &= D_{(\|u^i\|_2)_{i\in[m]}}\cdot\widetilde{f}_{a,b}([S_L(\widetilde{f}_{a,b}^{-1},\sqrt{c_{a,b}}\cdot\widehat{U})\cdot S_R(\widetilde{f}_{a,b}^{-1},\sqrt{c_{a,b}}\cdot\widehat{V})^T])\cdot D_{(\|v^j\|_2)_{j\in[n]}} \\ &\text{(by Eq. (5.2) and Remark 5.5.4)} \\ &= D_{(\|u^i\|_2)_{i\in[m]}}\cdot\widetilde{f}_{a,b}([\widetilde{f}_{a,b}^{-1}([c_{a,b}\cdot\widehat{U}\widehat{V}^T])])\cdot D_{(\|v^j\|_2)_{j\in[n]}} \\ &\text{(by Observation 5.5.1)} \\ &= D_{(\|u^i\|_2)_{i\in[m]}}\cdot c_{a,b}\cdot\widehat{U}\widehat{V}^T\cdot D_{(\|v^j\|_2)_{j\in[n]}} \\ &= c_{a,b}\cdot UV^T \end{aligned}$$

■

It remains to upper bound the denominator which we do using a straightforward convexity argument.

Lemma 5.5.6. $\mathbb{E}[\|\varphi(U)\,\mathbf{g}\|_q^b \cdot \|\psi(V)\,\mathbf{g}\|_{p^*}^a] \le \gamma_{p^*}^a\,\gamma_q^b$.

Proof.

$$
\begin{aligned}
& \mathbb{E}\left[\|\varphi(U)\,\mathbf{g}\|_q^b \cdot \|\psi(V)\,\mathbf{g}\|_{p^*}^a\right] \\
\le\ & \mathbb{E}\left[\|\varphi(U)\,\mathbf{g}\|_q^{q^*b}\right]^{1/q^*} \cdot \mathbb{E}\left[\|\psi(V)\,\mathbf{g}\|_{p^*}^{pa}\right]^{1/p} && \left(\frac{1}{p} + \frac{1}{q^*} \le 1\right) \\
=\ & \mathbb{E}\left[\|\varphi(U)\,\mathbf{g}\|_q^{q}\right]^{1/q^*} \cdot \mathbb{E}\left[\|\psi(V)\,\mathbf{g}\|_{p^*}^{p^*}\right]^{1/p} \\
=\ & \left[\sum_{i\in[m]} \mathbb{E}\left[|\mathcal{N}(0, \|u^i\|_2^{1/b})|^q\right]\right]^{1/q^*} \cdot \left[\sum_{j\in[n]} \mathbb{E}\left[|\mathcal{N}(0, \|v^j\|_2^{1/a})|^{p^*}\right]\right]^{1/p} && \text{(By Remark 5.5.4)} \\
=\ & \left[\sum_{i\in[m]} \|u^i\|_2^{q/b}\right]^{1/q^*} \cdot \left[\sum_{j\in[n]} \|v^j\|_2^{p^*/a}\right]^{1/p} \cdot \gamma_q^{q/q^*}\,\gamma_{p^*}^{p^*/p} \\
=\ & \left[\sum_{i\in[m]} \|u^i\|_2^{q^*}\right]^{1/q^*} \cdot \left[\sum_{j\in[n]} \|v^j\|_2^{p}\right]^{1/p} \cdot \gamma_q^{b}\,\gamma_{p^*}^{a} \\
=\ & \gamma_q^b\,\gamma_{p^*}^a && \text{(feasibility of } U, V) \quad \blacksquare
\end{aligned}
$$

We are now ready to prove our approximation guarantee.

Lemma 5.5.7. *Consider any* $1 \le q \le 2 \le p \le \infty$. *Then,*

$$\frac{\mathsf{CP}(A)}{\|A\|_{p\to q}} \ \le\ 1/(\gamma_{p^*}\,\gamma_q \cdot h_{a,b}^{-1}(1))$$

Proof.

$$
\begin{aligned}
\|A\|_{p\to q} \ &\ge\ \frac{\langle A\,, \mathbb{E}[\Psi_q(\varphi(U)\,\mathbf{g})\ \Psi_{p^*}(\psi(V)\,\mathbf{g})^T]\rangle}{\mathbb{E}[\|\Psi_q(\varphi(U)\,\mathbf{g})\|_{q^*} \cdot \|\Psi_{p^*}(\psi(V)\,\mathbf{g})\|_p]} \\
&=\ \frac{\langle A\,, \mathbb{E}[\Psi_q(\varphi(U)\,\mathbf{g})\ \Psi_{p^*}(\psi(V)\,\mathbf{g})^T]\rangle}{\mathbb{E}[\|\varphi(U)\,\mathbf{g}\|_q^b \cdot \|\psi(V)\,\mathbf{g}\|_{p^*}^a]} && \text{(by Remark 5.5.2)} \\
&=\ \frac{c_{a,b} \cdot \langle A\,, UV^T\rangle}{\mathbb{E}[\|\varphi(U)\,\mathbf{g}\|_q^b \cdot \|\psi(V)\,\mathbf{g}\|_{p^*}^a]} && \text{(by Lemma 5.5.5)} \\
&=\ \frac{c_{a,b} \cdot \mathsf{CP}(A)}{\mathbb{E}[\|\varphi(U)\,\mathbf{g}\|_q^b \cdot \|\psi(V)\,\mathbf{g}\|_{p^*}^a]} && \text{(by optimality of } U, V) \\
&\ge\ \frac{c_{a,b} \cdot \mathsf{CP}(A)}{\gamma_{p^*}^a\,\gamma_q^b} && \text{(by Lemma 5.5.6)}
\end{aligned}
$$

$$= \frac{\widetilde{h}_{a,b}^{-1}(1) \cdot \mathsf{CP}(A)}{\gamma_{p^*}^{a}\, \gamma_q^{b}}$$

$$= h_{a,b}^{-1}(1) \cdot \gamma_{p^*}\, \gamma_q \cdot \mathsf{CP}(A) \qquad \blacksquare$$

We next begin the primary technical undertaking of this chapter, namely proving upper bounds on $h_{p,q}^{-1}(1)$.

5.6 Hypergeometric Representation of $f_{a,b}(x)$

In this section, we show that $f_{a,b}(\rho)$ can be represented using the Gaussian hypergeometric function ${}_2F_1$. The result of this section can be thought of as a generalization of the so-called Grothendieck identity for hyperplane rounding which simply states that

$$f_{0,0}(\rho) \;=\; \frac{\pi}{2} \cdot \mathop{\mathbb{E}}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2} [\operatorname{sgn}(\mathbf{g}_1) \operatorname{sgn}(\mathbf{g}_2)] \;=\; \sin^{-1}(\rho)$$

We believe the result of this section and its proof technique to be of independent interest in analyzing generalizations of hyperplane rounding to convex bodies other than the hypercube.

Recall that $\widetilde{f}_{a,b}(\rho)$ is defined as follows:

$$\mathop{\mathbb{E}}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2} \left[\operatorname{sgn}(\mathbf{g}_1) |\mathbf{g}_1|^a \operatorname{sgn}(\mathbf{g}_2) |\mathbf{g}_1|^b \right]$$

where $a = p^* - 1$ and $b = q - 1$. Our starting point is the simple observation that the above expectation can be viewed as the noise correlation (under the Gaussian measure) of the functions $\widetilde{f}^{(a)}(\tau) := \operatorname{sgn} \tau \cdot |\tau|^a$ and $\widetilde{f}^{(b)}(\tau) := \operatorname{sgn} \tau \cdot |\tau|^b$. Elementary Hermite analysis then implies that it suffices to understand the Hermite coefficients of $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$ individually, in order to understand the Taylor coefficients of $f_{a,b}$. To understand the Hermite coefficients of $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$ individually, we use a generating function approach. More specifically, we derive an integral representation for the generating function of the (appropriately normalized) Hermite coefficients which fortunately turns out to be closely related to a well studied special function called the parabolic cylinder function.

Before proceeding, we require some preliminaries.

5.6.1 Hermite Analysis Preliminaries

Let γ denote the standard Gaussian probability distribution. For this section (and only for this section), the (Gaussian) inner product for functions $f, h \in (\mathbb{R}, \gamma) \to \mathbb{R}$ is defined as

$$\langle f\,, h \rangle := \int_{\mathbb{R}} f(\tau) \cdot h(\tau)\; d\gamma(\tau) = \mathop{\mathbb{E}}_{\tau \sim \mathcal{N}(0,1)} [f(\tau) \cdot h(\tau)]\;.$$

Under this inner product there is a complete set of orthonormal polynomials $(H_k)_{k \in \mathbb{N}}$ defined below.

Definition 5.6.1. *For a natural number k, then the k-th* Hermite *polynomial* $H_k : \mathbb{R} \to \mathbb{R}$

$$H_k(\tau) = \frac{1}{\sqrt{k!}} \cdot (-1)^k \cdot \mathrm{e}^{\tau^2/2} \cdot \frac{d^k}{d\tau^k} \mathrm{e}^{-\tau^2/2} .$$

Any function f satisfying $\int_{\mathbb{R}} |f(\tau)|^2 \, d\gamma(\tau) < \infty$ has a Hermite expansion given by $f = \sum_{k \geq 0} \widehat{f}_k \cdot H_k$ where $\widehat{f}_k = \langle f \, , H_k \rangle$.

We have

Fact 5.6.2. *$H_k(\tau)$ is an even (resp. odd) function when k is even (resp. odd).*

We also have the Plancherel Identity (as Hermite polynomials form an orthonormal basis):

Fact 5.6.3. *For two real valued functions f and h with Hermite coefficients $\widehat{f}_k$ and $\widehat{h}_k$, respectively, we have:*

$$\langle f \, , h \rangle = \sum_{k \geq 0} \widehat{f}_k \cdot \widehat{h}_k \, .$$

The generating function of appropriately normalized Hermite polynomials satisfies the following identity:

$$e^{\tau\lambda - \lambda^2/2} = \sum_{k \geq 0} H_k(\tau) \cdot \frac{\lambda^k}{\sqrt{k!}} \, . \tag{5.3}$$

Similar to the noise operator in Fourier analysis, we define the corresponding noise operator T_ρ for Hermite analysis:

Definition 5.6.4. *For $\rho \in [-1, 1]$ and a real valued function f, we define the function $T_\rho f$ as:*

$$(T_\rho f)(\tau) = \int_{\mathbb{R}} f\left(\rho \cdot \tau + \sqrt{1 - \rho^2} \cdot \theta\right) \, d\gamma(\theta) = \mathop{\mathbb{E}}_{\tau' \sim_\rho \tau} \left[f(\tau') \right] \, .$$

Again similar to the case of Fourier analysis, the Hermite coefficients admit the following identity:

Fact 5.6.5. $\widehat{(T_\rho f)}_k = \rho^k \cdot \widehat{f}_k$.

We recall that the $\widetilde{f}_{a,b}(\rho) = \mathbb{E}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2} [\widetilde{f}^{(a)}(\mathbf{g}_1) \cdot \widetilde{f}^{(b)}(\mathbf{g}_2))]$, where $\widetilde{f}^{(c)}(\tau) := \mathrm{sgn}(\tau) \cdot |\tau|^c$ for $c \in \{a, b\}$. As mentioned at the start of the section, we now note that $f_{a,b}(\rho)$ is the noise correlation of $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$. Thus we can relate the Taylor coefficients of $f_{a,b}(\rho)$, to the Hermite coefficients of $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$.

Claim 5.6.6 (Coefficients of $\widetilde{f}_{a,b}(\rho)$)**.** *For $\rho \in [-1, 1]$, we have:*

$$\widetilde{f}_{a,b}(\rho) = \sum_{k \geq 0} \rho^{2k+1} \cdot \widehat{f}^{(a)}_{2k+1} \cdot \widehat{f}^{(b)}_{2k+1} \, ,$$

where $\widehat{f}_i^{(a)}$ and $\widehat{f}_j^{(b)}$ are the i-th and j-th Hermite coefficients of $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$, respectively. Moreover, $\widehat{f}^{(a)}_{2k} = \widehat{f}^{(b)}_{2k} = 0$ for $k \geq 0$.

Proof. We observe that both $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$ are odd functions and hence Fact 5.6.2 implies that $\widehat{f}^{(a)}_{2k} = \widehat{f}^{(b)}_{2k} = 0$ for all $k \geq 0$ – as $\widetilde{f}^{(a)}(\tau) \cdot H_{2k}(\tau)$ is an odd function of τ.

$$\begin{aligned}
\widetilde{f}_{a,b}(\rho) &= \mathop{\mathbb{E}}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2}\left[\widetilde{f}^{(a)}(\mathbf{g}_1) \cdot \widetilde{f}^{(b)}(\mathbf{g}_2))\right] \\
&= \mathop{\mathbb{E}}_{\mathbf{g}_1}\left[\widetilde{f}^{(a)}(\mathbf{g}_1) \cdot T_\rho \widetilde{f}^{(b)}(\mathbf{g}_1)\right] && \text{(Definition 5.6.4)} \\
&= \langle \widetilde{f}^{(a)}, T_\rho \widetilde{f}^{(b)} \rangle \\
&= \sum_{k \geq 0} \widehat{f}^{(a)}_k \cdot (\widehat{T_\rho \widetilde{f}^{(b)}})_k && \text{(Fact 5.6.3)} \\
&= \sum_{k \geq 0} \widehat{f}^{(a)}_{2k+1} \cdot (\widehat{T_\rho \widetilde{f}^{(b)}})_{2k+1} \\
&= \sum_{k \geq 0} \rho^{2k+1} \cdot \widehat{f}^{(a)}_{2k+1} \cdot \widehat{f}^{(b)}_{2k+1} && \text{(Fact 5.6.5)}\,.
\end{aligned}$$

■

5.6.2 Hermite Coefficients of $\widetilde{f}^{(a)}$ and $\widetilde{f}^{(b)}$ via Parabolic Cylinder Functions

In this subsection, we use the generating function of Hermite polynomials to to obtain an integral representation for the generating function of the ($\sqrt{k!}$ normalized) odd Hermite coefficients of $\widetilde{f}^{(a)}$ (and similarly of $\widetilde{f}^{(b)}$) is closely related to a special function called the parabolic cylinder function. We then use known facts about the relation between parabolic cylinder functions and confluent hypergeometric functions, to show that the Hermite coefficients of $\widetilde{f}^{(c)}$ can be obtained from the Taylor coefficients of a confluent hypergeometric function.

Before we state and prove the main results of this subsection we need some preliminaries:

Gamma, Hypergeometric and Parabolic Cylinder Function Preliminaries

For a natural number k and a real number τ, we denote the rising factorial as $(\tau)_k := \tau \cdot (\tau+1) \cdot \dots \cdot (\tau+k-1)$. We now define the following fairly general classes of functions and we later use them we obtain a Taylor series representation of $\widetilde{f}_{a,b}(\tau)$.

Definition 5.6.7. *The confluent hypergeometric function with parameters α, β, and λ as:*

$$ {}_1F_1(\alpha\,;\beta\,;\lambda) := \sum_k \frac{(\alpha)_k}{(\beta)_k} \cdot \frac{\lambda^k}{k!}\,. $$

The (Gaussian) hypergeometric function is defined as follows:

Definition 5.6.8. *The hypergeometric function with parameters w, α, β and λ as:*

$$ {}_2F_1(w, \alpha\,;\beta\,;\lambda) := \sum_k \frac{(w)_k \cdot (\alpha)_k}{(\beta)_k} \cdot \frac{\lambda^k}{k!}\,. $$

Next we define the Γ function:

Definition 5.6.9. *For a real number τ, we define:*

$$\Gamma(\tau) := \int_0^\infty t^{\tau-1} \cdot \mathrm{e}^{-t}\, dt\,.$$

The Γ function has the following property:

Fact 5.6.10 (Duplication Formula)**.**

$$\frac{\Gamma(2\tau)}{\Gamma(\tau)} = \frac{\Gamma(\tau+1/2)}{2^{1-2\tau}\sqrt{\pi}}$$

We also note the relationship between Γ and γ_r:

Fact 5.6.11. *For $r \in [0, \infty)$,*

$$\gamma_r^r \;:=\; \mathop{\mathbb{E}}_{\mathbf{g}\sim\mathcal{N}(0,1)} [|\mathbf{g}|^r] \;=\; \frac{2^{r/2}}{\sqrt{\pi}} \cdot \Gamma\left(\frac{1+r}{2}\right)\,.$$

Proof.

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{g}\sim\mathcal{N}(0,1)} \left[|\mathbf{g}|^r\right] &= \frac{\sqrt{2}}{\sqrt{\pi}} \cdot \int_0^\infty |\mathbf{g}|^r \cdot \mathrm{e}^{-\mathbf{g}^2/2}\, d\mathbf{g} \\ &= \sqrt{\frac{2}{\pi}} \cdot 2^{(r-1)/2} \cdot \int_0^\infty \left|\frac{\mathbf{g}^2}{2}\right|^{(r-1)/2} \cdot \mathrm{e}^{-\mathbf{g}^2/2} \cdot \mathbf{g}\, d\mathbf{g} \\ &= \frac{2^{r/2}}{\sqrt{\pi}} \cdot \Gamma\left(\frac{1+r}{2}\right)\end{aligned}$$

∎

Next, we record some facts about parabolic cylinder functions:

Fact 5.6.12 (12.5.1 of [Loz03])**.** *Let U be the function defined as*

$$U(\alpha, \lambda) := \frac{\mathrm{e}^{\lambda^2/4}}{\Gamma\left(\frac{1}{2}+\alpha\right)} \int_0^\infty t^{\alpha-1/2} \cdot \mathrm{e}^{-(t+\lambda)^2/2}\, dt\,,$$

for all α such that $\mathbb{R}(\alpha) > -\frac{1}{2}$. The function $U(\alpha, \pm\lambda)$ is a parabolic cylinder function and is a standard solution to the differential equation: $\frac{d^2w}{d\lambda^2} - \left(\frac{\lambda^2}{4} + \alpha\right) w = 0$.

Next we quote the confluent hypergeometric representation of the parabolic cylinder function U defined above:

Fact 5.6.13 (12.4.1, 12.2.6, 12.2.7, 12.7.12, and 12.7.13 of [Loz03]).

$$U(\alpha,\lambda) = \frac{\sqrt{\pi}}{2^{\alpha/2+1/4}\cdot\Gamma\left(\frac{3}{4}+\frac{\alpha}{2}\right)}\cdot e^{\lambda^2/4}\cdot {}_1F_1\left(-\frac{1}{2}\alpha+\frac{1}{4};\frac{1}{2};-\frac{\lambda^2}{2}\right)$$
$$-\frac{\sqrt{\pi}}{2^{\alpha/2-1/4}\cdot\Gamma\left(\frac{1}{4}+\frac{\alpha}{2}\right)}\cdot\lambda\cdot e^{\lambda^2/4}\cdot {}_1F_1\left(-\frac{\alpha}{2}+\frac{3}{4};\frac{3}{2};-\frac{\lambda^2}{2}\right)$$

Combining the previous two facts, we get the following:

Corollary 5.6.14. *For all real $\alpha > -\frac{1}{2}$, we have:*

$$\int_0^\infty t^{\alpha-1/2}\cdot e^{-(t+\lambda)^2/2}\,dt \;=\; \frac{\sqrt{\pi}\cdot\Gamma\left(\frac{1}{2}+\alpha\right)}{2^{\alpha/2+1/4}\cdot\Gamma\left(\frac{3}{4}+\frac{\alpha}{2}\right)}\cdot {}_1F_1\left(-\frac{\alpha}{2}+\frac{1}{4};\frac{1}{2};-\frac{\lambda^2}{2}\right)$$
$$-\frac{\sqrt{\pi}\cdot\Gamma\left(\frac{1}{2}+\alpha\right)}{2^{\alpha/2-1/4}\cdot\Gamma\left(\frac{1}{4}+\frac{\alpha}{2}\right)}\cdot\lambda\cdot {}_1F_1\left(-\frac{\alpha}{2}+\frac{3}{4};\frac{3}{2};-\frac{\lambda^2}{2}\right).$$

Generating Function of Hermite Coefficients and its Confluent Hypergeometric Representation

Using the generating function of (appropriately normalized) Hermite polynomials, we derive an integral representation for the generating function of the (appropriately normalized) Hermite coefficients of $\widetilde{f}^{(a)}$ (and similarly $\widetilde{f}^{(b)}$):

Lemma 5.6.15. *For $c \in \{a,b\}$, let $\widehat{f}_k^{(c)}$ denote the k-th Hermite coefficient of $\widetilde{f}^{(c)}(\tau) := \mathrm{sgn}(\tau)\cdot|\tau|^c$. Then we have the following identity:*

$$\sum_{k\geq 0}\frac{\lambda^{2k+1}}{\sqrt{(2k+1)!}}\cdot\widehat{f}_{2k+1}^{(c)} = \frac{1}{\sqrt{2\pi}}\int_0^\infty \tau^c\cdot\left(e^{-(\tau-\lambda)^2/2}-e^{-(\tau+\lambda)^2/2}\right)d\tau\,.$$

Proof. We observe that for, $\widetilde{f}^{(c)}$ is an odd function and hence Fact 5.6.2 implies that $\widetilde{f}^{(c)}(\tau)\cdot H_{2k}(\tau)$ is an odd function and $\widetilde{f}^{(c)}(\tau)\cdot H_{2k+1}(\tau)$ is an even function. This implies for any $k\geq 0$, that $\widehat{f}_{2k}^{(c)} = 0$ and

$$\widehat{f}_{2k+1}^{(c)} = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^\infty \mathrm{sgn}(\tau)\cdot\tau^c\cdot H_{2k+1}(\tau)\cdot e^{-\tau^2/2}\,d\tau = \sqrt{\frac{2}{\pi}}\int_0^\infty \tau^c\cdot H_{2k+1}(\tau)\cdot e^{-\tau^2/2}\,d\tau\,.$$

Thus we have

$$\sum_{k\geq 0}\frac{\lambda^{2k+1}}{\sqrt{(2k+1)!}}\cdot\widehat{f}_{2k+1}^{(c)}$$
$$= \sqrt{\frac{2}{\pi}}\cdot\sum_{k\geq 0}\int_0^\infty \tau^c\cdot e^{-\tau^2/2}\cdot H_{2k+1}(\tau)\cdot\frac{\lambda^{2k+1}}{\sqrt{(2k+1)!}}\,d\tau$$

$$= \sqrt{\frac{2}{\pi}} \cdot \int_0^\infty \tau^c \cdot e^{-\tau^2/2} \sum_{k\geq 0} H_{2k+1}(\tau) \cdot \frac{\lambda^{2k+1}}{\sqrt{(2k+1)!}}\, d\tau \qquad \text{(see below)}$$

$$= \frac{1}{\sqrt{2\pi}} \cdot \int_0^\infty \tau^c \cdot e^{-\tau^2/2} \cdot \left(e^{\tau\lambda - \lambda^2/2} - e^{-\tau\lambda - \lambda^2/2} \right) d\tau \qquad \text{(by Eq. (5.3))}$$

$$= \frac{1}{\sqrt{2\pi}} \cdot \int_0^\infty \tau^c \cdot \left(e^{-(\tau-\lambda)^2/2} - e^{-(\tau+\lambda)^2/2} \right) d\tau$$

where the exchange of summation and integral in the second equality follows by Fubini's theorem. We include this routine verification for the sake of completeness. As a consequence of Fubini's theorem, if $(f_k : \mathbb{R} \to \mathbb{R})_k$ is a sequence of functions such that $\sum_{k\geq 0} \int_0^\infty |f_k| < \infty$, then $\sum_{k\geq 0} \int_0^\infty f_k = \int_0^\infty \sum_{k\geq 0} f_k$. Now for any fixed k, we have

$$\int_0^\infty \tau^c \cdot |H_k(x)|\, d\gamma(\tau) \;\leq\; \left(\int_0^\infty \tau^{2c}\, d\gamma(\tau) \right)^{1/2} \cdot \left(\int_0^\infty |H_k(x)|^2\, d\gamma(\tau) \right)^{1/2} \;\leq\; \gamma_{2c}^c < \infty\,.$$

Setting $f_k(\tau) := \tau^c \cdot e^{-\tau^2/2} \cdot H_{2k+1}(\tau) \cdot \lambda^{2k+1} / \sqrt{(2k+1)!}$, we get that $\sum_{k\geq 0} \int_0^\infty |f_k| < \infty$. This completes the proof. ∎

Finally using known results about parabolic cylinder functions, we are able to relate the aforementioned integral representation to a confluent hypergeometric function (whose Taylor coefficients are known).

Lemma 5.6.16. *For $\lambda \in [-1,1]$ and real valued $c > -1$, we have*

$$\frac{1}{\sqrt{2\pi}} \int_0^\infty \tau^c \left(e^{-(\tau-\lambda)^2/2} - e^{-(\tau+\lambda)^2/2} \right) d\tau \;=\; \gamma_{c+1}^{c+1} \cdot \lambda \cdot {}_1F_1\left(\frac{1-c}{2}; \frac{3}{2}; -\frac{\lambda^2}{2} \right)$$

Proof. We prove this by using the Corollary 5.6.14 with $a = c + \frac{1}{2}$. We note that $\alpha > -\frac{1}{2}$ and ${}_1F_1\left(\cdot, \cdot, -\lambda^2/2\right)$ is an even function of λ. So combining the two, we get:

$$\frac{1}{\sqrt{2\pi}} \int_0^\infty \tau^c \left(e^{-(\tau-\lambda)^2/2} - e^{-(\tau+\lambda)^2/2} \right) d\tau$$

$$= \frac{2}{\sqrt{2\pi}} \cdot \frac{\sqrt{\pi} \cdot \Gamma(c+1)}{2^{c/2} \cdot \Gamma\left(\frac{c+1}{2}\right)} \cdot \lambda \cdot {}_1F_1\left(-\frac{c}{2} + \frac{1}{2}; \frac{3}{2}; -\frac{1}{2}\lambda^2 \right)$$

$$= 2^{(1-c)/2} \cdot \frac{\Gamma\left(\frac{c+1}{2} + \frac{1}{2}\right)}{2^{-c} \cdot \sqrt{\pi}} \cdot \lambda \cdot {}_1F_1\left(\frac{1-c}{2}; \frac{3}{2}; -\frac{\lambda^2}{2} \right) \qquad \text{(by Fact 5.6.10)}$$

$$= \gamma_{c+1}^{c+1} \cdot \lambda \cdot {}_1F_1\left(\frac{1-c}{2}; \frac{3}{2}; -\frac{\lambda^2}{2} \right) \qquad \text{(by Fact 5.6.11)}$$ ∎

5.6.3 Taylor Coefficients of $\widetilde{f}_{a,b}(x)$ and Hypergeometric Representation

By Claim 5.6.6, we are left with understanding the function whose power series is given by a weighted coefficient-wise product of a certain pair of confluent hypergeometric functions. This turns out to be precisely the Gaussian hypergeometric function, as we will see below.

Observation 5.6.17. *Let $f_k := [\tau^k]\,{}_1F_1(a_1, 3/2, \tau)$ and $h_k := [\tau^k]\,{}_1F_1(b_1, 3/2, \tau)$. Further let $\mu_k := f_k \cdot h_k \cdot (2k+1)!/4^k$. Then for $\rho \in [-1,1]$,*

$$\sum_{k\geq 0} \mu_k \cdot \rho^n \;=\; {}_2F_1(a_1, b_1\,;\,3/2\,;\,\rho)\,.$$

Proof. The claim is equivalent to showing that $\mu_k = (a_1)_k\,(b_1)_k/((3/2)_k\,k!)$. Since we have $f_k = (a_1)_k/((3/2)_k\,k!)$ and $h_k = (b_1)_k/((3/2)_k\,k!)$, it is sufficient to show that $(2k+1)!/4^k = (3/2)_k \cdot k!$. Indeed we have,

$$\begin{aligned}(2k+1)! &= 2^k \cdot k! \cdot 1 \cdot 3 \cdot 5 \cdots (2k+1) \\ &= 4^k \cdot k! \cdot \frac{3}{2} \cdot \frac{5}{2} \cdots \left(\frac{3}{2} + k - 1\right) \\ &= 4^k \cdot k! \cdot (3/2)_k\,.\end{aligned}$$

■

We are finally equipped to put everything together.

Theorem 5.6.18. *For any $a, b \in (-1,\infty)$ and $\rho \in [-1,1]$, we have*

$$f_{a,b}(\rho) \;:=\; \frac{1}{\gamma_{a+1}^{a+1} \cdot \gamma_{b+1}^{b+1}} \cdot \mathop{\mathbb{E}}_{\mathbf{g}_1 \sim_\rho \mathbf{g}_2}\left[\mathrm{sgn}(\mathbf{g}_1)|\mathbf{g}_1|^a\,\mathrm{sgn}(\mathbf{g}_2)|\mathbf{g}_1|^b\right] \;=\; \rho \cdot {}_2F_1\left(\frac{1-a}{2}, \frac{1-b}{2}\,;\,\frac{3}{2}\,;\,\rho^2\right).$$

It follows that the $(2k+1)$-th Taylor coefficient of $f_{a,b}(\rho)$ is

$$\frac{((1-a)/2)_k\,((1-b)/2)_k}{((3/2)_k\,k!)}\,.$$

Proof. The claim follows by combining Claim 5.6.6, Lemmas 5.6.15 and 5.6.16, and Observation 5.6.17. ■

This hypergeometric representation immediately yields some non-trivial coefficient and monotonicity properties:

Corollary 5.6.19. *For any $a, b \in [0,1]$, the function $f_{a,b} : [-1,1] \to \mathbb{R}$ satisfies*

(M1) $[\rho]\,f_{a,b}(\rho) = 1$ *and* $[\rho^3]\,f_{a,b}(\rho) = (1-a)(1-b)/6$.

(M2) All Taylor coefficients are non-negative. Thus $f_{a,b}(\rho)$ is increasing on $[-1,1]$.

(M3) All Taylor coefficients are decreasing in a and in b. Thus for any fixed $\rho \in [-1,1]$, $f_{a,b}(\rho)$ is decreasing in a and in b.

(M4) Note that $f_{a,b}(0) = 0$ and by (M1) and (M2), $f_{a,b}(1) \geq 1$. By continuity, $f_{a,b}([0,1])$ contains $[0,1]$. Combining this with (M3) implies that for any fixed $\rho \in [0,1]$, $f_{a,b}^{-1}(\rho)$ is increasing in a and in b.

5.7 $\sinh^{-1}(1)/(1+\varepsilon_0)$ Bound on $h_{a,b}^{-1}(1)$

In this section we show that $p = \infty, q = 1$ (the Grothendieck case) is roughly the extremal case for the value of $h_{a,b}^{-1}(1)$, i.e., we show that for any $1 \leq q \leq 2 \leq p \leq \infty$, $h_{a,b}^{-1}(1) \geq \sinh^{-1}(1)/(1+\varepsilon_0)$ (recall that $h_{0,0}^{-1}(1) = \sinh^{-1}(1)$). While we were unable to establish as much, we conjecture that $h_{a,b}^{-1}(1) \geq \sinh^{-1}(1)$. Section 5.7.1 details some of the challenges involved in establishing that $\sinh^{-1}(1)$ is the worst case, and presents our approach to establish an approximate bound, which will be formally proved in Section 5.7.2.

5.7.1 Behavior of The Coefficients of $f_{a,b}^{-1}(z)$.

Krivine's upper bound on the real Grothendieck constant, Haagerup's upper bound [Haa81] on the complex Grothendieck constant and the work of Naor and Regev [NR14, BdOFV14] on the optimality of Krivine schemes are all closely related to our work in that each of the aforementioned papers needs to lower bound $(\text{abs}\,(f^{-1}))^{-1}(1)$ for an appropriate odd function f (the work of Briet et al. [BdOFV14] on the rank-constrained Grothendieck problem is also a generalization of Krivine's and Haagerup's work, however they did not derive a closed form upper bound on $(\text{abs}\,(f^{-1}))^{-1}(1)$ in their setting). In Krivine's setting $f = \sin^{-1} x$, implying $(\text{abs}\,(f^{-1}))^{-1} = \sinh^{-1}$ and hence the bound is immediate. In our setting, as well as in [Haa81] and [NR14, BdOFV14], f is given by its Taylor coefficients and is not known to have a closed form. In [NR14], all coefficients of f^{-1} subsequent to the third are negligible and so one doesn't incur much loss by assuming that $\text{abs}\,(f^{-1})\,(\rho) = c_1\rho + c_3\rho^3$. In [Haa81], the coefficient of ρ in $f^{-1}(\rho)$ is 1 and every subsequent coefficient is negative, which implies that $\text{abs}\,(f^{-1})\,(\rho) = 2\rho - f^{-1}(\rho)$. Note that if the odd coefficients of f^{-1} are alternating in sign like in Krivine's setting, then $\text{abs}\,(f^{-1})\,(\rho) = -i \cdot f^{-1}(i\rho)$. These structural properties of the coefficients help their analyses.

In our setting there does not appear to be such a strong relation between $(\text{abs}\,(f^{-1}))$ and f^{-1}. Consider $f(\rho) = f_{a,a}(\rho)$. For certain $a \in (0,1)$, the sign pattern of the coefficients of f^{-1} is unlike that of [Haa81] or $\sin\rho$. In fact empirical results suggest that the odd coefficients of f alternate in sign up to some term $K = K(a)$, and subsequently the coefficients are all non-positive (where $K(a) \to \infty$ as $a \to 0$), i.e., the sign pattern appears to be interpolating between that of $\sin\rho$ and that of $f^{-1}(\rho)$ in the case of Haagerup [Haa81].

Another source of difficulty is that for a fixed a, the coefficients of f^{-1} (with and without magnitude) are not necessarily monotone in k, and moreover for a fixed k, the k-th coefficient of f^{-1} is not necessarily monotone in a.

A key part of our approach is noting that certain milder assumptions on the coefficients are sufficient to show that $\sinh^{-1}(1)$ is the worst case. The proof crucially uses the monotonicity of $f_{a,b}(\rho)$ in a and b. The conditions are as follows:

Let $f_k^{-1} := [\rho^k]\, f_{a,b}^{-1}(\rho)$. Then

(C1) $f_k^{-1} \leq 1/k!$ if $k \pmod 4 \equiv 1$.

(C2) $f_k^{-1} \leq 0$ if $k \pmod 4 \equiv 3$.

To be more precise, we were unable to establish that the above conditions hold for all k (however we conjecture that it is true for all k), and instead use Mathematica to verify it for the fist few coefficients. We additionally show that the coefficients of $f_{a,b}^{-1}$ decay exponentially. Combining this exponential decay with a robust version of the previously advertised claim yields that $h_{a,b}^{-1}(1) \geq \sinh^{-1}(1)/(1+\varepsilon_0)$.

We next proceed to prove the claim that the aforementioned conditions are sufficient to show that $\sinh^{-1}(1)$ is the worst case. We will need the following definition. For an odd positive integer t, let

$$h_{err}(t,\rho) := \sum_{k\geq t} |f_k^{-1}| \cdot \rho^k$$

Lemma 5.7.1. *If t is an odd integer such that (C1) and (C2) are satisfied for all $k < t$, and $\rho = \sinh^{-1}(1 - 2h_{err}(t,\delta))$ for some $\delta \geq \rho$, then $h_{a,b}(\rho) \leq 1$.*

Proof. We have,

$$\begin{aligned}
& h_{a,b}(\rho) \\
&= \sum_{k\geq 1} |f_k^{-1}| \cdot \rho^k \\
&= -f_{a,b}^{-1}(\rho) + \sum_{k\geq 1} \max\{2f_k^{-1}, 0\} \cdot \rho^k \\
&= -f_{a,b}^{-1}(\rho) + \sum_{\substack{1\leq k<t \\ k \bmod 4 \equiv 1}} \max\{2f_k^{-1}, 0\} \cdot \rho^k + \sum_{k\geq t} \max\{2f_k^{-1}, 0\} \cdot \rho^k && \text{(by (C2))} \\
&\leq -f_{a,b}^{-1}(\rho) + \sum_{\substack{1\leq k<t \\ k \bmod 4 \equiv 1}} \max\{2f_k^{-1}, 0\} \cdot \rho^k + 2\, h_{err}(t,\rho) \\
&\leq -f_{a,b}^{-1}(\rho) + \sin(\rho) + \sinh(\rho) + 2\, h_{err}(t,\rho) && \text{(by (C1))} \\
&\leq -f_{a,b}^{-1}(\rho) + \sin(\rho) + 1 + 2(h_{err}(t,\rho) - h_{err}(t,\delta)) && (\rho = \sinh^{-1}(1 - 2h_{err}(t,\delta))) \\
&\leq -f_{a,b}^{-1}(\rho) + \sin(\rho) + 1 && (\rho \leq \delta) \\
&\leq -f_{0,0}^{-1}(\rho) + \sin(\rho) + 1 && \text{(Corollary 5.6.19 : (M4))} \\
&= 1 && (f_{0,0}^{-1}(\rho) = \sin(\rho))
\end{aligned}$$

■

Thus we obtain,

Theorem 5.7.2. *For any $1 \leq q \leq 2 \leq p \leq \infty$, let $a := p^* - 1, b = q - 1$. Then for any $m, n \in \mathbb{N}$ and $A \in \mathbb{R}^{m\times n}$, $\mathsf{CP}(A)/\|A\|_{p\to q} \leq 1/(h_{a,b}^{-1}(1) \cdot \gamma_q\, \gamma_{p^*})$ and moreover*

- $h_{1,b}^{-1}(1) = h_{a,1}^{-1}(1) = 1$.

- $h^{-1}_{a,b}(1) \geq \sinh^{-1}(1)/(1+\varepsilon_0)$ *where* $\varepsilon_0 = 0.00863$.

Proof. The first inequality follows from Lemma 5.5.7. As for the next item, If $p = 2$ or $q = 2$ (i.e., $a = 1$ or $b = 1$) we are trivially done since $h^{-1}_{a,b}(\rho) = \rho$ in that case (since for $k \geq 1$, $(0)_k = 0$). So we may assume that $a, b \in [0, 1)$.

We are left with proving the final part of the claim. Now using Mathematica we verify (exactly)[1] that $(C1)$ and $(C2)$ are true for $k \leq 29$. Now let $\delta = \sinh^{-1}(0.974203)$. Then by Lemma 5.7.7 (which states that f^{-1}_k decays exponentially and will be proven in the subsequent section),

$$h_{err}(31, \delta) := \sum_{k \geq 31} |f^{-1}_k| \cdot d^k \leq \frac{6.1831}{31} \cdot \frac{\delta^{31}}{1-\delta^2} \leq 0.0128991\ldots.$$

Now by Lemma 5.7.1 we know $h^{-1}_{a,b}(1) \geq \sinh^{-1}(1 - 2h_{err}(31, \delta))$. Thus, $h^{-1}_{a,b}(1) \geq \sinh^{-1}(0.974202) \geq \sinh^{-1}(1)/(1+\varepsilon_0)$ for $\varepsilon_0 = 0.00863$, which completes the proof. ∎

5.7.2 Bounding Inverse Coefficients

In this section we prove that f^{-1}_k decays as $1/c^k$ for some $c = c(a, b) > 1$, proving Lemma 5.7.7. Throughout this section we assume $1 \leq p^*, q < 2$, and $a = p^* - 1$, $b = q - 1$ (i.e., $a, b \in [0, 1)$). Via the power series representation, $f_{a,b}(z)$ can be analytically continued to the unit complex disk. Let $f^{-1}_{a,b}(z)$ be the inverse of $f_{a,b}(z)$ and recall f^{-1}_k denotes its k-th Taylor coefficient.

We begin by stating a standard identity from complex analysis that provides a convenient contour integral representation of the Taylor coefficients of the inverse of a function. We include a proof for completeness.

Lemma 5.7.3 (Inversion Formula)**.** *There exists $\delta > 0$, such that for any odd k,*

$$f^{-1}_k = \frac{2}{\pi k} \Im\left(\int_{C^+_\delta} f_{a,b}(z)^{-k}\, dz \right) \tag{5.4}$$

where C^+_δ denotes the first quadrant quarter circle of radius δ with counter-clockwise orientation.

Proof. Via the power series representation, $f_{a,b}(z)$ can be analytically continued to the unit complex disk. Thus by inverse function theorem for holomorphic functions, there exists $\delta_0 \in (0, 1]$ such that $f_{a,b}(z)$ has an analytic inverse in the open disk $|z| < \delta_0$. So for $\delta \in (0, \delta_0)$, $f_{a,b}(C_\delta)$ is a simple closed curve with winding number 1 (where C_δ is the complex circle of radius δ with the usual counter-clockwise orientation). Thus by Cauchy's integral formula we have

$$f^{-1}_k = \frac{1}{2\pi i} \int_{f_{a,b}(C_\delta)} \frac{f^{-1}_{a,b}(w)}{w^k}\, dw = \frac{1}{2\pi i} \int_{C_\delta} \frac{z \cdot f'_{a,b}(z)}{f_{a,b}(z)^{k+1}}\, dz$$

[1] We generated f^{-1}_k as a polynomial in a and b and maximized it over $a, b \in [0, 1]$ using the Mathematica "Maximize" function which is exact for polynomials.

where the second equality follows from substituting $w = f_{a,b}(z)$.

Now by Fact 5.7.5, $z/f_{a,b}(z)^k$ is holomorphic on the open set $|z| \in (0,1)$, which contains C_δ. Hence by the fundamental theorem of contour integration we have

$$\int_{C_\delta} \frac{d}{dz}\left(\frac{z}{f_{a,b}(z)^k}\right) dz = 0 \quad \Rightarrow \quad \int_{C_\delta} \frac{z \cdot f'_{a,b}(z)}{f_{a,b}(z)^{k+1}}\, dz = \frac{1}{k} \int_{C_\delta} \frac{1}{f_{a,b}(z)^k}\, dz$$

So we get,

$$f_k^{-1} = \frac{1}{2\pi i k} \int_{C_\delta} f_{a,b}(z)^{-k}\, dz = \frac{1}{2\pi k} \Im\left(\int_{C_\delta} f_{a,b}(z)^{-k}\, dz\right)$$

where the second equality follows since f_k^{-1} is purely real. Lastly, we complete the proof of the claim by using the fact that for odd k, $f_{a,b}(z)^{-k}$ is odd and that $\overline{f_{a,b}(z)} = f_{a,b}(\overline{z})$. ■

We next state a standard bound on the magnitude of a contour integral that we will use in our analysis.

Fact 5.7.4 (ML-inequality)**.** *If f is a complex valued continuous function on a contour Γ and $|f(z)|$ is bounded by M for every $z \in \Gamma$, then*

$$\left|\int_\Gamma f(z)\right| \leq M \cdot \ell(\Gamma)$$

where $\ell(\Gamma)$ is the length of Γ.

Unfortunately the integrand in Eq. (5.4) can be very large for small δ, and we cannot use the ML-inequality as is. To fix this, we modify the contour of integration (using Cauchy's integral theorem) so that the imaginary part of the integral vanishes when restricted to the sections close to the origin, and the integrand is small in magnitude on the sections far from the origin (thus allowing us to use the ML-inequality). To do this we will need some preliminaries.

$f_{a,b}(z)$ is defined on the closed complex unit disk. The domain is analytically extended to the region $\mathbb{C} \setminus ((-\infty,-1) \cup (1,\infty))$, using the Euler-type integral representation of the hypergeometric function.

$$f^+_{a,b}(z) := \mathrm{B}\left(\frac{1-b}{2}, 1+\frac{b}{2}\right)^{-1} \cdot \mathrm{I}(z)$$

where $\mathrm{B}(\tau_1, \tau_2)$ is the beta function and

$$\mathrm{I}(z) := z \int_0^1 \frac{(1-t)^{b/2}\, dt}{t^{(1+b)/2} \cdot (1-z^2 t)^{(1-a)/2}}.$$

Fact 5.7.5. *For any $a_1 > 0$, ${}_2F_1(a_1, b_1, c_1, z)$ has no non-zero roots in the region $\mathbb{C} \setminus (1,\infty)$. This implies that if $p^* < 2$, $f^+_{a,b}(z)$ has no non-zero roots in the region $\mathbb{C} \setminus ((-\infty,-1) \cup (1,\infty))$.*

Figure 5.1: The Contour $P(\alpha, \varepsilon)$

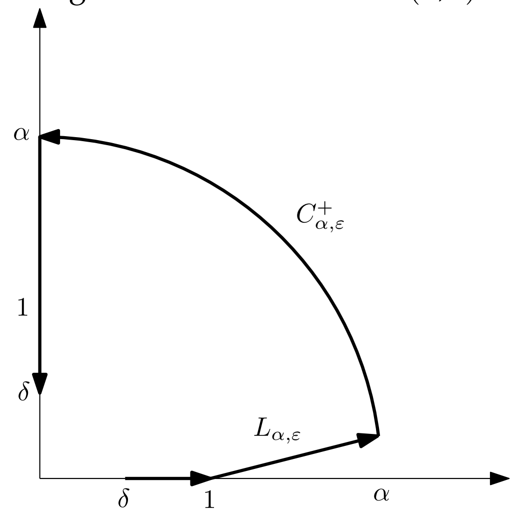

We are now equipped to expand the contour. Our choice of contour is inspired by that of Haagerup [Haa81] which he used in deriving an upper bound on the complex Grothendieck constant. The contour we choose has some differences for technical reasons related to the region to which hypergeometric functions can be analytically extended. The analysis is quite different from that of Haagerup since the functions in consideration behave differently. In fact the inverse function Haagerup considers has polynomially decaying coefficients while the class of inverse functions we consider have coefficients that have decay between exponential and factorial.

Observation 5.7.6 (Expanding Contour). *For any $\alpha \geq 1$ and $\varepsilon > 0$, let $P(\alpha, \varepsilon)$ be the four-part curve (see Fig. 5.1) given by*

- *the line segment $\delta \rightarrow (1-\varepsilon)$,*
- *the line segment $(1-\varepsilon) \rightarrow (\sqrt{\alpha-\varepsilon} + i\sqrt{\varepsilon})$ (henceforth referred to as $L_{\alpha,\varepsilon}$),*
- *the arc along C^+_α starting at $(\sqrt{\alpha-\varepsilon} + i\sqrt{\varepsilon})$ and ending at $i\alpha$ (henceforth referred to as $C^+_{\alpha,\varepsilon}$),*
- *the line segment $i\alpha \rightarrow i\delta$.*

By Cauchy's integral theorem, combining Lemma 5.7.3 with Fact 5.7.5 yields that for odd k,

$$f_k^{-1} = \frac{2}{\pi k} \Im\left(\int_{P(\alpha,\varepsilon)} f^+_{a,b}(z)^{-k}\, dz\right)$$

We will next see that the imaginary part of our contour integral vanishes on section of $P(\alpha, \varepsilon)$. Applying ML-inequality to the remainder of the contour, combined with lower bounds on $|f^+_{a,b}(z)|$ (proved below the fold in Section 5.7.2), allows us to derive an exponentially decaying upper bound on $|f_k^{-1}|$.

Lemma 5.7.7. *For any $1 \leq p^*, q < 2$, there exists $\varepsilon > 0$ such that*

$$|f_k^{-1}| \leq \frac{6.1831}{k(1+\varepsilon)^k}.$$

Proof. For a contour P, we define $V(P)$ as

$$V(P) := \frac{2}{\pi k} \Im\left(\int_P f_{a,b}^+(z)^{-k}\, dz \right)$$

As is evident from the integral representation, $f_{a,b}^+(z)$ is purely imaginary if z is purely imaginary, and as is evident from the power series, $f_{a,b}(z)$ is purely real if z lies on the real interval $[-1,1]$. This implies that $V(\delta \to (1-\varepsilon)) = V(i\alpha \to i\delta) = 0$.

Now combining Fact 5.7.4 (ML-inequality) with Lemma 5.7.9 and Lemma 5.7.12 (which state that the integrand is small in magnitude over $C_{6,\varepsilon}^+$ and $L_{6,\varepsilon}$ respectively), we get that for sufficiently small $\varepsilon > 0$,

$$\begin{aligned} |V(P(6,\varepsilon))| &\leq |V(C_{6,\varepsilon}^+)| + |V(L_{6,\varepsilon})| \\ &\leq \frac{2}{\pi k} \cdot \frac{3\pi/2}{(1+\varepsilon)^k} + \frac{2}{\pi k} \cdot \frac{6 - 1 + O(\sqrt{\varepsilon})}{(1+\varepsilon)^k} \\ &\leq \frac{6.1831}{k(1+\varepsilon)^k}. \qquad \text{(taking } \varepsilon \text{ sufficiently small)} \end{aligned}$$

■

Lower bounds on $|f_{a,b}^+(z)|$ Over $C_{\alpha,\varepsilon}^+$ and $L_{\alpha,\varepsilon}$

In this section we show that for sufficiently small ε, $|f_{a,b}^+(z)| > 1$ over $L_{\alpha,\varepsilon}$ (regardless of the value of α, Lemma 5.7.12), and over $C_{\alpha,\varepsilon}^+$ when α is a sufficiently large constant (Lemma 5.7.9).

We will first show the claim for $C_{\alpha,\varepsilon}^+$ by relating $|f_{a,b}^+(z)|$ to $|z|$. While the asymptotic behavior of hypergeometric functions for $|z| \to \infty$ has been extensively studied (see for instance [Loz03]), it appears that our desired estimates aren't immediate consequences of prior work for two reasons. Firstly, we require relatively precise estimates for moderately large but constant $|z|$. Secondly, due to the expressive power of hypergeometric functions, the estimates we derive can only be true for hypergeometric functions parameterized in a specific range. Indeed, our proof crucially uses the fact that $a, b \in [0,1)$. Our approach is to use the Euler-type integral representation of $f_{a,b}^+(z)$ which as a reminder to the reader is as follows:

$$f_{a,b}^+(z) := \mathrm{B}\left(\frac{1-b}{2}, 1 + \frac{b}{2}\right)^{-1} \cdot \mathrm{I}(z)$$

where $\mathrm{B}(x,y)$ is the beta function and

$$\mathrm{I}(z) := z \int_0^1 \frac{(1-t)^{b/2}\, dt}{t^{(1+b)/2} \cdot (1-z^2 t)^{(1-a)/2}}.$$

We start by making the simple observation that the integrand of $\mathrm{I}(z)$ is always in the positive complex quadrant — an observation that will come in handy multiple times in this section, in dismissing the possibility of cancellations. This is the part of our proof that makes the most crucial use of the assumption that $0 \le a < 1$ (equivalently $1 \le p^* < 2$).

Observation 5.7.8. *Let $z = re^{i\theta}$ be such that either one of the following two cases is satisfied:*

(A) $r < 1$ and $\theta = 0$.

(B) $\theta \in (0, \pi/2]$.

Then for any $0 \le a \le 1$ and any $t \in \mathbb{R}^+$,

$$\arg\left(\frac{z}{(1-tz^2)^{(1-a)/2}}\right) \in [0, \pi/2]$$

Proof. The claim is clearly true when $\theta = 0$ and $r < 1$. It is also clearly true when $\theta = \pi/2$. Thus we may assume $\theta \in (0, \pi/2)$.

$$\begin{aligned} \arg(z) \in (0,\pi/2) \quad &\Rightarrow \arg(-tz^2) \in (-\pi, 0) \quad \Rightarrow \Im(-tz^2) < 0 \\ \Rightarrow \Im(1-tz^2) < 0 \quad &\Rightarrow \arg(1-tz^2) \in (-\pi, 0) \end{aligned}$$

Moreover since $\arg(-tz^2) = 2\theta - \pi \in (-\pi, 0)$, we have $\arg(1-tz^2) > 2\theta - \pi$. Thus we have,

$$\begin{aligned} &\arg(1-tz^2) \in (2\theta-\pi, 0) \quad \Rightarrow \arg\left((1-tz^2)^{(1-a)/2}\right) \in ((1-a)(\theta-\pi/2), 0) \\ &\Rightarrow \arg\left(1/(1-tz^2)^{(1-a)/2}\right) \in (0, (1-a)(\pi/2-\theta)) \\ &\Rightarrow \arg\left(z/(1-tz^2)^{(1-a)/2}\right) \in (0, (1-a)(\pi/2-\theta)+\theta) \subseteq (0, \pi/2) \end{aligned}$$

■

We now show $|f^+_{a,b}(z)|$ is large over $C^+_{\alpha,\varepsilon}$. The main idea is to move from a complex integral to a real integral with little loss, and then estimate the real integral. To do this, we use Observation 5.7.8 to argue that the magnitude of $\mathrm{I}(z)$ is within $\sqrt{2}$ of the integral of the magnitude of the integrand.

Lemma 5.7.9 ($|f^+_{a,b}(z)|$ is large over $C^+_{\alpha,\varepsilon}$). *Assume $a, b \in [0,1)$ and consider any $z \in \mathbb{C}$ with $|z| \ge 6$. Then $|f^+_{a,b}(z)| > 1$.*

Proof. We start with a useful substitution.

$$\begin{aligned} \mathrm{I}(z) &= z \int_0^1 \frac{(1-t)^{b/2}\, dt}{t^{(1+b)/2} \cdot (1-z^2 t)^{(1-a)/2}} \\ &= r^b e^{i\theta} \int_0^{r^2} \frac{(1-s/r^2)^{b/2}\, ds}{s^{(1+b)/2} \cdot (1-e^{2i\theta}s)^{(1-a)/2}} \qquad \text{(Subst. } s = r^2 t, \text{ where } z = re^{i\theta}) \\ &= r^b \int_0^{r^2} \frac{w_a(s,\theta) \cdot (1-s/r^2)^{b/2}\, ds}{s^{1+(b-a)/2}} \end{aligned}$$

where

$$w_a(s,\theta) := \frac{e^{i\theta}}{(1/s - e^{2i\theta})^{(1-a)/2}}\,.$$

We next exploit the observation that the integrand is always in the positive complex quadrant by showing that $|\mathrm{I}(z)|$ is at most a factor of $\sqrt{2}$ away from the integral obtained by replacing the integrand with its magnitude.

$$\begin{aligned}
&|\mathrm{I}(z)| \\
=\ & \sqrt{\mathbb{R}(\mathrm{I}(z))^2 + \Im(\mathrm{I}(z))^2} \\
\geq\ & (|\mathbb{R}(\mathrm{I}(z))| + |\Im(\mathrm{I}(z))|)/\sqrt{2} && \text{(Cauchy-Schwarz)} \\
=\ & (\mathbb{R}(\mathrm{I}(z)) + \Im(\mathrm{I}(z)))/\sqrt{2} && \text{(by Observation 5.7.8)} \\
=\ & \frac{r^b}{\sqrt{2}} \int_0^{r^2} (\mathbb{R}(w_a(s,\theta)) + \Im(w_a(s,\theta))) \cdot \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} \\
=\ & \frac{r^b}{\sqrt{2}} \int_0^{r^2} (|\mathbb{R}(w_a(s,\theta))| + |\Im(w_a(s,\theta))|) \cdot \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} && \text{(by Observation 5.7.8)} \\
\geq\ & \frac{r^b}{\sqrt{2}} \int_0^{r^2} |w_a(s,\theta)| \cdot \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} && (\|v\|_1 \geq \|v\|_2) \\
\geq\ & \frac{r^b}{\sqrt{2}} \int_0^{r^2} \frac{1}{(1+1/s)^{(1-a)/2}} \cdot \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}}
\end{aligned}$$

We now break the integral into two parts and analyze them separately. We start by analyzing the part that's large when $b \to 0$.

$$\begin{aligned}
&\frac{r^b}{\sqrt{2}} \int_1^{r^2} \frac{1}{(1+1/s)^{(1-a)/2}} \cdot \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} \\
\geq\ & \frac{r^b}{2} \int_1^{r^2} \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} \\
\geq\ & \frac{r^b}{2} \int_1^{r^2/2} \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} \\
\geq\ & \frac{r^b}{2\sqrt{2}} \int_1^{r^2/2} \frac{ds}{s^{1+(b-a)/2}} && \text{(since } s \leq r^2/2\text{)} \\
\geq\ & \frac{r^b \cdot \min\{1, r^{a-b}\}}{2\sqrt{2}} \int_1^{r^2/2} \frac{ds}{s} \\
=\ & \frac{\min\{r^a, r^b\} \cdot \log(r^2/2)}{2\sqrt{2}} \\
\geq\ & \frac{\log(r/\sqrt{2})}{\sqrt{2}}
\end{aligned}$$

We now analyze the part that's large when $b \to 1$.

$$\begin{aligned}
&\frac{r^b}{\sqrt{2}} \int_0^1 \frac{1}{(1+1/s)^{(1-a)/2}} \cdot \frac{(1-s/r^2)^{b/2}\,ds}{s^{1+(b-a)/2}} \\
= \;& \frac{r^b}{\sqrt{2}} \int_0^1 \frac{(1-s/r^2)^{b/2}}{(1+s)^{(1-a)/2}} \cdot \frac{ds}{s^{(1+b)/2}} \\
\geq \;& \frac{r^b \cdot \sqrt{1-1/r^2}}{2} \int_0^1 \frac{ds}{s^{(1+b)/2}} && \text{(since } s \leq 1\text{)} \\
= \;& \frac{r^b \cdot \sqrt{1-1/r^2}}{1-b}
\end{aligned}$$

Combining the two estimates above yields that if $r > \sqrt{2}$,

$$|f^+_{a,b}(z)| \geq \mathrm{B}\left(\frac{1-b}{2}, 1+\frac{b}{2}\right)^{-1} \cdot \left(\frac{\log(r/\sqrt{2})}{\sqrt{2}} + \frac{r^b \cdot \sqrt{1-1/r^2}}{1-b}\right)$$

Lastly, the proof follows by using the following estimate:

Fact 5.7.10. *Via Mathematica, for $0 \leq b < 1$ we have*

$$\mathrm{B}\left(\frac{1-b}{2}, 1+\frac{b}{2}\right)^{-1} \cdot \left(\frac{\log(6/\sqrt{2})}{\sqrt{2}} + \frac{6^b \cdot \sqrt{1-1/6^2}}{1-b}\right) \;\geq\; 1.003 \qquad \blacksquare$$

Remark 5.7.11. *The preceding proof can be used to derive the precise asymptotic behavior of $|f^+_{a,b}(z)|$ in r. Specifically, it grows as $r^a \log r$ if $a = b$ and as $r^{\max\{a,b\}}$ if $a \neq b$.*

We now show that $|f^+_{a,b}(z)| > 1$ over $L_{\alpha,\varepsilon}$. To do this, it is insufficient to assume that $|z| \geq 1$ since there exist points z (for instance $z = i$) of unit length such that $|f^+_{a,b}(z)| < 1$. To show the claim, we observe that $|f^+_{a,b}(z)|$ is large when z is close to the real line and use the fact that $L_{\alpha,\varepsilon}$ is close to the real line. Formally, we show that if z is of length at least 1 and is sufficiently close to the real line, $|f^+_{a,b}(z)|$ is close to $f^+_{a,b}(1)$. Lastly, we use the power series representation of the hypergeometric function to obtain a sufficiently accurate lower bound on $f^+_{a,b}(1)$.

Lemma 5.7.12 ($|f^+_{a,b}(z)|$ is large over $L_{\alpha,\varepsilon}$). *Assume $a, b \in [0, 1)$ and consider any $\gamma \geq 1 - \varepsilon_1$. Let $\varepsilon_2 := \sqrt{\varepsilon_1}$ and $z := \gamma(1 + i\varepsilon_1)$. Then for $\varepsilon_1 > 0$ sufficiently small, $|f^+_{a,b}(z)| > 1$.*

Proof. Below the fold we will show

$$|\mathrm{I}(z)| \;\geq\; (1 - O(\sqrt{\varepsilon_1})) \int_0^{1-\varepsilon_2} \frac{(1-s)^{b/2}\,ds}{s^{(1+b)/2} \cdot (1-s)^{(1-a)/2}} \tag{5.5}$$

But we know (LHS, RHS refer to Eq. (5.5))

$$\mathrm{B}\left(\frac{1-b}{2}, 1+\frac{b}{2}\right)^{-1} \cdot \mathrm{LHS} \;=\; f^+_{a,b}(z) \quad \text{and}$$

$$\mathrm{B}\left(\frac{1-b}{2}, 1+\frac{b}{2}\right)^{-1} \cdot \mathrm{RHS} \;\to\; f_{a,b}(1) \text{ as } \varepsilon_1 \to 0$$

Also by Corollary 5.6.19 : (M1), (M2), $f_{a,b}(1) \geq 1 + (1-a)(1-b)/6 > 1$. Thus for ε_1 sufficiently small, we must have $|f^+_{a,b}(z)| > 1$.

We now show Eq. (5.5), by comparing integrands point-wise. To do this, we will assume the following closeness estimate that we will prove below the fold:

$$\mathbb{R}\left(\frac{1+i\varepsilon_1}{(1-s(1+i\varepsilon_1)^2)^{(1-a)/2}}\right) = \frac{1-O(\varepsilon_2)}{(1-s)^{(1-a)/2}}. \tag{5.6}$$

We will also need the following inequality. Since $\gamma \geq 1-\varepsilon_1 = 1-\varepsilon_2^2$, for any $0 \leq s \leq 1-\varepsilon_2$, we have

$$(1-s/\gamma^2)^{b/2} \geq (1-O(\varepsilon_2)) \cdot (1-s)^{b/2}. \tag{5.7}$$

Given, these estimates, we can complete the proof of Eq. (5.5) as follows:

$$\begin{aligned}
&\mathbb{R}(\mathrm{I}(z)) \\
&= \mathbb{R}\left(z\int_0^1 \frac{(1-t)^{b/2}\,dt}{t^{(1+b)/2}\cdot(1-tz^2)^{(1-a)/2}}\right) \\
&= \mathbb{R}\left(\gamma^b(1+i\varepsilon_1)\int_0^{\gamma^2} \frac{(1-s/\gamma^2)^{b/2}\,ds}{s^{(1+b)/2}\cdot(1-s(1+i\varepsilon_1)^2)^{(1-a)/2}}\right) && \text{(subst. } s \leftarrow \gamma^2 t\text{)} \\
&\geq \mathbb{R}\left(\gamma^b(1+i\varepsilon_1)\int_0^{1-\varepsilon_2} \frac{(1-s/\gamma^2)^{b/2}\,ds}{s^{(1+b)/2}\cdot(1-s(1+i\varepsilon_1)^2)^{(1-a)/2}}\right) && \text{(by Observation 5.7.8)} \\
&= \gamma^b\int_0^{1-\varepsilon_2} \mathbb{R}\left(\frac{1+i\varepsilon_1}{(1-s(1+i\varepsilon_1)^2)^{(1-a)/2}}\right)\frac{(1-s/\gamma^2)^{b/2}\,ds}{s^{(1+b)/2}} \\
&\geq (1-O(\varepsilon_2))\cdot\gamma^b\int_0^{1-\varepsilon_2} \frac{(1-s/\gamma^2)^{b/2}\,ds}{s^{(1+b)/2}\cdot(1-s)^{(1-a)/2}} && \text{(by Eq. (5.6))} \\
&\geq (1-O(\varepsilon_2))\int_0^{1-\varepsilon_2} \frac{(1-s)^{b/2}\,ds}{s^{(1+b)/2}\cdot(1-s)^{(1-a)/2}} && \text{(by Eq. (5.7), } \gamma \geq 1-\varepsilon_1\text{)}
\end{aligned}$$

It remains to establish Eq. (5.6), which we will do by considering the numerator and reciprocal of the denominator separately and subsequently using the fact that $\mathbb{R}(z_1z_2) = \mathbb{R}(z_1)\mathbb{R}(z_2) - \Im(z_1)\Im(z_2)$. In doing this, we need to show that the respective real parts are large and respective imaginary parts are small for which the following simple facts will come in handy.

Fact 5.7.13. *Let $z = re^{i\theta}$ be such that $\mathbb{R}z \geq 0$ (i.e. $-\pi/2 \leq \theta \leq \pi/2$). Then for any $0 \leq \alpha \leq 1$,*

$$\mathbb{R}(1/z^\alpha) = \cos(-\alpha\theta)/r^\alpha = \cos(\alpha\theta)/r^\alpha \geq \cos(\theta)/r^\alpha = \mathbb{R}(z)/r^{1+\alpha}$$

Fact 5.7.14. *Let $z = re^{-i\theta}$ be such that $\mathbb{R}z \geq 0, \Im z \leq 0$ (i.e. $0 \leq \theta \leq \pi/2$). Then for any $0 \leq \alpha \leq 1$,*

$$\Im(1/z^\alpha) = \sin(\alpha\theta)/r^\alpha \leq \sin(\theta)/r^\alpha = -\Im(z)/r^{1+\alpha}$$

We are now ready to prove the claimed properties of the reciprocal of the denominator from Eq. (5.6). For any $0 \le s \le 1 - \varepsilon_2$ we have,

$$
\begin{aligned}
&\mathbb{R}\left(\frac{1}{(1-s(1+i\varepsilon_1)^2)^{(1-a)/2}}\right) \\
&= \mathbb{R}\left(\frac{1}{(1-s+s\varepsilon_1^2-2is\varepsilon_1)^{(1-a)/2}}\right) \\
&= \frac{1}{(1-s)^{(1-a)/2}} \cdot \mathbb{R}\left(\frac{1}{(1+s\varepsilon_1^2/(1-s)-2i\varepsilon_1/(1-s))^{(1-a)/2}}\right) \\
&\ge \frac{1}{(1-s)^{(1-a)/2} \cdot (1+O(\varepsilon_1^2/\varepsilon_2^2))^{(3-a)/4}} && \text{(by Fact 5.7.13, and } 1-s \ge \varepsilon_2\text{)} \\
&= \frac{1-O(\varepsilon_1^2/\varepsilon_2^2)}{(1-s)^{(1-a)/2}} \\
&= \frac{1-O(\varepsilon_2)}{(1-s)^{(1-a)/2}} && (5.8)
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
&\Im\left(\frac{1}{(1-s+s\varepsilon_1^2-2is\varepsilon_1)^{(1-a)/2}}\right) \\
&= \frac{1}{(1-s)^{(1-a)/2}} \cdot \Im\left(\frac{1}{(1+s\varepsilon_1^2/(1-s)-2i\varepsilon_1/(1-s))^{(1-a)/2}}\right) \\
&\le \frac{2\varepsilon_1}{(1-s)^{(1-a)/2}} && \text{(by Fact 5.7.14)} \quad (5.9)
\end{aligned}
$$

Combining Eq. (5.8) and Eq. (5.9) with the fact that $\mathbb{R}(z_1 z_2) = \mathbb{R}(z_1)\mathbb{R}(z_2) - \Im(z_1)\Im(z_2)$ yields,

$$
\mathbb{R}\left(\frac{1+i\varepsilon_1}{(1-s(1+i\varepsilon_1)^2)^{(1-a)/2}}\right) = \frac{1-O(\varepsilon_2)}{(1-s)^{(1-a)/2}} .
$$

This completes the proof. ■

Challenges of Proving (C1) and (C2) for all k

For certain values of a and b, the inequalities in (C1) and (C2) leave very little room for error. In particular, when $a = b = 0$, (C1) holds at equality and (C2) has $1/k!$ additive slack. In this special case, it would mean that one cannot analyze the contour integral (for the k-th coefficient of $f_{a,b}^{-1}(\rho)$) by using ML-inequality on any section of the contour that is within a distance of $\exp(k)$ from the origin. Analytic approaches would require extremely precise estimates on the value of the contour integral on parts close to the origin. Other challenges to naive approaches come from the lack of monotonicity properties for f_k^{-1} (both in k and in a, b - see Section 5.7.1)

5.8 Factorization of Linear Operators

In this section we will show that our approximation results imply improved bounds on $\Phi(\ell_p^n, \ell_q^m)$ for certain values of p and q. (see Section 3.8 for a recap on factorization).

5.8.1 Integrality Gap Implies Factorization Upper Bound

Known upper bounds on $\Phi(X, Y)$ involve Hahn-Banach separation arguments. In this section we see that for a special class of Banach spaces admitting a convex programming relaxation, $\Phi(X, Y)$ is bounded by the integrality gap of the relaxation as an immediate consequence of Convex programming duality (which of course uses a separation argument under the hood). A very similar observation had already been made by Tropp [Tro09] in the special case of $X = \ell_\infty^n, Y = \ell_1^m$ with a slightly different convex program.

For norms X over $\mathbb{R}^n$, Y over $\mathbb{R}^m$ and an operator $A : X \to Y$, we define

$$\Phi_3(A) := \inf_{D_1 B D_2 = A} \frac{\|D_2\|_{X\to 2} \cdot \|B\|_{2\to 2} \cdot \|D_1\|_{2\to Y}}{\|A\|_{X\to Y}} \qquad \Phi_3(X, Y) := \sup_{A:X\to Y} \Phi_3(A)$$

where the infimum runs over diagonal matrices D_1, D_2 and $B \in \mathbb{R}^{m\times n}$. Clearly, $\Phi(A) \le \Phi_3(A)$ and therefore $\Phi(X, Y) \le \Phi_3(X, Y)$.

Henceforth let X be exactly an exactly 2-convex norm over $\mathbb{R}^n$ and Y^* be an exactly 2-convex norm over $\mathbb{R}^m$ (i.e., Y is exactly 2-concave). As was the approach of Grothendieck, we give an upper bound on $\Phi(X, Y)$ by giving an upper bound on $\Phi_3(X, Y)$, which we do by showing

Lemma 5.8.1. *Let X be an exactly 2-convex (sign-invariant) norm over $\mathbb{R}^n$ and Y^* be an exactly 2-convex (sign-invariant) norm over $\mathbb{R}^m$. Then for any $A : X \to Y$, $\Phi_3(A) \le \mathrm{DP}(A)/\|A\|_{X\to Y}$.*

Proof. Consider an optimal solution to $\mathrm{DP}(A)$. We will show

$$\inf_{D_1 B D_2 = A} \|D_2\|_{X\to 2} \cdot \|B\|_{2\to 2} \cdot \|D_1\|_{2\to Y} \le \mathrm{DP}(A)$$

by taking $D_1 := D_s^{1/2}$, $D_2 := D_t^{1/2}$ and $B := \left(D_s^{1/2}\right)^\dagger A \left(D_t^{1/2}\right)^\dagger$ (where for a diagonal matrix D, $D^\dagger$ only inverts the non-zero diagonal entries and zero-entries remain the same). Note that $s_i = 0$ (resp. $t_i = 0$) implies the i-th row (resp. i-th column) of A is all zeroes, since otherwise one can find a 2×2 principal submatrix (of the block matrix in the relaxation) that is not PSD. This implies that $D_1 B D_2 = A$.

It remains to show that $\|D_2\|_{X\to 2} \cdot \|B\|_{2\to 2} \cdot \|D_1\|_{2\to Y} \le \mathrm{DP}(A)$. Now we have,

$$\|D_t^{1/2}\|_{X\to 2} = \sup_{x\in \mathrm{Ball}(X)} \|D_t^{1/2} x\|_2 = \sup_{x\in \mathrm{Ball}(X)} \sqrt{\langle t\,, [x]^2\rangle} = \sup_{\widetilde{x}\in \mathrm{Ball}(X^{(1/2)})} \sqrt{|\langle t\,, \widetilde{x}\rangle|} = \sqrt{\|t\|_{X^{(1/2)}}}\,.$$

Similarly, since $\|D_1\|_{2\to Y} = \|D_1\|_{Y^*\to 2}$ we have

$$\|D_1\|_{Y^*\to 2} \le \sqrt{\|s\|_{Y^{*(1/2)}}}\,.$$

Thus it suffices to show $\|B\|_{2\to 2} \le 1$ since

$$\|D_2\|_{X\to 2} \cdot \|D_1\|_{2\to Y} \le \sqrt{\|t\|_{X^{(1/2)}} \cdot \|s\|_{Y^{*(1/2)}}} \le (\|s\|_{Y^{*(1/2)}} + \|t\|_{X^{(1/2)}})/2 = \mathrm{DP}(A)\,.$$

We have,

$$\begin{aligned}
&\begin{bmatrix} D_s & -A \\ -A^T & D_t \end{bmatrix} \succeq 0 \\
\Rightarrow &\begin{bmatrix} \left(D_s^{1/2}\right)^{\dagger} & 0 \\ 0 & \left(D_t^{1/2}\right)^{\dagger} \end{bmatrix} \begin{bmatrix} D_s & -A \\ -A^T & D_t \end{bmatrix} \begin{bmatrix} \left(D_s^{1/2}\right)^{\dagger} & 0 \\ 0 & \left(D_t^{1/2}\right)^{\dagger} \end{bmatrix} \succeq 0 \\
\Rightarrow &\begin{bmatrix} D_{\bar{s}} & -B \\ -B^T & D_{\bar{t}} \end{bmatrix} \succeq 0 \qquad \text{for some } \bar{s} \in \{0,1\}^m,\ \bar{t} \in \{0,1\}^n \\
\Rightarrow &\begin{bmatrix} \mathrm{I} & -B \\ -B^T & \mathrm{I} \end{bmatrix} \succeq 0 \\
\Rightarrow &\ \|B\|_{2\to 2} \le 1
\end{aligned}$$

■

5.8.2 Improved Factorization Bounds for Certain ℓ_p^n, ℓ_q^m

Let $1 \le q \le 2 \le p \le \infty$. Then taking $\mathcal{F}_X$ to be the $\ell_{p/2}^n$ unit ball and $\mathcal{F}_Y$ to be the $\ell_{q^*/2}^m$ unit ball, we have $\sqrt{\mathcal{F}}_X$ and $\sqrt{\mathcal{F}}_Y$ are respectively the unit balls in ℓ_p^n and $\ell_{q^*}^m$. Therefore X and Y as defined above are the spaces ℓ_p^n and ℓ_q^m respectively. Hence we obtain

Theorem 5.8.2 ($\ell_p^n \to \ell_q^m$ factorization)**.** *If $1 \le q \le 2 \le p \le \infty$, then for any $m, n \in \mathbb{N}$ and $\varepsilon_0 = 0.00863$,*

$$\Phi(\ell_p^n, \ell_q^m) \ \le\ \frac{1+\varepsilon_0}{\sinh^{-1}(1)\cdot \gamma_{p^*}\,\gamma_q} \ \le\ \frac{1+\varepsilon_0}{\sinh^{-1}(1)} \cdot C_2(\ell_{p^*}^n)\cdot C_2(\ell_q^m).$$

This improves upon Pisier's bound and for a certain range of (p,q), improves upon K_G as well as the bound of Kwapień-Maurey.

Krivine and independently Nesterov[Nes98] observed that the integrality gap of $\mathsf{CP}(A)$ for any pair of convex sets $\mathcal{F}_X, \mathcal{F}_Y$ is bounded by K_G (Grothendieck's constant). This provides a class of Banach space pairs for which K_G is an upper bound on the factorization constant. We include a proof for completeness.

5.8.3 K_G Bound on Integrality Gap

In this subsection, we prove the observation that for exactly 2-convex X, Y^*, the integrality gap for $X \to Y$ operator norm is always bounded by K_G.

Lemma 5.8.3. *Let X be an exactly 2-convex (sign-invariant) norm over $\mathbb{R}^n$ and Y^* be an exactly 2-convex (sign-invariant) norm over $\mathbb{R}^m$. Then for any $A : X \to Y$, $\mathsf{CP}(A)/\|A\|_{X\to Y} \ \le\ K_G$.*

Proof. Let $B := \frac{1}{2}\begin{bmatrix} 0 & A \\ A^T & 0 \end{bmatrix}$. The main intuition of the proof is to decompose $x \in X$ as $x = |[x]| \circ \operatorname{sgn}[x]$ (where $\circ$ denotes Hadamard/entry-wise multiplication), and then use Grothendieck's inequality on $\operatorname{sgn}[x]$ and $\operatorname{sgn}[y]$. Another simple observation is that for any convex set $\mathcal{F}$, the feasible set we optimize over is invariant under *factoring out the magnitudes of the diagonal entries*. In other words,

$$\begin{aligned} &\{D_d\, \Sigma\, D_d : d \in \sqrt{|[\mathcal{F}]|} \cap \mathbb{R}^n_{\geq 0},\ \Sigma \succeq 0,\ \operatorname{diag}(\Sigma) = \mathbb{1}\} \\ =&\{\mathbb{X} : \operatorname{diag}(\mathbb{X}) \in \mathcal{F},\ \mathbb{X} \succeq 0\} \end{aligned} \tag{5.10}$$

We will apply the above fact for $\mathcal{F} = \operatorname{Ball}(X^{(1/2)}) \oplus \operatorname{Ball}(Y^{*(1/2)})$. Let X^+ denote $X \cap \mathbb{R}^n_{\geq 0}$ (analogous for $(Y^*)^+$). Now simple algebraic manipulations yield

$$\begin{aligned}
& \|A\|_{X\to Y} \\
= \quad & \sup_{x\in X,\ y\in Y^*} (y\oplus x)^T B (y \oplus x) \\
= \quad & \sup_{\substack{d_x\in X^+,\ \sigma_x\in\{\pm1\}^n, \\ d_y\in (Y^*)^+,\ \sigma_y\in\{\pm1\}^m}} ((d_y\circ\sigma_y)\oplus(d_x\circ\sigma_x))^T\, B\, ((d_y\circ\sigma_y)\oplus(d_x\circ\sigma_x)) \\
= \quad & \sup_{\substack{d_x\in X^+,\ \sigma_x\in\{\pm1\}^n, \\ d_y\in (Y^*)^+,\ \sigma_y\in\{\pm1\}^m}} (\sigma_y\oplus\sigma_x)^T (D_{d_y\oplus d_x}\, B\, D_{d_y\oplus d_x})(\sigma_y\oplus\sigma_x) \\
\geq \quad & (1/K_G)\cdot \sup_{\substack{d_x\in X^+,\ d_y\in(Y^*)^+, \\ \Sigma:\ \operatorname{diag}(\Sigma)=\mathbb{1},\ \Sigma\succeq 0}} \left\langle \Sigma\ ,\ D_{d_y\oplus d_x}\, B\, D_{d_y\oplus d_x} \right\rangle && \text{(Grothendieck)} \\
= \quad & (1/K_G)\cdot \sup_{\substack{d_x\in X^+,\ d_y\in(Y^*)^+, \\ \Sigma:\ \operatorname{diag}(\Sigma)=\mathbb{1},\ \Sigma\succeq 0}} \left\langle D_{d_y\oplus d_x}\, \Sigma\, D_{d_y\oplus d_x}\ ,\ B \right\rangle \\
= \quad & (1/K_G)\cdot \mathsf{CP}(A) && \text{(by Eq. (5.10))}
\end{aligned}$$

■

Chapter 6

Hardness Results for $2 \to (2\text{-convex})$ Norms

In this chapter we reduce from random label cover (admitting polynomial level SoS gaps) to show polynomial factor gaps for certain $2 \to$(2-convex) operator norms (specifically $2 \to$(mixed ℓ_p norm)). Informally we show the following:

Theorem (informal)**.** *There exist constants $\varepsilon, \delta, \delta', \delta'', \overline{\delta} > 0$, a family of random label cover instances $\mathcal{L}$ that w.h.p. have satisfiability fraction at most $1/n^{\delta}$, and a family X_n of exactly 2-convex norms such that if there exists an $n^{\delta''}$-approximation algorithm for $2 \to X_n$ norm with runtime $R(n)$, then there exists an algorithm with runtime $R(n^{O(1)})$ that can certify $\mathcal{L}$ is at most $1/n^{\overline{\delta}}$ satisfiable w.h.p.*
Moreover $n^{\delta'}$ levels of SoS w.h.p. cannot certify a bound better than 1 on the satisfiability of $\mathcal{L}$.

6.1 Mixed ℓ_p Norms

For a vector $x \in \mathbb{R}^{[n_1]\times[n_2]}$ we define the $\ell_{p_1}(\ell_{p_2})$ norm (or $p_1(p_2)$ for short) as

$$\left(\sum_{i\in[n_1]}\left(\sum_{j\in[n_2]}|x_{i,j}|^{p_2}\right)^{p_1/p_2}\right)^{1/p_1}$$

i.e., the p_1 norm aggregates the p_2 norms of each of the "buckets".

If $X = \ell_{p_1}(\ell_{p_2})$ then $X^{(1/2)}$ is easily verified to be $\ell_{p_1/2}(\ell_{p_2/2})$ which is a sign invariant norm as long as $p_1, p_2 \geq 2$. Thus $p_1(p_2)$ is an exactly 2-convex norm whenever $p_1, p_2 \geq 2$ (see Section 3.6 for a reminder on 2-convexity). In this section we will give hardness results for $2 \to q(q')$ where $q > 2$ is a constant very close to 2 and $2 < q' < \infty$ is some large constant. Before we state our key technical theorem we first state some structural assumptions on the label cover instances we reduce from.

6.2 Structure of Label Cover

Let $\mathcal{L}$ be a bipartite 2-CSP instance (i.e., each constraint involves a variable from each side) with

- Variables $U \cup V$ with $|U| = |V| = n$. They are Δ-regular. Let R be the label size.
- Let $A \in \mathbb{R}^{n \times n}$ be its bipartite adjacency matrix (we allow parallel edges so $A_{u,v}$ denotes the number of constraints containing u and v). A has exactly $n\Delta$ edges. Let $p = \Delta/n$ be the "density" of this matrix. Then the largest singular value of A is Δ with the all-ones vector as the corresponding left and right singular vectors. Let λ be the second largest singular value. We define E so that

$$A = \Delta\left(\frac{\mathbb{1}_n}{\sqrt{n}}\right)\left(\frac{\mathbb{1}_n}{\sqrt{n}}\right)^T + E = pJ + E,$$

where $\|E\|_2 \leq \lambda$ and J is the all-ones matrix.

We are now ready to state our main technical result:

Theorem 6.2.1. *Let $\mathcal{L}$ be a label cover instance as above and let $q = 2 + \varepsilon$. Then for sufficiently small $\varepsilon > 0$, there is a deterministic polynomial time reduction from $\mathcal{L}$ to a matrix $\overline{A} \in \mathbb{R}^{nR \times nR}$ such that*

- $\|\overline{A}\|_{q^*(1) \to q(\infty)} \geq \mathrm{OPT}(\mathcal{L})$.
- *If $\mathcal{L}$ has value at most s, then $\|\overline{A}\|_{q^*(1) \to q(\infty)} \leq \alpha$ for any α satisfying $\lambda \leq \alpha \cdot p \cdot n^{2/q}/10$ and $\alpha \geq (800s)^{1/7}$.*

6.3 Reduction from Label Cover to $\|\cdot\|_{q^*(1) \to q(\infty)}$

- Fix $1/3 > \varepsilon \geq 0$ and let $q = 2 + \varepsilon$.
- Given $\mathcal{L}$, we produce a matrix $\overline{A} \in \mathbb{R}^{nR \times nR}$ where $((u,i),(v,j))$'s entry is t where t denotes the number of constraints that contain (u,v) and are satisfied by the assignment $u \leftarrow i$, $v \leftarrow j$ (i.e., the bipartite adjacency matrix of the label-extended graph of $\mathcal{L}$).
- Given $\overline{x} \in \mathbb{R}^{[n] \times [R]}$, let $x^u \in \mathbb{R}^R$ be the vector defined as $(x^u)_i := \overline{x}_{u,i}$ and let $x \in \mathbb{R}^n$ be the vector defined as $x_u := \|x^u\|_1$. We define $y^v \in \mathbb{R}^R$ and $y \in \mathbb{R}^n$ similarly.
- We are reducing to the problem

$$\|A\|_{q^*(1) \to q(\infty)} = \sup_{\|\overline{x}\|_{q^*(1)} \leq \|\overline{y}\|_{q^*(1)} \leq 1} \overline{x}^T \overline{A} \overline{y}$$

where $\|\overline{x}\|_{q^*(1)} := \|x\|_{q^*}$.

proof of Theorem 6.2.1.
Completeness. Suppose $\mathcal{L}$ has a labeling that satisfies every constraint. Then let $\overline{x}$ and $\overline{y}$ be the indicator vectors of these labels, in which case x and y equal the all-ones vector, and for every $(u,v) \in \mathcal{L}$ ($\Leftrightarrow A_{u,v} = 1$), we have $(x^u)^T \overline{A}_{u,v}(y^v) = 1$, where $\overline{A}_{u,v}$ denotes the (u,v)'th block submatrix of $\overline{A}$. Thus we have

$$(\overline{x})^T \overline{A}\,\overline{y} = n\Delta.$$

Since $\|\overline{x}\|_{q^*(1)} = \|\overline{y}\|_{q^*(1)} = n^{1/q^*}$, the completeness value is at least

$$c := n\Delta/(n^{2/q^*}) = p \cdot n^{2-2/q^*} = p \cdot n^{2/q}.$$

Soundness. Fix $\overline{x}$ and $\overline{y}$ such that $\|\overline{x}\|_{q^*(1)} = \|\overline{y}\|_{q^*(1)} = 1$ (which implies that $\|x\|_{q^*} = \|y\|_{q^*} = 1$). Since $\overline{A}$ has nonnegative entries, we may assume without loss of generality that all entries of $\overline{x}$ and $\overline{y}$ are nonnegative.

Consider the randomized strategy where every vertex $u \in U$ is assigned a label $i \in [R]$ with probability $(x^u)_i/\|x^u\|_1$ independently of all other vertices and all vertices $v \in V$ are assigned similarly according to $\overline{y}$. Let $B \in \mathbb{R}^{n\times n}$ be defined so that $B_{u,v}$ denotes the expected number of constraints involving (u,v) that are satisfied by this labeling strategy.

Then B has nonnegative entries and $B_{u,v} \le A_{u,v}$ for all u, v. Assuming $\mathcal{L}$ has soundness s, we have

$$\mathbb{1}^T B \mathbb{1} \le s \mathbb{1}^T A \mathbb{1}.$$

Also by definition,

$$\overline{x}^T \overline{A}\,\overline{y} = x^T B y.$$

Let $\alpha \le 1$ be such that $\overline{x}^T \overline{A}\,\overline{y} = x^T B y = c \cdot \alpha$. Our goal is to prove that α goes to 0 with s.

Since $x^T A y \ge x^T B y = c \cdot \alpha$ and $\|x\|_2 \le \|x\|_{q^*} = 1$ (same for y), we have

$$c \cdot \alpha \le x^T A y = p\|x\|_1\|y\|_1 + x^T E y \le p\|x\|_1\|y\|_1 + \lambda. \tag{6.1}$$

Recall we assume that $\lambda \le c \cdot \alpha/10$. Then Eq. (6.1) implies

$$p\|x\|_1\|y\|_1 \ge c \cdot \alpha/2 = \alpha \cdot p \cdot n^{2/q}/2.$$

Further since $\|x\|_1, \|y\|_1 \le n^{1-1/q^*} = n^{1/q}$, we now have $\|x\|_1, \|y\|_1 \ge \alpha \cdot n^{1/q}/2$.

Let $\gamma = 20/\alpha^3$, and let x' be such that $x'_u = x_u$ if $x_u \le \gamma/n^{1/q^*}$ and 0 otherwise. (Note that if x is well-spread then $x_u = 1/n^{1/q^*}$ for all u.)

Claim 6.3.1.

$$\|x'\|_1 \ge \|x\|_1 - \frac{n^{1/q}}{\gamma^{q^*-1}}.$$

Proof. Consider the vector $(x - x')$. Since $(x - x')_u = 0$ if $x \le \gamma/n^{1/q^*}$ and $(x - x')_u = x_u$ otherwise, all nonzero entries of $(x - x')$ are at least $\gamma/n^{1/q^*}$. Since both x and x' are nonnegative, $\|x - x'\|_{q^*} \le \|x\|_{q^*} \le 1$. Given the q^*-norm bound and the lower bound on each nonzero entry, $\|x - x'\|_1$ is maximized when each nonzero entry is equal to $\gamma/n^{1/q^*}$ and there are n/γ^{q^*} entries, so we have

$$\|x\|_1 - \|x'\|_1 \le \|x - x'\|_1 \le (\gamma/n^{1/q^*}) \cdot (n/\gamma^{q^*}) = n^{1/q}/\gamma^{q^*-1}. \qquad \blacksquare$$

Since $\|x\|_1 \ge \alpha \cdot n^{1/q}/2$, substituting $\gamma = 20/\alpha^3$ in the above claim implies that

$$\|x'\|_1 \ge \|x\|_1(1 - 2/(\alpha \cdot \gamma^{q^*-1})) \ge (1 - \alpha/10)\|x\|_1,$$

and similarly $\|y'\|_1 \ge (1 - \alpha/10)\|y\|_1$. Again using the fact that $\|x'\|_2, \|y'\|_2 \le 1$ and the spectral bound on E, we get

$$(x')^T A y' \ge p\|x'\|_1\|y'\|_1 - \lambda \ge (1 - \alpha/5)p\|x\|_1\|y\|_1 - \lambda.$$

Finally, we have

$$x^T Ay - (x')^T Ay' \le p\|x\|_1\|y\|_1 + \lambda - \big((1 - \alpha/5)p\|x\|_1\|y\|_1 - \lambda\big) = (\alpha/5)p\|x\|_1\|y\|_1 + 2\lambda.$$

The final quantity is at most $c \cdot \alpha/2$ because $p\|x\|_1\|y\|_1 \le c$ for any x, y with $\|x\|_{q^*} = \|y\|_{q^*} = 1$ and since we assumed $\lambda \le c \cdot \alpha/10$.

Since x', y' are entrywise dominated by x, y, and B is dominated by A, and every entry is nonnegative, $x^T By - (x')^T By' \le x^T Ay - (x')^T Ay' \le c \cdot \alpha/2$. Since $x^T By = c \cdot \alpha$, $(x')^T By' \ge c \cdot \alpha/2$. Since each entry of x' and y' is bounded by $\gamma/n^{1/q^*}$, $(n^{1/q^*}/\gamma)x'$ and $(n^{1/q^*}/\gamma)y'$ are dominated (entry-wise) by the all ones vectors. This implies that

$$\mathbb{1}^T B\mathbb{1} \ge (n^{2/q^*}/\gamma^2)c \cdot \alpha/2 = n^2 \cdot p \cdot \alpha/(2\gamma^2) = \mathbb{1}^T A\mathbb{1} \cdot \alpha^7/800.$$

Since the soundness of Label Cover implies $\mathbb{1}^T B\mathbb{1} \le s\mathbb{1}^T A\mathbb{1}$, this implies that our problem has soundness at most $\alpha \cdot c$, where $\alpha = (800s)^{1/7}$. $\blacksquare$

It is straightforward to check that we also obtain the following corollary:

Corollary 6.3.2. *Let $\mathcal{L}$ be a label cover instance as above, let $q = 2 + \varepsilon$ and let $\delta > 0$ be some constant. Then for sufficiently small $\varepsilon > 0$ and sufficiently large $2 < q' < \infty$, there is a deterministic polynomial time reduction from $\mathcal{L}$ to a matrix $\overline{A} \in \mathbb{R}^{nR \times nR}$ such that*

- $\|\overline{A}\|_{q^*(q'^*) \to q(q')} \ge \mathrm{OPT}(\mathcal{L})$.
- *if $\mathcal{L}$ has value at most $1/n^\delta$ and $10\lambda/(p \cdot n^{2/q}) \le 1/n^\delta$, then $\|\overline{A}\|_{q^*(q'^*) \to q(q')} \le 1/n^{\delta'}$ for some constant δ'.*

We next define the distribution over label cover instances that we shall reduce from.

6.4 Distribution over Label Cover Instances

We first define a map from point vs. line CSP instances to label cover instances. A constraint of the point vs. line predicate has arity k and alphabet $\mathbb{F}_k$ (the finite field of size k where k is a prime power) and is defined as follows:
given shifts $s_0, \ldots s_{k-1} \in \mathbb{F}_k$, the constraint accepts an assignment $x_0, \ldots x_{k-1} \in \mathbb{F}_k$, if there exist $a, b \in \mathbb{F}_k$ such that $a \cdot i + b = x_i + s_i$ for all $i \in \mathbb{F}_k$.

Given a point vs. line instance φ, we generate a label cover instance $\mathcal{L}(\varphi)$ as follows:

- Let $\mathcal{L}'$ be the label cover instance obtained as the constraint-variable game of φ.

- We square $\mathcal{L}'$ to obtain $\mathcal{L}(\varphi)$: i.e., consider the CSP instance where the two sides are both copies of V (say V_1 and V_2) and we have a constraint (v, v') for each pair of Label Cover constraints $((u, v), (u, v'))$. The new constraint is satisfied by a labeling $(\ell(v), \ell(v'))$ if there exists $\ell(u)$ such that $(\ell(v), \ell(u))$ and $(\ell(v'), \ell(u))$ satisfy both (u, v) and (u, v'). If $\overline{B} \in \mathbb{R}^{nL \times mR}$ denotes the bipartite adjacency matrix of the label extended graph of $\mathcal{L}'$, then the bipartite adjacency matrix of the label extended graph of $\mathcal{L}(\varphi)$ is precisely $\overline{A} = \overline{B}^T \overline{B}$.

Our final distribution over label cover instances as follows:

- Sample a random instance φ of the point vs. line CSP with $O(n)$ constraints, where we take the field size to be $k = n^\delta$ for a sufficiently small constant $\delta > 0$. More specifically we sample $O(n)$ random constraints where the subset of k participating variables and the shifts are chosen uniformly at random for each constraint.

- $\mathcal{L}(\varphi)$ is the random label cover instance. It is well known that $\mathrm{OPT}(\mathcal{L}(\varphi)) = \widetilde{O}(1/k)$ w.h.p.

6.5 The Result

Before stating our main result we will describe the hardness assumption we make for random label cover. This can be seen as an analogue of Feige's Random 3-SAT hypothesis [Fei02] and other related hypotheses [Ale03, AAM+11, BKS12] that were used to derive impressive hardness results. Informally, certifying an upper bound of (say) 0.99 w.h.p. on $\mathrm{OPT}(\mathcal{L}(\varphi))$ requires time 2^{n^δ}, where $\mathcal{L}(\varphi)$ is the random label cover instance described above. Formally,

Assumption 6.5.1 (Random Label Cover Hardness)**.** *For sufficiently small $\delta > 0$, there is no 2^{n^δ}-runtime algorithm* ALG *mapping label cover instances to* $[0, 1]$ *and satisfying*

- *for any point vs. line instance φ,* $\mathrm{ALG}(\mathcal{L}(\varphi)) \geq \mathrm{OPT}(\mathcal{L}(\varphi))$.

- *for a random label cover instance $\mathcal{L}(\varphi)$ defined as above,*
  $\mathbf{Pr}[\mathrm{ALG}(\mathcal{L}(\varphi)) \leq 0.99] \geq 1 - o(1)$.

Remark. It has been established in [BCG+12] that the SoS hierarchy satisfies Assumption 6.5.1. More precisely, upto $n^{\delta'}$ levels of the SoS relaxation for the random label cover instance $\mathcal{L}(\varphi)$ output a value of 1 w.h.p. (even though $\text{OPT}(\mathcal{L}(\varphi)) = \widetilde{O}(1/k)$ w.h.p.).

Combining Corollary 6.3.2 with the fact that $\|B^T B\|_{X^* \to X} = \|B^T\|_{2 \to X}^2$ (see Fact 3.4.2 for the non-trivial direction), we obtain our main result:

Theorem 6.5.2. *Let $q = 2 + \varepsilon$. Then for sufficiently small constants $\varepsilon, \delta'' > 0$ and sufficiently large $2 < q' < \infty$, there is a deterministic polynomial time reduction from (point vs. line) label cover instances to matrices ($\mathcal{L}(\varphi) \mapsto \overline{B}^T$) such that*

- *for any point vs. line instance φ, $\|\overline{B}^T\|_{2 \to q(q')}^2 \geq \text{OPT}(\mathcal{L}(\varphi))$.*

- *for a random instance $\mathcal{L}(\varphi)$ defined as above, $\|\overline{B}^T\|_{2 \to q(q')}^2 \leq 1/n^{\delta''}$ w.h.p.*

Consequently there exists $\delta' > 0$, such that any $1/n^{\delta''}$-approximation algorithm for $2 \to q(q')$ norm with runtime $2^{n^{\delta'}}$ would violate Assumption 6.5.1.

Part III

Optimizing Polynomials of Degree ≥ 3

Chapter 7

Additional Preliminaries (Sum of Squares)

In Chapter 8 and Chapter 9 we give use the sum of squares hierarchy to obtain results for polynomial optimization over the sphere. In this section we introduce the necessary SoS preliminaries. Consequently the following preliminaries are relevant only to those chapters.

7.1 Polynomials

We use $\mathbb{R}_d[x]$ to denote the set of all homogeneous polynomials of degree (exactly) d. Similarly, $\mathbb{R}_d^+[x]$ is used to denote the set of polynomials with non-negative coefficients. All polynomials considered in subsequent chapters will be n-variate and homogeneous (with x denoting the set of n variables $x_1, \ldots, x_n$) unless otherwise stated.

A multi-index is defined as sequence $\alpha \in \mathbb{N}^n$. We use $|\alpha|$ to denote $\sum_{i=1}^n \alpha_i$ and $\mathbb{N}_d^n$ (resp. $\mathbb{N}_{\leq d}^n$) to denote the set of all multi-indices α with $|\alpha| = d$ (resp. $|\alpha| \leq d$). Thus, a polynomial $f \in \mathbb{R}_d[x]$ can be expressed in terms of its coefficients as

$$f(x) \;=\; \sum_{\alpha \in \mathbb{N}_d^n} f_\alpha \cdot x^\alpha \,,$$

where x^α is used to denote the monomial corresponding to α. A polynomial is multilinear if $\alpha \leq \mathbb{1}$ whenever $f_\alpha \neq 0$, where $\mathbb{1}$ denotes the multi-index 1^n. We use the notation α^r to denote the vector $(\alpha_1^r, \ldots, \alpha_n^r)$ for $r \in \mathbb{R}$. Exclusively in Chapter 8, with the exception of absolute-value, any scalar function when applied to a vector/multi-index returns the vector obtained by applying the function entry-wise. We also use $\circ$ to denote the Hadamard (entry-wise) product of two vectors.

Sphere Maximization. For a homogeneous polynomial f, let $f_{\max}$ denote $\sup_{\|x\|_2=1} f(x)$ and let $\|f\|_2$ denote $\sup_{\|x\|_2=1} |f(x)|$.

7.2 Matrices

For $k \in \mathbb{N}$, we will consider $n^k \times n^k$ matrices M with real entries. All square matrices considered in this document should be taken to be symmetric (unless otherwise stated). We index entries of the matrix M as $M[I, J]$ by tuples $I, J \in [n]^k$.

A tuple $I = (i_1, \ldots, i_k)$ naturally corresponds to a multi-index $\alpha(I) \in \mathbb{N}^n_k$ with $|\alpha(I)| = k$, i.e. $\alpha(I)_j = |\{\ell \mid i_\ell = j\}|$. For a tuple $I \in [n]^k$, we define $\mathcal{O}(I)$ the set of all tuples J which correspond to the same multi-index i.e., $\alpha(I) = \alpha(J)$. Thus, any multi-index $\alpha \in \mathbb{N}^n_k$, corresponds to an equivalence class in $[n]^k$. We also use $\mathcal{O}(\alpha)$ to denote the class of all tuples corresponding to α.

Note that a matrix of the form $\left(x^{\otimes k}\right)\left(x^{\otimes k}\right)^T$ has many additional symmetries, which are also present in solutions to programs given by the SoS hierarchy. To capture this, consider the following definition:

[SoS-Symmetry] A matrix M which satisfies $\mathsf{M}[I, J] = \mathsf{M}[K, L]$ whenever $\alpha(I) + \alpha(J) = \alpha(K) + \alpha(L)$ is referred to as SoS-symmetric.

Remark. It is easily seen that every homogeneous polynomial has a unique SoS-Symmetric matrix representation.

For a matrix M, we will henceforth use the shorthand $\|M\|_2$ to denote $\|M\|_{2\to 2}$ (i.e., spectral norm/maximum singular value).

7.3 Pseudoexpectations and Moment Matrices

Let $\mathbb{R}[x]_{\leq q}$ be the vector space of polynomials with real coefficients in variables $x = (x_1, \ldots, x_n)$, of degree at most q. For an even integer q, the degree-q pseudo-expectation operator is a linear operator $\widetilde{\mathbf{E}} : \mathbb{R}[x]_{\leq q} \mapsto \mathbb{R}$ such that

1. $\widetilde{\mathbf{E}}[1] = 1$ for the constant polynomial 1.
2. $\widetilde{\mathbf{E}}[p_1 + p_2] = \widetilde{\mathbf{E}}[p_1] + \widetilde{\mathbf{E}}[p_2]$ for any polynomials $p_1, p_2 \in \mathbb{R}[x]_{\leq q}$.
3. $\widetilde{\mathbf{E}}\left[p^2\right] \geq 0$ for any polynomial $p \in \mathbb{R}[x]_{\leq q/2}$.

The pseudo-expectation operator $\widetilde{\mathbf{E}}$ can be described by a moment matrix $\widehat{\mathsf{M}} \in \mathbb{R}^{\mathbb{N}^n_{\leq q/2} \times \mathbb{N}^n_{\leq q/2}}$ defined as $\widehat{\mathsf{M}}[\alpha, \beta] = \widetilde{\mathbf{E}}\left[x^{\alpha+\beta}\right]$ for $\alpha, \beta \in \mathbb{N}^n_{\leq q/2}$.

For each fixed $t \leq q/2$, we can also consider the principal minor of $\widehat{\mathsf{M}}$ indexed by $\alpha, \beta \in \mathbb{N}^n_t$. This also defines a matrix $M \in \mathbb{R}^{[n]^t \times [n]^t}$ with $M[I, J] = \widetilde{\mathbf{E}}\left[x^{\alpha(I)+\alpha(J)}\right]$. Note that this new matrix M satisfies $M[I, J] = M[K, L]$ whenever $\alpha(I) + \alpha(J) = \alpha(K) + \alpha(L)$. Recall that a matrix in $\mathbb{R}^{[n]^t \times [n]^t}$ with this symmetry is said to be SoS-symmetric.

We will use the following facts about the operator $\widetilde{\mathbf{E}}$ given by the SoS hierarchy.

Claim 7.3.1 (Pseudo-Cauchy-Schwarz [BKS14]). $\widetilde{\mathbf{E}}[p_1 p_2] \leq (\widetilde{\mathbf{E}}\left[p_1^2\right] \widetilde{\mathbf{E}}\left[p_2^2\right])^{1/2}$ *for any* p_1, p_2 *of degree at most* $q/2$.

7.3.1 Constrained Pseudoexpectations

For a system of polynomial constraints $C = \{f_1 = 0, \ldots, f_m = 0, g_1 \geq 0, \ldots, g_r \geq 0\}$, we say $\widetilde{\mathbf{E}}_C$ is a pseudoexpectation operator respecting C, if in addition to the above conditions, it also satisfies

1. $\widetilde{\mathbf{E}}_C[p \cdot f_i] = 0, \forall i \in [m]$ and $\forall p$ such that $\deg(p \cdot f_i) \leq q$.
2. $\widetilde{\mathbf{E}}_C\left[p^2 \cdot \prod_{i \in S} g_i\right] \geq 0, \forall S \subseteq [r]$ and $\forall p$ such that $\deg(p^2 \cdot \prod_{i \in S} g_i) \leq q$.

It is well-known that such constrained pseudoexpectation operators can be described as solutions to semidefinite programs of size $n^{O(q)}$ [BS14, Lau09]. This hierarchy of semidefinite programs for increasing q is known as the SoS hierarchy.

7.4 Matrix Representations of Polynomials and relaxations of $\|f\|_2$

For a homogeneous polynomial h of even degree q, a matrix $M_h \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ is called a matrix representation of h if $(x^{\otimes q/2})^T \cdot M_h \cdot x^{\otimes q/2} = h(x) \;\; \forall x \in \mathbb{R}^n$. The following well known fact will come in handy later.

Fact 7.4.1. *For polynomials p_1, p_2, let $p_1 \succeq p_2$ denote that $p_1 - p_2$ is a sum of squares. It is easy to verify that if p_1, p_2 are homogeneous degree d polynomials and there exist matrix representations M_{p_1} and M_{p_2} of p_1 and p_2 respectively, such that $M_{p_1} - M_{p_2} \succeq 0$, then $p_1 - p_2 \succeq 0$.*

Next we define a quantity that is closely related to our final relaxation:

$$\Lambda(h) \;:=\; \inf \left\{ \sup_{\|z\|_2 = 1} z^T M_h \, z \;\middle|\; M \text{ is a representation of } h \right\}. \tag{7.1}$$

Clearly, $h_{\max} \leq \Lambda(h)$, i.e. $\Lambda(h)$ is a relaxation of $h_{\max}$. However, this does not imply that $\Lambda(h)$ is a relaxation of $\|h\|_2$, since it can be the case that $h_{\max} \neq \|h\|_2$. To remedy this, one can instead consider $\sqrt{\Lambda(h^2)}$ which is a relaxation of $\|h\|_2$, since $(h^2)_{\max} = \|h^2\|_2$. More generally, for a degree-d homogeneous polynomial f and an integer q divisible by $2d$, we have the upper estimate

$$\|f\|_2 \;\leq \Lambda\left(f^{q/d}\right)^{d/q}$$

Let $\mathsf{M}_f \in \mathbb{R}^{n^{d/2} \times n^{d/2}}$ denote the unique SoS-symmetric matrix representation of f. Figure Fig. 7.1 gives the primal and dual forms of the relaxation computing $\Lambda(f)$. It is easy to check that strong duality holds in this case, since the solution $\widetilde{\mathbf{E}}_C[x^\alpha] = (1/\sqrt{n})^{|\alpha|}$ for all $\alpha \in \mathbb{N}^n_{\leq d}$, is strictly feasible and in the relative interior of the domain. Thus, the objective values of the two programs are equal.

Primal

$$\Lambda(f) \ := \ \inf \left\{ \sup_{\|z\|=1} z^T M z \ \middle|\ M \in \mathrm{Sym}\, n^{d/2},\ \ (x^{\otimes d/2})^T \cdot M \cdot x^{\otimes d/2} = f(x)\ \ \forall x \in \mathbb{R}^n \right\}$$

Dual I

$$\begin{array}{lr} \text{maximize} & \langle \mathsf{M}_f, \mathsf{X} \rangle \\ \text{subject to :} & \mathrm{Tr}(\mathsf{X}) = 1 \\ & \mathsf{X} \text{ is SoS symmetric} \\ & \mathsf{X} \succeq 0 \end{array}$$

Dual II

$$\begin{array}{lr} \text{maximize} & \widetilde{\mathbf{E}}_C[f] \\ \text{subject to :} & \widetilde{\mathbf{E}}_C \text{ is a degree-}d \\ & \text{pseudoexpectation} \\ & \widetilde{\mathbf{E}}_C \text{ respects } C \equiv \left\{ \|x\|_2^d = 1 \right\} \end{array}$$

Figure 7.1: Primal and dual forms for the relaxation computing $\Lambda(f)$

Primal

$$\|f\|_{sp} \ := \ \inf \left\{ \|M\|_2 \ \middle|\ M \in \mathrm{Sym}\, n^{d/2},\ \ (x^{\otimes d/2})^T \cdot M \cdot x^{\otimes d/2} = f(x)\ \ \forall x \in \mathbb{R}^n \right\}$$

Dual

$$\begin{array}{lr} \text{maximize} & \langle \mathsf{M}_f, \mathsf{X} \rangle \\ \text{subject to :} & \|\mathsf{X}\|_{S_1} = 1 \\ & \mathsf{X} \text{ is SoS symmetric} \end{array}$$

Figure 7.2: Primal and dual forms for the relaxation computing $\|f\|_{sp}$

We will also consider a weaker relaxation of $\|f\|_2$, which we denote by $\|f\|_{sp}$. A somewhat weaker version of this was used as the reference value in the work of [BKS14]. Figure Fig. 7.2 gives the primal and dual forms of this relaxation.

We will also need to consider constraint sets $C = \left\{\|x\|_2^2 = 1, x^{\beta_1} \geq 0, \ldots, x^{\beta_m} \geq 0\right\}$. We refer to the non-negativity constraints here as moment non-negativity constraints. When considering the maximum of $\widetilde{\mathbf{E}}_C[f]$, for constraint sets C containing moments non-negativity constraints in addition to $\|x\|_2^2 = 1$, we refer to the optimum value as $\Lambda_C(f)$. Note that the maximum is still taken over degree-d pseudoexpectations. Also, strong duality still holds in this case since $\widetilde{\mathbf{E}}_C[x^\alpha] = (1/\sqrt{n})^{|\alpha|}$ is still a strictly feasible solution.

7.4.1 Properties of relaxations obtained from constrained pseudoexpectations

We use the following claim, which is an easy consequence of the fact that the sum-of-squares algorithm can produce a certificate of optimality (see [OZ13]). In particular, if $\max_{\widetilde{\mathbf{E}}_C} \widetilde{\mathbf{E}}_C[f] = \Lambda_C(f)$ for a degree-q_1 pseudoexpectation operator respecting C containing $\|x\|_2^2 = 1$ and moment non-negativity constraints for $\beta_1, \ldots, \beta_m$, then for every $\lambda > \Lambda_C(f)$, we have that $\lambda - f$ can be certified to be positive by showing that $\lambda - f \in \Sigma_C^{q_1}$. Here $\Sigma_C^{(q_1)}$ is the set of all expressions of the form

$$\lambda - f \;=\; \sum_j p_j \cdot \left(\|x\|_2^2 - 1\right) + \sum_{S \subseteq [m]} h_S(x) \cdot \prod_{i \in S} x^{\beta_i} \,,$$

where each h_S is a sum of squares of polynomials and the degree of each term is at most q_1.

Lemma 7.4.2. *Let $\Lambda_C(f)$ denote the maximum of $\widetilde{\mathbf{E}}_C[f]$ over all degree-d pseudoexpectation operators respecting C. Then, for a pseudoexpectation operator of degree d' (respecting C) and a polynomial p of degree at most $(d' - d)/2$, we have that*

$$\widetilde{\mathbf{E}}_C\left[p^2 \cdot f\right] \;\leq\; \widetilde{\mathbf{E}}_C\left[p^2\right] \cdot \Lambda_C(f) \;.$$

Proof. As described above, for any $\lambda > \Lambda_C(f)$, we can write $\lambda - f = g$ for $g \in \Sigma_C^{(d)}$. Since the degree of each term in $p^2 \cdot g$ is at most d', we have by the properties of pseudoexpectation operators (of degree d') that

$$\lambda \cdot \widetilde{\mathbf{E}}_C\left[p^2\right] - \widetilde{\mathbf{E}}_C\left[p^2 \cdot f)\right] \;=\; \widetilde{\mathbf{E}}_C\left[p^2 \cdot (\lambda - f)\right] \;=\; \widetilde{\mathbf{E}}_C\left[p^2 \cdot g\right] \;\geq\; 0\,.$$

∎

The following monotonicity claim for non-negative coefficient polynomials will come in handy in later sections.

Lemma 7.4.3. *Let C be a system of polynomial constraints containing $\{\forall \beta \in \mathbb{N}_t^n, x^\beta \geq 0\}$. Then for any non-negative coefficient polynomials f and g of degree t, and such that $f \geq g$ (coefficient-wise, i.e. $f - g$ has non-negative coefficients), we have $\Lambda_C(f) \geq \Lambda_C(g)$.*

Proof. For any pseudo-expectation operator $\widetilde{\mathbf{E}}_C$ respecting C, we have $\widetilde{\mathbf{E}}_C[f - g] \geq 0$ because of the moment non-negativity constraints and by linearity.

So let $\widetilde{\mathbf{E}}_C$ be a pseudo-expectation operator realizing $\Lambda_C(g)$. Then we have,

$$\Lambda_C(f) \geq \widetilde{\mathbf{E}}_C[f] = \widetilde{\mathbf{E}}_C[g] + \widetilde{\mathbf{E}}_C[f - g] = \Lambda_C(g) + \widetilde{\mathbf{E}}_C[f - g] \geq 0.$$

∎

Chapter 8

Worst Case Polynomial Optimization over the Sphere

In this chapter, we study the problem of optimizing homogeneous polynomials over the unit sphere. Formally, given an n-variate degree-d homogeneous polynomial f, the goal is to compute

$$\|f\|_2 \;:=\; \sup_{\|x\|=1} |f(x)|\,. \tag{8.1}$$

We develop general techniques to design and analyze algorithms for polynomial optimization over the sphere. The sphere constraint is one of the simplest constraints for polynomial optimization and thus is a good testbed for techniques. Indeed, we believe these techniques will also be useful in understanding polynomial optimization for other constrained settings.

8.1 Algorithmic Results

The following result shows that $\Lambda\big(f^{q/d}\big)^{d/q}$ approximates $\|f\|_2$ within polynomial factors, and also gives an algorithm to approximate $\|f\|_2$ with respect to the upper bound $\Lambda\big(f^{q/d}\big)^{d/q}$. In the statements below and the rest of this section, $O_d(\cdot)$ and $\Omega_d(\cdot)$ notations hide $2^{O(d)}$ factors. Our algorithmic results are as follows:

Theorem 8.1.1. *Let f be an n-variate homogeneous polynomial of degree-d, and let $q \le n$ be an integer divisible by $2d$. Then,*

Arbitrary f:	$\left(\Lambda\left(f^{q/d}\right)\right)^{d/q} \;\le\; O_d\left((n/q)^{d/2-1}\right) \cdot \|f\|_2$
f with Non-neg. Coefficients:	$\left(\Lambda_C\left(f^{q/d}\right)\right)^{d/q} \;\le\; O_d\left((n/q)^{d/4-1/2}\right) \cdot \|f\|_2$
f with Sparsity m:	$\left(\Lambda\left(f^{q/d}\right)\right)^{d/q} \;\le\; O_d\left(\sqrt{m/q}\right) \cdot \|f\|_2.$

(where $\Lambda_C(\cdot)$ is a a related efficiently computable quantity that we define in Section 7.4)

Furthermore, there is a deterministic algorithm that runs in $n^{O(q)}$ time and returns x such that

$$|f(x)| \geq \frac{\Lambda\left(f^{q/d}\right)^{d/q}}{O_d(c(n,d,q))}$$

where $c(n,d,q)$ is $(n/q)^{d/2-1}$, $(n/q)^{d/4-1/2}$ and $\sqrt{m/q}$ respectively, for each of the above cases (the inequality uses $\Lambda_C(\cdot)$ in the case of polynomials with non-negative coefficients).

Remark 8.1.2. *Interestingly, our deterministic algorithms only involve computing the maximum eigenvectors of $n^{O(q)}$ different matrices in $\mathbb{R}^{n\times n}$, and actually don't require computing $\Lambda\left(f^{q/d}\right)^{d/q}$ (even though this quantity can also be computed in $n^{O(q)}$ time by the sum-of-squares SDP; see Section 8.2). The quantity $\Lambda\left(f^{q/d}\right)^{d/q}$ is only used in the analysis.*

Remark 8.1.3. *If $m = n^{\rho\cdot d}$ for $\rho < 1/3$, then for all $q \leq n^{1-\rho}$, the $\sqrt{m/q}$-approximation for sparse polynomials is better than the $(n/q)^{d/2-1}$ arbitrary polynomial approximation.*

Remark 8.1.4. *In cases where $\|f\|_2 = f_{\max}$ (such as when d is odd or f has non-negative coefficients), the above result holds whenever q is even and divisible by d, instead of $2d$.*

A key technical ingredient en route establishing the above results is a method to reduce the problem for arbitrary polynomials to a list of *multilinear* polynomial problems (over the same variable set). We believe this to be of independent interest, and describe its context and abstract its consequence (Theorem 8.1.5) next.

Let M_g be a matrix representation of a degree-q homogeneous polynomial g, and let $K = (I, J) \in [n]^{q/2} \times [n]^{q/2}$ have all distinct elements. Observe that there are $q!$ distinct entries of M_g including K across which, one can arbitrarily assign values and maintain the property of representing g, as long as the sum across all $q!$ entries remains the same (specifically, this is the set of all permutations of K). In general for $K' = (I', J') \in [n]^{q/2} \times [n]^{q/2}$, we define the orbit of K' denoted by $\mathcal{O}(K')$, as the set of permutations of K', i.e. the number of entries to which 'mass' from $M_g[I', J']$ can be moved while still representing g.

As q increases, the orbit sizes of the entries increase, and to show better bounds on $\Lambda\left(f^{q/d}\right)$, one must exploit these additional "degrees of freedom" in representations of $f^{q/d}$. However, a big obstacle is that the orbit sizes of different entries can range anywhere from 1 to $q!$, two extremal examples being $((1,\ldots 1),(1,\ldots 1))$ and $((1,\ldots q/2),(q/2+1,\ldots q))$. This makes it hard to exploit the additional freedom afforded by growing q. Observe that if g were multilinear, all matrix entries corresponding to non-zero coefficients have a span of $q!$ and indeed it turns out to be easier to analyze the approximation factor in the multilinear case as a function of q since the representations of g can be highly symmetrized. However, we are still faced with the problem of $f^{q/d}$ being highly non-multilinear. The natural symmetrization strategies that work well for multilinear polynomials fail on general polynomials, which motivates the following result:

Theorem 8.1.5 (Informal version of Theorem 8.7.13)**.** *For even q, let $g(x)$ be a degree-q homogeneous polynomial. Then there exist multilinear polynomials $g_1(x), \ldots, g_m(x)$ of degree at most q, such that*

$$\frac{\Lambda(g)}{\|g\|_2} \;\leq\; 2^{O(q)} \cdot \max_{i\in[m]} \frac{\Lambda(g_i)}{\|g_i\|_2}$$

By combining Theorem 8.1.5 (or an appropriate generalization) with the appropriate analysis of the multilinear polynomials induced by $f^{q/d}$, we obtain the aforementioned results for various classes of polynomials.

Weak decoupling lemmas. A common approach for reducing to the multilinear case is through more general "decoupling" or "polarization" lemmas (see e.g., Lemma 8.7.6), which also have variety of applications in functional analysis and probability [DlPG12]. However, such methods increase the number of variables to nq, which would completely nullify any advantage obtained from the increased degrees of freedom.

Our proof of Theorem 8.1.5 (and its generalizations) requires only a decoupling with somewhat weaker properties than given by the above lemmas. However, we need it to be very efficient in the number of variables. In analogy with "weak regularity lemmas" in combinatorics, which trade structural control for complexity of the approximating object, we call these results "weak decoupling lemmas" (see Section 8.6.1 and Lemma 8.7.12). They provide a milder form of decoupling but only increase the number of variables to $2n$ (independently of q).

We believe these could be more generally applicable; in particular to other constrained settings of polynomial optimization as well as in the design of sub-exponential algorithms. Our techniques might also be able to yield a full tradeoff between the number of variables and quality of decoupling.

8.2 Connection to Sum-of-Squares hierarchy

The *Sum of Squares Hierarchy* (SoS) is one of the the most canonical and well-studied approaches to attack polynomial optimization problems. Algorithms based on this framework are parameterized by the degree or level q of the SoS relaxation. For the case of optimization of a homogenous polynomial h of even degree q (with some matrix representation M_h) over the unit sphere, the level q SoS relaxes the non-convex program of maximizing $(x^{\otimes q/2})^T \cdot M_h \cdot x^{\otimes q/2} = h(x)$ over $x \in \mathbb{R}^n$ with $\|x\|_2 = 1$, to the semidefinite program of maximizing $\mathsf{Tr}(M_h^T X)$ over all positive semidefinite matrices $X \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ with $\mathsf{Tr}(X) = 1$. (This is a relaxation because $X = x^{\otimes q/2}(x^{\otimes q/2})^T$ is psd with $\mathsf{Tr}(X) = \|x\|_2^q$.)

It is well known (see for instance [Lau09]) that the quantity $\Lambda(h)$ from Eq. (7.1) is the dual value of this SoS relaxation. Further, strong duality holds for the case of optimization on the sphere and therefore $\Lambda(h)$ equals the optimum of the SoS SDP and can be computed in time $n^{O(q)}$. (See Section 7.3 for more detailed SoS preliminaries.) In light of this, our results from Theorem 8.1.1 can also be viewed as a convergence analysis of the SoS hierarchy for optimization over the sphere, as a function of the number of levels q. Such results are of significant interest in the optimization community, and have been studied for example in [DW12, dKLS14] (see Section 8.3 for a comparison of results).

SoS Lower Bounds. While the approximation factors in our upper bounds of Theorem 8.1.1 are modest, there is evidence to suggest that this is inherent.

When h is a degree-q polynomial with *random* i.i.d ± 1 coefficients, we will show in Chapter 9 that there is a constant c such that w.h.p. $\left(\frac{n}{q^{c+o(1)}}\right)^{q/4} \leq \Lambda(h) \leq \left(\frac{n}{q^{c-o(1)}}\right)^{q/4}$. On the other hand, $\|h\|_2 \leq O(\sqrt{nq\log q})$ w.h.p. Thus the ratio between $\Lambda(h)$ and $\|h\|_2$ can be as large as $\Omega_q(n^{q/4-1/2})$.

Hopkins et al. [HKP+17] recently proved that degree-d polynomials with random coefficients achieve a degree-q SoS gap of roughly $(n/q^{O(1)})^{d/4-1/2}$ (provided $q > n^{\varepsilon}$ for some constant $\varepsilon > 0$). This is also a lower bound on the ratio between $\Lambda\left(f^{q/d}\right)^{d/q}$ and $\|f\|_2$ for the case of *arbitrary* polynomials. Note that this lower bound is roughly square root of our upper bound from Theorem 8.1.1. Curiously, our upper bound for the case of polynomials with non-negative coefficients essentially matches this lower bound for random polynomials.

Non-Negative Coefficient Polynomials. In this chapter, we give a new lower bound construction for the case of non-negative polynomials, To the best of out knowledge, the only previous lower bound for this problem, was known through Nesterov's reduction [DK08], which only rules out a PTAS. We give the following polynomially large lower bound. The gap applies for random polynomials associated with a novel distribution of 4-uniform hypergraphs, and is analyzed using subgraph counts in a random graph.

Theorem 8.2.1. *There exists an n variate degree-4 homogeneous polynomial f with non-negative coefficients such that*

$$\|f\|_2 \;\leq\; (\log n)^{O(1)} \qquad \textit{and} \qquad \Lambda(f) \;\geq\; \tilde{\Omega}(n^{1/6})\,.$$

For larger degree t, we prove an $n^{\Omega(t)}$ gap between $\|h\|_2$ and a quantity $\|h\|_{sp}$ that is closely related to $\Lambda(h)$. Specifically, $\|h\|_{sp}$ is defined by replacing the largest eigenvalue of matrix representations M_h of h in Eq. (7.1) by the *spectral norm* $\|M_h\|_2$. (See Fig. 7.2 for a formal definition.) Note that $\|h\|_{sp} \geq \max\{\Lambda(h), \Lambda(-h)\}$. Like $\Lambda(\cdot)$, $\|\cdot\|_{sp}$ suggests a natural hierarchy of relaxations for the problem of approximating $\|h\|_2$, obtained by computing $\|h^{q/t}\|_{sp}^{t/q}$ as the q-th level of the hierarchy.

We prove a lower bound of $n^{q/24}/\,(q\cdot\log n)^{O(q)}$ on $\|f^{q/4}\|_{sp}$ where f is as in Section 8.9. This not only gives $\|\cdot\|_{sp}$ gaps for the degree-q optimization problem on polynomials with non-negative coefficients, but also an $n^{1/6}/(q\log n)^{O(1)}$ gap on higher levels of the aforementioned $\|\cdot\|_{sp}$ hierarchy for optimizing degree-4 polynomials with non-negative coefficients. Formally we show:

Theorem 8.2.2. *Let $g := f^{q/4}$ where f is the degree-4 polynomial as in Section 8.9. Then*

$$\frac{\|g\|_{sp}}{\|g\|_2} \;\geq\; \frac{n^{q/24}}{(q\log n)^{O(q)}}\,.$$

Our lower bound on $\|f^{q/4}\|_{sp}$ is based on a general tool that allows one to "lift" level-4 $\|\cdot\|_{sp}$ gaps, that meet one additional condition, to higher levels. While we derive final results only for the weaker relaxation $\|\cdot\|_{sp}$, the underlying structural result may prove useful in lifting SoS lower bounds (i.e. gaps for $\Lambda(\cdot)$) as well. Recently, the insightful pseudo-calibration approach [BHK+16] has provided a recipe to give higher level SoS lower bounds for certain *average-case* problems. We believe our lifting result might similarly be useful in the context of *worst-case* problems, where in order to get higher degree lower bounds, it suffices to give lower bounds for constant degree SoS with some additional structural properties.

8.3 Related Work

Polynomial optimization is a vast area with several previous results. Below, we collect the results most relevant for comparison with the ones in this chapter, grouped by the class of polynomials. Please see the excellent monographs [Lau09, Las09] for a survey.

Arbitrary Polynomials. For general homogeneous polynomials of degree-d, an $O_d\left(n^{d/2-1}\right)$ approximation was given by He et al. [HLZ10], which was improved to $O_d\left((n/\log n)^{d/2-1}\right)$ by So [So11]. The convergence of SDP hierarchies for polynomial optimization over the sphere was analyzed by Faybusovich [Fay04] and Doherty and Wehner [DW12] and subsequent to the publication of our work by deKlerk and Laurent [dKL19]. However, for $q \ll n$ levels, these convergence results would only imply a weaker bound of $(O(n)/q)^{d/2}$.

Thus, our results can be seen as giving a stronger interpolation between the polynomial time algorithms obtained by [HLZ10, So11] and the exponential time algorithms given by $\Omega(n)$ levels of SoS, although the bounds obtained by [Fay04, DW12, dKL19] are tighter (by a factor of $2^{O(d)}$) for $q = \Omega(n)$ levels.

For the case of arbitrary polynomials, we believe a tradeoff between running time and approximation quality similar to ours can also be obtained by considering the tradeoffs for the results of Brieden et al. [BGK+01] used by So [So11]. However, to the best of our knowledge, this is not published. In particular, So uses the techniques of Khot and Naor [KN08] to reduce degree-d polynomial optimization to $d-2$ instances of the problem of optimizing the ℓ_2 diameter of a convex body. This is solved by [BGK+01], who give an $O((n/k)^{1/2})$ approximation in time $2^k \cdot n^{O(1)}$. We believe this can be combined with proof of So, to yield a $O_d\left((n/q)^{d/2-1}\right)$ approximation in time 2^q. We note here that the method of Khot and Naor [KN08] cannot be improved further (up to polylog) for the case $d = 3$ (see Section 8.12). Our results for the case of arbitrary polynomials show that similar bounds can also be obtained by a very generic algorithm given by the SoS hierarchy. Moreover, the general techniques developed here are versatile and demonstrably applicable to various other cases (like polynomials with non-negative coefficients, sparse polynomials, worst-case sparse PCA) where no alternate proofs are available. The techniques of [KN08, So11] are oblivious to the structure in the polynomials and it appears to be unlikely that similar results can be obtained by using diameter estimation techniques.

Polynomials with Non-negative Coefficients. The case of polynomials with non-negative coefficients was considered by Barak, Kelner, and Steurer [BKS14] who proved that the relaxation obtained by $\Omega(d^3 \cdot \log n/\varepsilon^2)$ levels of the SoS hierarchy provides an $\varepsilon \cdot \|f\|_{BKS}$ additive approximation to the quantity $\|f\|_2$. Here, the parameter we denote by $\|f\|_{BKS}$

corresponds to a relaxation for $\|f\|_2$ that is weaker than the one given by $\|f\|_{sp}$.[1] Their results can be phrased as showing that a relaxation obtained by q levels of the SoS hierarchy gives an approximation ratio of

$$1 + \left(\frac{d^3 \cdot \log n}{q}\right)^{1/2} \cdot \frac{\|f\|_{BKS}}{\|f\|_2}.$$

Motivated by connections to quantum information theory, they were interested in the special case where $\|f\|_{BKS}/\|f\|_2$ is bounded by a constant. However, this result does not imply strong multiplicative approximations outside of this special case since in general $\|f\|_{BKS}$ and $\|f\|_2$ can be far apart. In particular, we are able to establish that there exist polynomials f with non-neg. coefficients such that $\|f\|_{BKS}/\|f\|_2 \geq n^{d/24}$. Moreover we conjecture that the worst-case gap between $\|f\|_{BKS}$ and $\|f\|_2$ for polynomials with non-neg. coefficients is as large as $\widetilde{\Omega}_d((n/d)^{d/4-1/2})$ (note that the conjectured $(n/d)^{d/4-1/2}$ gap for non-negative coefficient polynomials is realizable using arbitrary polynomials, i.e. we show in Chapter 9 that polynomials with i.i.d. ± 1 coefficients achieve this gap w.h.p.).

Our results show that q levels of SOS gives an $(n/q)^{d/4-1/2}$ approximation to $\|f\|_2$ which has a better dependence on q and consequently, converges to a constant factor approximation after $\Omega(n)$ levels.

2-to-4 norm. It was proved in [BKS14] that for any matrix A, q levels of the SoS hierarchy approximates $\|A\|_{2\to 4}^4$ within a factor of

$$1 + \left(\frac{\log n}{q}\right)^{1/2} \cdot \frac{\|A\|_{2\to 2}^2 \|A\|_{2\to\infty}^2}{\|A\|_{2\to 4}^4}.$$

Brandao and Harrow [BH15] also gave a nets based algorithm with runtime 2^q that achieves the same approximation as above. Here again, the cases of interest were those matrices for which $\|A\|_{2\to 2}^2 \|A\|_{2\to\infty}^2$ and $\|A\|_{2\to 4}^4$ are at most constant apart.

We would like to bring attention to an open problem in this line of work. It is not hard to show that for an $m \times n$ matrix A with i.i.d. Gaussian entries, $\|A\|_{2\to 2}^2 = \Theta(m+n)$, $\|A\|_{2\to\infty}^2 = \Theta(n)$, and $\|A\|_{2\to 4}^2 = \Theta(m+n^2)$ which implies the worst case approximation factor achieved above is $\Omega(n/\sqrt{q})$ when we take $m = \Omega(n^2)$.

Our result for arbitrary polynomials of degree-4, achieves an approximation factor of $O(n/q)$ after q levels of SoS which implies that the current best known approximation 2-to-4 norm is oblivious to the structure of the 2-to-4 polynomial and seems to suggest that this problem can be better understood for arbitrary tall matrices. For instance, can one get a $\sqrt{m}/q$ approximation for $(m \times n)$ matrices (note that [BH15] already implies a $\sqrt{m/q}$-approximation for all m, and our result implies a $\sqrt{m}/q$-approximation when $m = \Omega(n^2)$).

Random Polynomials. For the case when f is a degree-d homogeneous polynomial with i.i.d. random ± 1 coefficients the concurrent works [BGL16, RRS16] showed that

[1] Specifically, $\|f\|_{BKS}$ minimizes the spectral norm over a smaller set of matrix representations of f than $\|f\|_{sp}$ which allows all matrix representations.

degree-q SoS certifies an upper bound on $\|f\|_2$ that is with high probability at most $\widetilde{O}((n/q)^{d/4-1/2}) \cdot \|f\|_2$. Curiously, this matches our approximation guarantee for the case of *arbitrary* polynomials with non-negative coefficients. This problem was also studied for the case of sparse random polynomials in [RRS16] motivated by applications to refuting random CSPs.

8.4 Organization

We cover some preliminaries in Section 8.5 and provide an overview of our proofs and techniques in Section 8.6. Chapter 7 provides details of various relaxations used in this document, and their duals in terms of the Sum-of-Squares hierarchy. We first give a basic version of the reduction from general to multilinear polynomials in Section 8.7, which only obtains a weaker result (without the additive term in the exponent). Section 8.8 gives a generalization of this reduction, which yields Theorem 8.1.1. We prove an SoS lower bound for degree-4 polynomials with non-negative coefficients in Section 8.9. In Section 8.10, we provide a general technique for lifting lower bounds for the slightly weaker relaxation given by $\|f\|_{sp}$, to relaxation higher level relaxations.

8.5 Preliminaries and Notation

Definition 8.5.1 (Folded Polynomials)**.** *A degree-(d_1, d_2)* folded polynomial *$f \in \left(\mathbb{R}_{d_2}[x]\right)_{d_1}[x]$ is defined to be a polynomial of the form*

$$f(x) = \sum_{\alpha \in \mathbb{N}^n_{d_1}} \overline{f}_\alpha(x) \cdot x^\alpha ,$$

where each $\overline{f}_\alpha(x) \in \mathbb{R}_{d_2}[x]$ is a homogeneous polynomial of degree d_2. Folded polynomials over $\mathbb{R}^+$ are defined analogously.

- *We refer to the polynomials $\overline{f}_\alpha$ as the* folds *of f and the terms x^α as the monomials in f.*
- *A folded polynomial can also be used to define a degree $d_1 + d_2$ polynomial by multiplying the monomials with the folds (as polynomials in $\mathbb{R}[x]$). We refer to this polynomial in $\mathbb{R}_{d_1+d_2}[x]$ as the* unfolding *of f, and denote it by $\mathsf{U}(f)$.*
- *For a degree (d_1, d_2)-folded polynomial f and $r \in \mathbb{N}$, we take f^r to be a degree-$(r \cdot d_1, r \cdot d_2)$ folded polynomial, obtained by multiplying the folds as coefficients.*

An additional operation on folded polynomials. We define the following operation (and it's folded counterpart) which in the case of a multilinear polynomial corresponds (up to scaling) to the sum of a row of the SoS-symmetric matrix representation of the polynomial. This will be useful in our result for non-negative polynomials.

Definition 8.5.2 (Collapse). *Let $f \in \mathbb{R}_d[x]$ be a polynomial. The k-*collapse *of f, denoted as $\mathsf{C}_k(f)$ is the degree $d-k$ polynomial g given by,*

$$g(x) = \sum_{\gamma \in \mathbb{N}^n_{d-k}} g_\gamma \cdot x^\gamma \quad \textit{where} \quad g_\gamma = \sum_{\alpha \in \mathbb{N}^n_k} f_{\gamma+\alpha} \,.$$

For a degree-(d_1, d_2) folded polynomial f, we define $\mathsf{C}_k(f)$ similarly as the degree-$(d_1 - k, d_2)$ folded polynomial g given by,

$$g(x) = \sum_{\gamma \in \mathbb{N}^n_{d_1-k}} \bar{g}_\gamma(x) \cdot x^\gamma \quad \textit{where} \quad \bar{g}_\gamma = \sum_{\alpha \in \mathbb{N}^n_k} \bar{f}_{\gamma+\alpha} \,.$$

8.6 Overview of Proofs and Techniques

In the interest of clarity, we shall present all techniques for the special case where f is an arbitrary degree-4 homogeneous polynomial. We shall further assume that $\|f\|_2 = f_{\max}$ just so that $\Lambda(f)$ is a relaxation of $\|f\|_2$. Summarily, the goal of this section is to give an overview of an $O(n/q)$-approximation of $\|f\|_2$, i.e.

$$\Lambda\left(f^{q/4}\right)^{4/q} \leq O(n/q) \cdot \|f\|_2.$$

Many of the high level ideas remain the same when considering higher degree polynomials and special classes like polynomials with non-negative coefficients, or sparse polynomials.

8.6.1 Warmup: (n^2/q^2)-Approximation

We begin with seeing how to analyze constant levels of the $\Lambda(\cdot)$ relaxation and will then move onto higher levels in the next section. The level-4 relaxation actually achieves an n-approximation, however we will start with n^2 as a warmup and cover the n-approximation a few sections later.

n^2-Approximation using level-4 relaxation

We shall establish that $\Lambda(f) \leq O(n^2) \cdot \|f\|_2$. Let M_f be the SoS-symmetric representation of f, let $x_{i_1} x_{i_2} x_{i_3} x_{i_4}$ be the monomial whose coefficient in f has the maximum magnitude, and let B be the magnitude of this coefficient. Now by Gershgorin circle theorem, we have $\Lambda(f) \leq \|M_f\|_2 \leq n^2 \cdot B$.

It remains to establish $\|f\|_2 = \Omega(B)$. To this end, define the decoupled polynomial $\mathcal{F}(x,y,z,t) := (x \otimes y)^T \cdot M_f \cdot (z \otimes t)$ and define the decoupled two-norm as

$$\|\mathcal{F}\|_2 := \sup_{\|x\|,\|y\|,\|z\|,\|t\|=1} \mathcal{F}(x,y,z,t).$$

It is well known that $\|f\|_2 = \Theta(\|\mathcal{F}\|_2)$ (see Lemma 8.6.1). Thus, we have,

$$\|f\|_2 \;=\; \Omega(\|\mathcal{F}\|_2) \;\geq\; \Omega\left(|\mathcal{F}(e_{i_1}, e_{i_2}, e_{i_3}, e_{i_4})|\right) \;=\; \Omega(B) \;=\; \Omega\left(\Lambda(f)\,/n^2\right).$$

In order to better analyze $\Lambda\left(f^{q/4}\right)^{4/q}$ we will need to introduce some new techniques.

(n^2/q^2)-Approximation Assuming Theorem 8.1.5

We will next show that $\Lambda\left(f^{q/4}\right)^{4/q} \leq O(n^2/q^2) \cdot \|f\|_2$ (for q divisible by 4). In fact, one can show something stronger, namely that for every homogeneous polynomial g of degree-q, $\Lambda(g) \leq 2^{O(q)} \cdot (n/q)^{q/2} \cdot \|g\|_2$ which clearly implies the above claim (also note that for the target $O(n^2/q^2)$-approximation to $\|f\|_2$, losses of $2^{O(q)}$ in the estimate of $\|g\|_2$ are negligible, while factors of the order $q^{\Omega(q)}$ are crucial).

Given the additional freedom in choice of representation (due to the polynomial having higher degree), a first instinct would be to completely symmetrize, i.e. take the SoS-symmetric representation of g, and indeed this works for multilinear g (see Theorem 8.7.16 for details).

However, the above approach of taking the SoS-symmetric representation breaks down when the polynomial is non-multilinear. To circumvent this issue, we employ Theorem 8.1.5 which on combining with the aforementioned multilinear polynomial result, yields that for every homogeneous polynomial g of degree-q, $\Lambda(g) \leq (n/q)^{q/2} \cdot \|g\|_2$. The proofs of Theorem 8.1.5 and it's generalizations (that will be required for the n/q approximation), are quite non-trivial and are the most technically involved sections of our upper bound results. We shall next give an outline of the proof of Theorem 8.1.5.

Reduction to Optimization of Multi-linear Polynomials

One of the main techniques we develop in this work, is a way of reducing the optimization problem for general polynomials to that of multi-linear polynomials, which *does not increase the number of variables*. While general techniques for reduction to the multi-linear case have been widely used in the literature [KN08, HLZ10, So11] (known commonly as decoupling/polarization techniques), these reduce the problem to optimizing a multi-linear polynomial in $n \cdot d$ variables (when the given polynomial h is of degree d). Below is one example:

Lemma 8.6.1 ([HLZ10]). *Let $\mathcal{A}$ be a SoS-symmetric d-tensor and let $h(x) := \langle \mathcal{A}, x^{\otimes d}\rangle$. Then $\|h\|_2 \;\geq\; 2^{-O(d)} \cdot \max_{\|x^i\|=1} \langle \mathcal{A}, x^1 \otimes \cdots \otimes x^d\rangle$.*

Since we are interested in the improvement in approximation obtained by considering $f^{q/4}$ for a large q, applying these would yield a multi-linear polynomial in $n \cdot q$ variables. For our analysis, this increase in variables exactly cancels the advantage we obtain by considering $f^{q/4}$ instead of f (i.e., the advantage obtained by using q levels of the SoS hierarchy).

We can uniquely represent a homogeneous polynomial g of degree q as

$$g(x) \;=\; \sum_{|\alpha|\le q/2} x^{2\alpha}\cdot G_{2\alpha}(x) \;=\; \sum_{r=0}^{q/2}\sum_{|\alpha|=r} x^{2\alpha}\cdot G_{2\alpha}(x) \;=\; \sum_{r=0}^{q/2} g_r(x)\,, \tag{8.2}$$

where each $G_{2\alpha}$ is a multi-linear polynomial and $g_r(x) := \sum_{|\alpha|=r} x^{2\alpha}\cdot G_{2\alpha}(x)$. We reduce the problem to optimizing $\|G_{2\alpha}\|_2$ for each of the polynomials $G_{2\alpha}$. More formally, we show that

$$\frac{\Lambda(g)}{\|g\|_2} \;\le\; \max_{\alpha\in\mathbb{N}^n_{\le q/2}} \frac{\Lambda(G_{2\alpha})}{\|G_{2\alpha}\|_2}\cdot 2^{O(q)} \tag{8.3}$$

As a simple and immediate example of its applicability, Eq. (8.3) provides a simple proof of a polytime constant factor approximation for optimization over the simplex (actually this case is known to admit a PTAS [dKLP06, dKLS15]). Indeed, observe that a simplex optimization problem for a degree-$q/2$ polynomial in the variable vector y can be reduced to a sphere optimization by substituting $y_i = x_i^2$. Now since every variable present in a monomial has even degree in that monomial, each $G_{2\alpha}$ is constant, which implies a constant factor approximation (dependent on q) on applying Eq. (8.3).

Returning to our overview of the proof, note that given representations of each of the polynomials $G_{2\alpha}$, each of the polynomials g_r can be represented as a block-diagonal matrix with one block corresponding to each α. Combining this with triangle inequality and the fact that the maximum eigenvalue of a block-diagonal matrix is equal to the maximum eigenvalue of one of the blocks, gives the following inequality:

$$\Lambda(g) \;\le\; (1+q/2)\cdot \max_{\alpha\in\mathbb{N}^n_{\le q/2}} \Lambda(G_{2\alpha})\,. \tag{8.4}$$

We can further strengthen Eq. (8.4) by averaging the "best" representation of $G_{2\alpha}$ over $|\mathcal{O}(\alpha)|$ diagonal-blocks which all correspond to $x^{2\alpha}$. This is the content of Lemma 8.7.2 wherein we show

$$\Lambda(g) \;\le\; (1+q/2)\cdot \max_{\alpha\in\mathbb{N}^n_{\le q/2}} \frac{\Lambda(G_{2\alpha})}{|\mathcal{O}(\alpha)|}\,. \tag{8.5}$$

Since $|\mathcal{O}(\alpha)|$ can be as large as $q^{\Omega(q)}$, the above strengthening is crucial. We then prove the following inequality, which shows that the decomposition in Eq. (8.2) not only gives a block-diagonal decomposition for matrix representations of g, but can in fact be thought of as a "block-decomposition" of the *tensor* corresponding to g (with regards to computing $\|g\|_2$). Just as the maximum eigenvalue of a block-diagonal matrix is at least the maximum eigenvalue of a block, we show that

$$\|g\|_2 \;\ge\; 2^{-O(q)}\cdot \max_{\alpha\in\mathbb{N}^n_{\le q/2}} \frac{\|G_{2\alpha}\|_2}{|\mathcal{O}(\alpha)|}\,. \tag{8.6}$$

The above inequality together with Eq. (8.5), implies Eq. (8.3).

Bounding $\|g\|_2$ via a new weak decoupling lemma

Recall that the expansion of $g(x)$ in Eq. (8.2), contains the term $x^{2\alpha} \cdot G_{2\alpha}(x)$. The key part of proving the bound in Eq. (8.6) is to show the following "weak decoupling" result for $x^{2\alpha}$ and $G_{2\alpha}$.

$$\forall \alpha \qquad \|g\|_2 \;\geq\; \max_{\|y\|=\|x\|=1} y^{2\alpha} \cdot G_{2\alpha}(x) \cdot 2^{-O(q)} \;=\; \max_{\|y\|=1} y^{2\alpha} \cdot \|G_{2\alpha}\|_2 \cdot 2^{-O(q)}.$$

The proof of Eq. (8.6) can then be completed by considering the unit vector $y := \sqrt{\alpha}/\sqrt{|\alpha|}$, i.e. $y := \sum_{i\in[n]} \frac{\sqrt{\alpha_i}}{\sqrt{|\alpha|}} \cdot e_i$. A careful calculation shows that $y^{2\alpha} \;\geq\; 2^{-O(q)}/\,|\mathcal{O}(\alpha)|$ which finishes the proof.

The primary difficulty in establishing the above decoupling is the possibility of cancellations. To see this, let x^* be the vector realizing $\|G_{2\alpha}\|_2$ and substitute $z = (x^* + y)$ into g. Clearly, $y^{2\alpha} \cdot G_{2\alpha}(x^*)$ is a term in the expansion of $g(z)$, however there is no guarantee that the other terms in the expansion don't cancel out this value. To fix this our proof relies on multiple delicate applications of the first-moment method, i.e. we consider a complex vector random variable $Z(x^*, y)$ that is a function of x^* and y, and argue about $\mathbb{E}[|g(Z)|]$.

The base case of $\alpha = 0^n$. We first consider the base case with $\alpha = 0^n$, where we define $y^{2\alpha} = 1$. This amounts to showing that for every homogeneous polynomial h of degree t, $\|h\|_2 \geq \|h_m\|_2 \cdot 2^{-O(t)}$ where h_m is the restriction of h to it's multilinear monomials.

Given the optimizer x^* of $\|h_m\|_2$, let z be a random vector such that each $Z_i = x_i^*$ with probability p and $Z_i = 0$ otherwise. Then, $\mathbb{E}[h(Z)]$ is a *univariate* degree-t polynomial in p with the coefficient of p^t equal to $h_m(x^*)$. An application of Chebyshev's extremal polynomial inequality (Lemma 8.7.5) then gives that there exists a value of the probability p such that

$$\|h\|_2 \;\geq\; \mathbb{E}\left[|h(Z)|\right] \;\geq\; \left|\mathbb{E}\left[h(Z)\right]\right| \;\geq\; 2^{-O(t)} \cdot |h_m(x^*)| \;=\; 2^{-O(t)} \cdot \|h_m\|_2\,.$$

For the case of general α, we first pass to the *complex version* of $\|g\|_2$ defined as

$$\|g\|_2^c \;:=\; \sup_{z\in\mathbb{C}^n, \|z\|=1} |g(z)|\;.$$

We use another averaging argument together with an application of the polarization lemma (Lemma 8.6.1) to show that we do not loose much by considering $\|g\|_2^c$. In particular, $\|g\|_2 \;\leq\; \|g\|_2^c \;\leq\; 2^{O(q)} \cdot \|g\|_2$.

The case of $g = g_r$. In this case, the problem reduces to showing that for all $\alpha \in \mathbb{N}_r^n$ and for all $y \in \mathbb{S}^{n-1}$,

$$\|g_r\|_2^c \;\;\geq\;\; y^{2\alpha} \cdot \|G_{2\alpha}\|_2 \cdot 2^{-O(q)}.$$

Fix any $\alpha \in \mathbb{N}_r^n$, and let $\omega \in \mathbb{C}^n$ be a complex vector random variable, such that ω_i is an independent and uniformly random $(2\alpha_i + 1)$-th root of unity. Let Ξ be a random $(q - 2r + 1)$-th root of unity, and let x^* be the optimizer of $\|G_{2\alpha}\|_2$. Let $Z := \omega \circ y + \Xi \cdot x^*$, where $\omega \circ y$ denotes the coordinate-wise product. Observe that for any α', γ such that $|\alpha'| = r,\ |\gamma| = q - 2r,\ \gamma \leq \mathbf{1}$,

$$\mathbb{E}\left[\prod_i \omega_i \cdot \Xi \cdot Z^{2\alpha'+\gamma}\right] \;=\; \begin{cases} y^{2\alpha} \cdot (x^*)^\gamma & \text{if } \alpha' = \alpha \\ 0 & \text{otherwise} \end{cases}$$

By linearity, this implies $\mathbb{E}[\prod_i \omega_i \cdot \Xi \cdot g_r(Z)] = y^{2\alpha} \cdot G_{2\alpha}(x^*)$. The claim then follows by noting that

$$\|g_r\|_2^c \;\geq\; \mathbb{E}\left[|g_r(Z)|\right] = \;\mathbb{E}\left[\left|\prod_i \omega_i \cdot \Xi \cdot g_r(Z)\right|\right] \;\geq\; \left|\mathbb{E}\left[\prod_i \omega_i \cdot \Xi \cdot g_r(Z)\right]\right| \;\geq\; y^{2\alpha} \cdot \|G_{2\alpha}\|_2.$$

The general case. The two special cases considered here correspond to the cases when we need to extract a specific g_r (for $r = 0$), and when we need to extract a fixed α from a given g_r. The argument for the general case uses a combination of the arguments for both these cases. Moreover, to get an $O(n/q)$ approximation, we also need versions of such decoupling lemmas for folded polynomials to take advantage of “easy substructures” as described next.

8.6.2 Exploiting Easy Substructures via Folding and Improved Approximations

To obtain the desired n/q-approximation to $\|f\|_2$, we need to use the fact that the problem of optimizing quadratic polynomials can be solved in polynomial time, and moreover that SoS captures this. More generally, in this section we consider the problem of getting improved approximations when a polynomial contains "easy substructures". It is not hard to obtain improved guarantees when considering constant levels of SoS. The second main technical contribution of our work is in giving sufficient conditions under which higher levels of SoS improve on the approximation of constant levels of SoS, when considering the optimization problem over polynomials containing "easy substructures".

As a warmup, we shall begin with seeing how to exploit easy substructures at constant levels by considering the example of degree-4 polynomials that trivially "contain" quadratics.

n-Approximation using Degree-4 SoS

Given a degree-4 homogeneous polynomial f (assume f is multilinear for simplicity), we consider a degree-$(2,2)$ folded polynomial h, whose unfolding yields f, chosen so that $\max_{\|y\|=1} \|h(y)\|_2 = \Theta(\|f\|_2)$ (recall that an evaluation of a folded polynomial returns a polynomial, i.e., for a fixed y, $h(y)$ is a quadratic polynomial in the indeterminate x). Such an h always exists and is not hard to find based on the SoS-symmetric representation of f. Also recall,

$$h(x) = \sum_{|\beta|=2,\, \beta \leq \mathbf{1}} \bar{h}_\beta(x) \cdot x^\beta\,,$$

where each $\bar{h}_\beta$ is a quadratic polynomial (the aforementioned phrase "easy substructures" is referencing the folds: $\bar{h}_\beta$ which are easy to optimize). Now by assumption we have,

$$\|f\|_2 \geq \max_{|\beta|=2,\, \beta\leq\mathbf{1}} \|h(\beta/\sqrt{2})\|_2 = \max_{|\beta|=2,\, \beta\leq\mathbf{1}} \|\bar{h}_\beta\|_2/2.$$

We then apply the block-matrix generalization of Gershgorin circle theorem to the SoS-symmetric matrix representation of f to show that

$$\Lambda(f) \leq \|f\|_{sp} \leq n \cdot \max_{|\beta|=2,\, \beta\leq\mathbf{1}} \|\bar{h}_\beta\|_{sp} \;=\; n \cdot \max_{|\beta|=2,\, \beta\leq\mathbf{1}} \|\bar{h}_\beta\|_2\,,$$

where the last step uses the fact that $\bar{h}_\beta$ is a quadratic, and $\|\cdot\|_{sp}$ is a tight relaxation of $\|\cdot\|_2$ for quadratics. This yields the desired n-approximation.

n/q-approximation using Degree-q SoS

Following the cue of the n^2/q^2-approximation, we derive the desired n/q bound by proving a folded-polynomial analogue of every claim in the previous section (including the multilinear reduction tools), a notable difference being that when we consider a power $f^{q/4}$ of f, we need to consider degree-$(q-2q/4, 2q/4)$ folded polynomials, since we want to use the fact that any ***product of*** $q/4$ ***quadratic polynomials*** is "easy" for SoS (in contrast to Section 8.6.2 where we only used the fact quadratic polynomials are easy for SoS). We now state an abstraction of the general approach we use to leverage the tractability of the folds.

Conditions for Exploiting "Easy Substructures" at Higher Levels of SoS. Let $d := d_1 + d_2$ and $f := \mathsf{U}(h)$ where h is a degree-(d_1, d_2) folded polynomial that satisfies

$$\sup_{\|y\|=1} \|h(y)\|_2 = \Theta_d(\|f\|_2)\,.$$

Implicit in Section 8.8, is the following theorem we believe to be of independent interest:

Theorem 8.6.2. *Let h, f be as above, and let*

$$\Gamma := \;\min\left\{ \frac{\Lambda(p)}{\|p\|_2} \;\middle|\; p(x) \in span\Big(\bar{h}_\beta \;\Big|\; \beta \in \mathbb{N}^n_{d_2}\Big) \right\}.$$

Then for any q divisible by $2d$, $\Lambda\big(f^{q/d}\big)^{d/q} \leq \; O_d\big(\Gamma \cdot (n/q)^{d_1/2}\big) \cdot \|f\|_2$.

In other words, if degree-d_2 SoS gives a good approximation for every polynomial in the subspace spanned by the folds of h, then higher levels of SoS give an improving approximation that exploits this. In this work, we only apply the above with $\Gamma = 1$, where exact optimization is easy for the space spanned by the folds.

While we focus on general polynomials for the overview, let us remark that in the case of polynomials with non-negative coefficients, the approximation factor in Theorem 8.6.2 becomes $O_d\big(\delta \cdot (n/q)^{d_1/4}\big)$.

8.6.3 Lower Bounds for Polynomials with Non-negative Coefficients

Degree-4 Lower Bound for Polynomials with Non-Negative Coefficients

We discuss some of the important ideas from the proof of Section 8.9. The lower bound in Chapter 9 proves $\frac{\Lambda(f)}{\|f\|_2}$ is large by considering a random polynomial f where each coefficient of f is an independent (Gaussian) random variable with bounded variance. The most natural adaptation of the above strategy to degree-4 polynomials with non-negative coefficients is to consider a random polynomial f where each coefficient f_α is independently sampled such that $f_\alpha = 1$ with probability p and $f_\alpha = 0$ with probability $1 - p$. However, this construction fails for every choice of p. If we let $\mathsf{A} \in \mathbb{R}^{[n]^2 \times [n]^2}$ be the natural matrix representation of f (i.e., each coefficient f_α is distributed uniformly among the corresponding entries of A), the Perron-Frobenius theorem shows that $\|\mathsf{A}\|_2$ is less than the maximum row sum $\max(\tilde{O}(n^2 p), 1)$ of M, which is also an upper bound on $\Lambda(f)$. However, we can match this bound by (within constant factors) choosing $x = (\frac{1}{\sqrt{n}}, \ldots, \frac{1}{\sqrt{n}})$ when $p \geq 1/n^2$. Also, when $p < 1/n^2$, we can take any α with $f_\alpha = 1$ and set $x_i = 1/2$ for all i with $\alpha_i > 0$, which achieves a value of $1/16$.

We introduce another natural distribution of random non-negative polynomials that bypasses this problem. Let $G = (V, E)$ be a random graph drawn from the distribution $G_{n,p}$ (where we choose $p = n^{-1/3}$ and $V = [n]$. Let $\mathcal{C} \subseteq \binom{V}{4}$ be the set of 4-cliques in G. The polynomial f is defined as

$$f(x_1, \ldots, x_n) := \sum_{\{i_1,i_2,i_3,i_4\} \in \mathcal{C}} x_{i_1} x_{i_2} x_{i_3} x_{i_4}.$$

Instead of trying $\Theta(n^4)$ p-biased random bits, we use $\Theta(n^2)$ of them. This limited independence bypasses the problem above, since the rows of A now have significantly different row sums: $\Theta(n^2 p)$ rows that correspond to an edge of G will have row sum $\Theta(n^2 p^5)$, and all other rows will be zeros. Since these $n^2 p$ rows (edges) are chosen independently from $\binom{[n]}{2}$, they still reveal little information that can be exploited to find a n-dimensional vector x with large $f(x)$. However, the proof requires a careful analysis of the trace method (to bound the spectral norm of an "error" matrix).

It is simple to prove that $\|f\|_{sp} \geq \Omega\left(\sqrt{n^2 p^5}\right) = \Omega(n^{1/6})$ by considering the Frobenius norm of the $n^2 p \times n^2 p$ principal submatrix, over any matrix representation (indeed, A is the minimizer). To prove $\Lambda(f) \geq \tilde{\Omega}(n^{1/6})$, we construct a moment matrix M that is SoS-symmetric, positive semidefinite, and has a large $\langle \mathsf{M}, \mathsf{A} \rangle$ (see the dual form of $\Lambda(f)$ in Chapter 7). It turns out that the $n^2 p \times n^2 p$ submatrix of A shares spectral properties of the adjacency matrix of a random graph $G_{n^2 p, p^4}$, and taking $\mathsf{M} := c_1 \mathsf{A} + c_2 \mathsf{I}$ for some identity-like matrix I proves $\Lambda(f) \geq \tilde{\Omega}(n^{1/6})$. An application of the trace method is needed to bound c_2.

To upper bound $\|f\|_2$, we first observe that $\|f\|_2$ is the same as the following natural combinatorial problem up to an $O(\log^4 n)$ factor: find four sets $S_1, S_2, S_3, S_4 \subseteq V$ that maximize

$$\frac{|\mathcal{C}_G(S_1,S_2,S_3,S_4)|}{\sqrt{|S_1||S_2||S_3||S_4|}}$$

where $|\mathcal{C}_G(S_1,S_2,S_3,S_4)|$ is the number of 4-cliques $\{v_1,\ldots,v_4\}$ of G with $v_i \in S_i$ for $i = 1,\ldots,4$. The number of copies of a fixed subgraph H in $G_{n,p}$, including its tail bound, has been actively studied in probabilistic combinatorics [Vu01, KV04, JOR04, Cha12, DK12a, DK12b, LZ16], though we are interested in bounding the 4-clique density of *every* 4-tuple of subsets simultaneously. The previous results give a strong enough tail bound for a union bound, to prove that the optimal value of the problem is $O(n^2p^6 \cdot \log^{O(1)} n)$ when $|S_1| = \cdots = |S_4|$, but this strategy inherently does not work when the set sizes become significantly different. However, we give a different analysis for the above asymmetric case, showing that the optimum is still no more than $O(n^2p^6 \cdot \log^{O(1)} n)$.

Lifting Stable Degree-4 Lower Bounds

For a degree-t (t even) homogeneous polynomial f, note that $\max\{|\Lambda(f)|, |\Lambda(-f)|\}$ is a relaxation of $\|f\|_2$. $\|f\|_{sp}$ is a slightly weaker (but still quite natural) relaxation of $\|f\|_2$ given by

$$\|f\|_{sp} := \inf\{\|M\|_2 \mid M \text{ is a matrix representation of } f\}.$$

As in the case of $\Lambda(f)$, for a (say) degree-4 polynomial f, $\|f^{q/4}\|_{sp}^{4/q}$ gives a hierarchy of relaxations for $\|f\|_2$, for increasing values of q.

We give an overview of a general method of "lifting" certain "stable" low degree gaps for $\|\cdot\|_{sp}$ to gaps for higher levels with at most $q^{O(1)}$ loss in the gap. While we state our techniques for lifting degree-4 gaps, all the ideas are readily generalized to higher levels. We start with the observation that the dual of $\|f\|_{sp}$ is given by the following "nuclear norm" program. Here M_f the canonical matrix representation of f, and $\|\mathsf{X}\|_{S_1}$ is the Schatten 1-norm (nuclear norm) of X, which is the sum of it's singular values.

$$\begin{array}{ll} \text{maximize} & \langle \mathsf{M}_f, \mathsf{X}\rangle \\ \text{subject to :} & \|\mathsf{X}\|_{S_1} = 1 \\ & \mathsf{X} \text{ is SoS symmetric} \end{array}$$

Now let X be a solution realizing a gap of δ between $\|f\|_{sp}$ and $\|f\|_2$. We shall next see how assuming reasonable conditions on X and M_f, one can show that $\|f^{q/4}\|_{sp}/\|f^{q/4}\|_2$ is at least $\delta^{q/4}/q^{O(q)}$.

In order to give a gap for the program corresponding to $\|f^{q/4}\|_{sp}$, a natural choice for a solution is the symmetrized version of the matrix $\mathsf{X}^{\otimes q/4}$ normalized by its Schatten-1 norm i.e., for $Y = \mathsf{X}^{\otimes q/4}$, we take

$$\mathsf{Z} := \frac{Y^S}{\|Y^S\|_{S_1}} \qquad \text{where} \qquad Y^S[K] = \mathop{\mathbb{E}}_{\pi \in \mathbb{S}_q}[Y[\pi(K)]] \quad \forall K \in [n]^q.$$

To get a good gap, we are now left with showing that $\|Y^S\|_{S_1}$ is not too large. Note that symmetrization can drastically change the spectrum of a matrix as for different permutations π, the matrices $Y^\pi[K] := Y[\pi(K)]$ can have very different ranks (while $\|Y\|_F = \|Y^\pi\|_F$). In particular, symmetrization can increase the maximum eigenvalue of a matrix by polynomial factors, and thus one must carefully count the number of such large eigenvalues in order to get a good upper bound on $\|Y^S\|_{S_1}$. Such an upper bound is a consequence of a structural result about Y^S that we believe to be of independent interest.

To state the result, we will first need some notation. For a matrix $M \in \mathbb{R}^{[n]^2 \times [n]^2}$ let $T \in \mathbb{R}^{[n]^4}$ denote the tensor given by, $T[i_1, i_2, i_3, i_4] = M[(i_1, i_2), (i_3, i_4)]$. Also for any non-negative integers x, y satisfying $x + y = 4$, let $M_{x,y} \in \mathbb{R}^{[n]^x \times [n]^y}$ denote the (rectangular) matrix given by, $M[(i_1, \ldots, i_x), (j_1, \ldots j_y)] = T[i_1, \ldots, i_x, j_1, \ldots j_y]$. Let $M \in \mathbb{R}^{[n]^2 \times [n]^2}$ be a degree-4 SoS-Symmetric matrix, let $M_A := M_{1,3} \otimes M_{4,0} \otimes M_{1,3}$, let $M_B := M_{1,3} \otimes M_{3,1}$, let $M_C := M$ and let $M_D := \mathsf{Vec}(M)\,\mathsf{Vec}(M)^T = M_{0,4} \otimes M_{4,0}$.

We show (see Theorem 8.10.4) that $(M^{\otimes q/4})^S$ can be written as the sum of $2^{O(q)}$ terms of the form:

$$C(a, b, c, d) \cdot P \cdot (M_A^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d}) \cdot P$$

where $12a + 8b + 4c + 8d = q$, P is a matrix with spectral norm 1 and $C(a, b, c, d) = 2^{O(q)}$. This implies that controlling the spectrum of M_A, M_B, M and M_D is sufficient to control on the spectrum of $(M^{\otimes q/4})^S$.

Using this result with $M := \mathsf{X}$, we are able to establish that if X satisfies the additional condition of $\|\mathsf{X}_{1,3}\|_{S_1} \leq 1$ (note that we already know $\|X\|_{S_1} \leq 1$), then $\|Y^S\|_{S_1} = 2^{O(q)}$. Thus Z realizes a $\langle M_f^{\otimes q/4}, Y^S \rangle / 2^{O(q)}$ gap for $\|f^{q/4}\|_{sp}$. On composing this result with the degree-4 gap from the previous section, we obtain an $\|\cdot\|_{sp}$ gap of $n^{q/24} / (q \cdot \log n)^{O(q)}$ for degree-q polynomials with non-neg. coefficients. We also show the q-th level $\|\cdot\|_{sp}$ gap for degree-4 polynomials with non-neg. coefficients is $\tilde{\Omega}(n^{1/6}) / q^{O(1)}$.

There are by now quite a few results giving near-tight lower bounds on the performance of higher level SoS relaxations for *average-case* problems [BHK+16, KMOW17, HKP+17]. However, there are few examples in the literature of matching SoS upper/lower bounds on *worst-case* problems. We believe our lifting result might be especially useful in such contexts, where in order to get higher degree lower bounds, it suffices to give stable lower bounds for constant degree SoS.

8.7 Results for Polynomials in $\mathbb{R}_d[x]$ and $\mathbb{R}_d^+[x]$

8.7.1 Reduction to Multilinear Polynomials

Lemma 8.7.1. *Any homogeneous n-variate degree-d polynomial $f(x)$ has a unique representation of the form*

$$\sum_{\alpha \in \mathbb{N}^n_{\leq d/2}} F_{2\alpha}(x) \cdot x^{2\alpha}$$

where for any $\alpha \in \mathbb{N}^n_{\leq d/2}$, $F_{2\alpha}$ is a homogeneous multilinear degree-$(d-2|\alpha|)$ polynomial.

We would like to approximate $\|f\|_2$ by individually approximating $\|F_{2\alpha}\|_2$ for each multilinear polynomial $F_{2\alpha}$. This section will establish the soundness of this goal.

Upper Bounding $\Lambda(f)$ in terms of $\Lambda(F_{2\alpha})$

We first bound $\Lambda(f)$ in terms of $\max_{\alpha \in \mathbb{N}^n_{\leq d/2}} \Lambda(F_{2\alpha})$. The basic intuition is that any matrix M_f such that $\left(x^{\otimes(d/2)}\right)^T \cdot M_f \cdot x^{\otimes(d/2)}$ for all x (called a matrix representation of f) can be written as a sum of matrices $M_{t,f}$ for each $t \leq d/2$, each of which is block-diagonal matrix with blocks corresponding to matrix representations of the polynomials $M_{F_{2\alpha}}$ for each α with $|\alpha| = 2t$.

Lemma 8.7.2. *Consider any homogeneous n-variate degree-d polynomial $f(x)$. We have,*

$$\Lambda(f) \;\leq\; \max_{\alpha \in \mathbb{N}^n_{\leq d/2}} \frac{\Lambda(F_{2\alpha})}{|\mathcal{O}(\alpha)|}\,(1+d/2)$$

Proof. We shall start by constructing an appropriate matrix representation M_f of f that will give us the desired upper bound on $\Lambda(f)$. To this end, for any $\alpha \in \mathbb{N}^n_{\leq d/2}$, let $M_{F_{2\alpha}}$ be the matrix representation of $F_{2\alpha}$ realizing $\Lambda(F_{2\alpha})$. For any $0 \leq t \leq d/2$, we define $M_{(t,f)}$ so that for any $\alpha \in \mathbb{N}^n_t$ and $I \in \mathcal{O}(\alpha)$, $M_{(t,f)}[I,I] := M_{F_{2\alpha}}/|\mathcal{O}(\alpha)|$, and $M_{(t,f)}$ is zero everywhere else. Now let $M_f := \sum_{t \in [d/2]} M_{(t,f)}$. As for validity of M_f as a representation of f we have,

$$\begin{aligned}
\langle M_f, x^{\otimes d/2}(x^{\otimes d/2})^T \rangle &= \sum_{0 \leq t \leq \frac{d}{2}} \langle M_{(t,f)}, x^{\otimes d/2}(x^{\otimes d/2})^T \rangle \\
&= \sum_{\alpha \in \mathbb{N}^n_{\leq d/2}} \sum_{I \in \mathcal{O}(\alpha)} \langle M_{(|\alpha|,f)}[I,I], x^{\otimes(d/2-|\alpha|)}(x^{\otimes(d/2-|\alpha|)})^T \rangle x^{2\alpha} \\
&= \sum_{\alpha \in \mathbb{N}^n_{\leq d/2}} \sum_{I \in \mathcal{O}(\alpha)} \frac{1}{|\mathcal{O}(\alpha)|} \langle M_{F_{2\alpha}}, x^{\otimes(d/2-|\alpha|)}(x^{\otimes(d/2-|\alpha|)})^T \rangle x^{2\alpha} \\
&= \sum_{\alpha \in \mathbb{N}^n_{\leq d/2}} x^{2\alpha} \cdot \langle M_{F_{2\alpha}}, x^{\otimes(d/2-|\alpha|)}(x^{\otimes(d/2-|\alpha|)})^T \rangle \\
&= \sum_{\alpha \in \mathbb{N}^n_{\leq d/2}} F_{2\alpha}(x) x^{2\alpha} \\
&= f(x)
\end{aligned}$$

Now observe that $M_{(t,f)}$ is a block-diagonal matrix (up to simultaneous permutation of it's rows and columns). Thus we have $\|M_{(t,f)}\| \leq \max_{\alpha \in \mathbb{N}^n_t} \|M_{F_{2\alpha}}\|/|\mathcal{O}(\alpha)|$. Thus on applying triangle inequality, we obtain $\|M_f\| \leq \max_{\alpha \in \mathbb{N}^n_{\leq d/2}} (1+d/2)\,\|M_{F_{2\alpha}}\|/|\mathcal{O}(\alpha)|$ ∎

Lower Bounding $\|f\|_2$ in terms of $\|F_{2\alpha}\|_2$ (non-negative coefficients)

We first bound $\|f\|_2$ in terms of $\max_{\alpha\in\mathbb{N}^n_{\le d/2}} \|F_{2\alpha}\|_2$, when every coefficient of f is non-negative. If x^* is the optimizer of $F_{2\alpha}$, then it is easy to see that $x^* \ge 0$. Setting $y = x^* + \frac{\sqrt{\alpha}}{|\alpha|}$ ensures that $\|y\|_2 \le 2$ and $f(y)$ is large, since $f(y)$ recovers a significant fraction (up to a $2^{O(d)} \cdot |\mathcal{O}(\alpha)|$ factor) of $F_{2\alpha}(x^*)$.

Lemma 8.7.3. *Let $f(x)$ be a homogeneous n-variate degree-d polynomial with non-negative coefficients. Consider any $\alpha \in \mathbb{N}^n_{\le d/2}$. Then*

$$\|f\|_2 \ \ge\ \frac{\|F_{2\alpha}\|_2}{2^{O(d)}\,|\mathcal{O}(\alpha)|}.$$

Proof. Consider any $0 \le t \le d/2$, and any $\alpha \in \mathbb{N}^n_t$. Let $x^*_\alpha := \operatorname{argmax}\|F_{2\alpha}\|_2$ (note x^*_α must be non-negative). Let

$$y^* := x^*_\alpha \ +\ \frac{\sqrt{\alpha}}{\sqrt{t}}$$

and let $x^* := y^*/\|y^*\|$. The second term is a unit vector since $\|\sqrt{\alpha}\|_2^2 = t$. Thus $\|y^*\| = \Theta(1)$ since y^* is the sum of two unit vectors. This implies $f(x^*) \ge f(y^*)/2^{O(d)}$. Now we have,

$$\begin{aligned}
f(y^*) &= \sum_{\beta\in\mathbb{N}^n_{\le d/2}} F_{2\beta}(y^*)\cdot(y^*)^{2\beta} && \text{(by Lemma 8.7.1)}\\
&\ge F_{2\alpha}(y^*)\cdot(y^*)^{2\alpha} && \text{(by non-negativity of coefficients)}\\
&\ge F_{2\alpha}(y^*)\,\frac{1}{t^t}\prod_{\ell\in\operatorname{Supp}\alpha}\alpha_\ell^{\alpha_\ell} && \left(y^* \ge \frac{\sqrt{\alpha}}{\sqrt{t}}\text{ entry-wise}\right)\\
&\ge F_{2\alpha}(y^*)\,\frac{1}{2^{O(t)}\,t!}\prod_{\ell\in\operatorname{Supp}\alpha}\alpha_\ell^{\alpha_\ell}\\
&\ge F_{2\alpha}(y^*)\,\frac{\prod_{\ell\in\operatorname{Supp}\alpha}\alpha_\ell!}{2^{O(t)}\,t!}\\
&\ge F_{2\alpha}(y^*)\,\frac{1}{2^{O(t)}\,|\mathcal{O}(\alpha)|}\\
&\ge F_{2\alpha}(x^*)\,\frac{1}{2^{O(t)}\,|\mathcal{O}(\alpha)|} && \text{(y^* is entry-wise at least x^*)}\\
&= \frac{\|F_{2\alpha}\|_2}{2^{O(t)}\,|\mathcal{O}(\alpha)|}.
\end{aligned}$$

This completes the proof. ∎

Theorem 8.7.4. *Consider any homogeneous n-variate degree-d polynomial $f(x)$ with non-negative coefficients. Then*

$$\frac{\Lambda(f)}{\|f\|_2} \ \le\ 2^{O(d)} \max_{\alpha\in\mathbb{N}^n_{\le d/2}} \frac{\Lambda(F_{2\alpha})}{\|F_{2\alpha}\|_2}.$$

Proof. Combining Lemma 8.7.2 and Lemma 8.7.3 yields the claim. ∎

We will next generalize Theorem 8.7.4 by proving a more general version of Lemma 8.7.3.

Lower Bounding $\|f\|_2$ in terms of $\|F_{2\alpha}\|_2$ (general case)

We lower bound $\|f\|_2$ in terms of $\|F_{2\alpha}\|_2$ for all polynomials. We will first recollect and establish some polynomial identities that will be used in the proof of the generalized version of Lemma 8.7.3 (i.e. Lemma 8.7.11).

Polynomial Identities

Lemma 8.7.5 (Chebyshev's Extremal Polynomial Inequality)**.** *Let $p(x)$ be a univariate degree-d polynomial and let c_d be it's leading coefficient. Then we have,* $\max_{x\in[0,1]} |p(x)| \geq 2|c_d|/4^d$.

Lemma 8.7.6 ([HLZ10])**.** *Let $x^1, x^2, \dots x^d \in \mathbb{R}^n$ be arbitrary, let $\mathcal{A} \in \mathbb{R}^{[n]^d}$ be a SoS-symmetric d-tensor, and let $\xi_1, \dots, \xi_d$ be independent Rademacher random variables. Then*

$$\mathbb{E}\left[\prod_{i\in[d]} \xi_i \,\langle \mathcal{A}, (\xi_1 x^1 + \cdots + \xi_d x^d)^{\otimes d}\rangle\right] = d!\,\langle \mathcal{A}, x^1 \otimes \cdots \otimes x^d\rangle.$$

This lemma implies:

Lemma 8.7.7 ([HLZ10])**.** *Let $\mathcal{A}$ be a SoS-symmetric d-tensor and let $f(x) := \langle \mathcal{A}, x^{\otimes d}\rangle$. Then*

$$\|f\|_2 \;\geq\; \frac{1}{2^{O(d)}} \max_{\|x^i\|=1} \langle \mathcal{A}, x^1 \otimes \cdots \otimes x^d\rangle.$$

Lemma 8.7.8. *Let f be an n-variate degree-d homogeneous polynomial. Let* $\|f\|_2^c := \max_{\substack{z\in\mathbb{C}^n \\ \|z\|=1}} |f(z)|$,

then

$$\frac{\|f\|_2^c}{2^{O(d)}} \leq \|f\|_2 \leq \|f\|_2^c.$$

Proof. Let $\mathcal{A}$ be the SoS-symmetric tensor representing f. Let $z^* = a^* + ib^*$ be the complex unit vector realizing $f(z^*) = \|f\|_2^c$. Then we have,

$$\begin{aligned} f(z^*) &= \langle \mathcal{A}, (z^*)^{\otimes d}\rangle \\ &= \langle \mathcal{A}, (a^* + ib^*)^{\otimes d}\rangle \\ &= \sum_{c^1,\dots c^d \in \{a^*, ib^*\}} \langle \mathcal{A}, \bigotimes_{j\in[d]} c^j\rangle \\ \Rightarrow \mathrm{Re}(f(z^*)) &= \sum_{\substack{c^1,\dots c^d \in \{a^*, b^*\}, \\ |\{j | c^j = b^*\}| \%4 = 0}} \langle \mathcal{A}, \bigotimes_{j\in[d]} c^j\rangle \;-\; \sum_{\substack{c^1,\dots c^d \in \{a^*, b^*\}, \\ |\{j | c^j = b^*\}| \%4 = 2}} \langle \mathcal{A}, \bigotimes_{j\in[d]} c^j\rangle, \end{aligned}$$

$$\operatorname{im} f(z^*) = \sum_{\substack{c^1,\dots c^d \in \{a^*,b^*\},\\ |\{j|c^j=b^*\}|\%4=1}} \langle \mathcal{A}, \bigotimes_{j\in[d]} c^j \rangle - \sum_{\substack{c^1,\dots c^d \in \{a^*,b^*\},\\ |\{j|c^j=b^*\}|\%4=3}} \langle \mathcal{A}, \bigotimes_{j\in[d]} c^j \rangle$$

which implies that there exists $c^1, \dots, c^d \in \{a^*, b^*\}$ such that $|\langle \mathcal{A}, \bigotimes_{j\in[d]} c^j \rangle| \geq \|f\|_2^c / 2^{O(d)}$. Lastly, applying Lemma 8.7.7 implies the claim. ■

Some Probability Facts

Lemma 8.7.9. *Let $X_1, \dots X_k$ be i.i.d.* Bernoulli(p) *random variables. Then for any $t_1, \dots, t_k \in \mathbb{N}$,*
$\mathbb{E}[X_1^{t_1} \dots X_k^{t_k}] = p^k$.

Lemma 8.7.10. *Let ζ be a uniformly random p-th root of unity. Then for any $t \in [p-1]$, $\mathbb{E}[\zeta^t] = 0$. Also, clearly $\mathbb{E}[\zeta^p] = 1$.*

We finally lower bound $\|f\|_2$ in terms of $F_{2\alpha}$. Fix $\alpha \in \mathbb{N}^n_{\leq d/2}$ and, let x^* be the optimizer of $F_{2\alpha}$. Setting $y = x^* + \frac{\sqrt{\alpha}}{|\alpha|}$ as in the non-negative coefficient case does not work since terms from $F_{2\beta}$ may be negative. We bypass this issue by first lower bounding $\|f\|_2^c$ in terms of $F_{2\alpha}$ and using Lemma 8.7.8. For $\|f\|_2^c$, we use random roots of unity and Bernoulli random variables, together with Lemma 8.7.5, to extract nonzero contribution only from the monomials that are multiples of x^α times multilinear parts.

Lemma 8.7.11. *Let $f(x)$ be a homogeneous n-variate degree-d polynomial. Then for any $\alpha \in \mathbb{N}^n_{\leq d/2}$,*

$$\|f\|_2 \;\geq\; \frac{\|F_{2\alpha}\|_2}{2^{O(d)}\,|\mathcal{O}(\alpha)|}.$$

Proof. Fix any any $\alpha \in \mathbb{N}^n_{\leq d/2}$, let $t := |\alpha|$ and let $k := d - 2t$. For any $i \in [n]$, let ζ_i be an independent and uniformly randomly chosen $(2\alpha_i + 1)$-th root of unity, and let Ξ be an independent and uniformly randomly chosen $(k+1)$-th root of unity.

Let $\bar{x} := \operatorname{argmax}\|F_{2\alpha}\|_2$. Let $p \in [0,1]$ be a parameter to be fixed later, let $b_1, \dots, b_n$ be i.i.d. Bernoulli(p) random variables, let $\zeta := (\zeta_1, \dots, \zeta_n)$, $b := (b_1, \dots, b_n)$ and finally let

$$z \quad := \quad \Xi \cdot b \circ \frac{1}{2\alpha + \mathbf{1}} \circ \bar{x} \;+\; \frac{\zeta \circ \sqrt{\alpha}}{\sqrt{t}}.$$

Since $\sum_{\ell \in \operatorname{Supp}\alpha} \alpha_\ell = t$ and roots of unity have magnitude one, z has length $O(1)$. Now consider any fixed $\gamma \in \{0,1\}^n_k$. We have,

$$\mathbb{E}\left[z^{2\alpha+\gamma} \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]$$
$$= \text{coefficient of } \;\Xi^k \cdot \prod_{i\in[n]} \zeta_i^{2\alpha_i} \;\text{ in }\; \mathbb{E}\left[z^{2\alpha+\gamma}\right] \qquad\qquad \text{(by Lemma 8.7.10)}$$

$$= \text{coefficient of } \Xi^k \cdot \prod_{i\in[n]} \zeta_i^{2\alpha_i} \text{ in } \mathbb{E}\left[\prod_{i\in[n]}\left(\zeta_i \cdot \frac{\sqrt{\alpha_i}}{\sqrt{t}} + \Xi \cdot \frac{b_i \cdot \bar{x}_i}{2\alpha_i+1}\right)^{2\alpha_i+\gamma_i}\right]$$

$$= \prod_{i\in[n]} \text{coefficient of } \Xi^{\gamma_i} \cdot \zeta_i^{2\alpha_i} \text{ in } \mathbb{E}\left[\left(\zeta_i \cdot \frac{\sqrt{\alpha_i}}{\sqrt{t}} + \Xi \cdot \frac{b_i \cdot \bar{x}_i}{2\alpha_i+1}\right)^{2\alpha_i+\gamma_i}\right] \qquad (\text{since } \gamma \in \{0,1\}^n_k)$$

$$= p^k \cdot \prod_{i\in \operatorname{Supp}\alpha} \frac{\alpha_i^{\alpha_i}}{t^{\alpha_i}} \cdot \bar{x}_i^{\gamma_i} \qquad (\text{by Lemma 8.7.9})$$

$$= p^k \cdot \bar{x}^{\gamma} \cdot \prod_{i\in \operatorname{Supp}\alpha} \frac{\alpha_i^{\alpha_i}}{t^{\alpha_i}}$$

Thus we have,

$$\mathbb{E}\left[f(z) \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]$$

$$= \sum_{\beta\in\mathbb{N}^n_d} f_\beta \cdot \mathbb{E}\left[z^\beta \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]$$

$$= \sum_{\substack{\beta\in\mathbb{N}^n_d\\ \beta\geq 2\alpha}} f_\beta \cdot \mathbb{E}\left[z^\beta \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right] \qquad (\text{by Lemma 8.7.10})$$

$$= \sum_{\gamma\in\{0,1\}^n_k} f_{2\alpha+\gamma} \cdot \mathbb{E}\left[z^{2\alpha+\gamma} \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right] + \sum_{\substack{\gamma\in\mathbb{N}^n_k\\ \gamma\not\leq \mathbf{1}}} f_{2\alpha+\gamma} \cdot \mathbb{E}\left[z^{2\alpha+\gamma} \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]$$

$$= \sum_{\gamma\in\{0,1\}^n_k} f_{2\alpha+\gamma} \cdot \mathbb{E}\left[z^{2\alpha+\gamma} \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right] + r(p) \qquad (\text{by Lemma 8.7.9})$$

where $r(p)$ is some univariate polynomial in p, s.t. $\deg(r) < k$

$$= \sum_{\gamma\in\{0,1\}^n_k} f_{2\alpha+\gamma} \cdot p^k \cdot \bar{x}^{\gamma} \cdot \prod_{i\in \operatorname{Supp}\alpha} \frac{\alpha_i^{\alpha_i}}{t^{\alpha_i}} + r(p)$$

$$= p^k \cdot F_{2\alpha}(\bar{x}) \cdot \prod_{i\in \operatorname{Supp}\alpha} \frac{\alpha_i^{\alpha_i}}{t^{\alpha_i}} + r(p) \qquad (\text{where } \deg(r) < k)$$

Lastly we have,

$$\begin{aligned} \|f\|_2 &\geq \|f\|_2^c \cdot 2^{-O(d)} && \text{by Lemma 8.7.8} \\ &\geq \max_{p\in[0,1]} \mathbb{E}\left[|f(z)|\right] \cdot 2^{-O(d)} && (\|z\| = O(1)) \end{aligned}$$

$$
\begin{aligned}
&= \max_{p\in[0,1]} \mathbb{E}\left[\left|f(z)\cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right|\right] \cdot 2^{-O(d)} \\
&\geq \max_{p\in[0,1]} \left|\mathbb{E}\left[f(z)\cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]\right| \cdot 2^{-O(d)} \\
&\geq |F_{2\alpha}(\bar{x})| \cdot \prod_{i\in \mathrm{Supp}\,\alpha} \frac{\alpha_i^{\alpha_i}}{t^{\alpha_i}} \cdot 2^{-O(d)} && \text{(by Chebyshev: Lemma 8.7.5)} \\
&= \|F_{2\alpha}\|_2 \cdot \prod_{i\in \mathrm{Supp}\,\alpha} \frac{\alpha_i^{\alpha_i}}{t^{\alpha_i}} \cdot 2^{-O(d)} \\
&\geq \frac{\|F_{2\alpha}\|_2}{|\mathcal{O}(\alpha)|} \cdot 2^{-O(d)}
\end{aligned}
$$

This completes the proof. ∎

In fact, the proof of Lemma 8.7.11 yields a more general result:

Lemma 8.7.12 (Weak Decoupling). *Let $f(x)$ be a homogeneous n-variate degree-d polynomial. Then for any $\alpha \in \mathbb{N}^n_{\leq d/2}$ and any unit vector y,*

$$
\|f\|_2 \;\geq\; y^{2\alpha} \cdot \|F_{2\alpha}\|_2 \cdot 2^{-O(d)}.
$$

We are finally able to establish the multilinear reduction result that is the focus of this section.

Theorem 8.7.13. *Let $f(x)$ be a homogeneous n-variate degree-d (for even d) polynomial. Then*

$$
\frac{\Lambda(f)}{\|f\|_2} \;\leq\; 2^{O(d)} \max_{\alpha\in\mathbb{N}^n_{\leq d/2}} \frac{\Lambda(F_{2\alpha})}{\|F_{2\alpha}\|_2}.
$$

Proof. Combining Lemma 8.7.2 and Lemma 8.7.11 yields the claim. ∎

8.7.2 $(n/q)^{d/4}$-Approximation for Non-negative Coefficient Polynomials

Theorem 8.7.14. *Consider any homogeneous multilinear n-variate degree-d polynomial $f(x)$ with non-negative coefficients. We have,*

$$
\frac{\Lambda(f)}{\|f\|_2} \;\leq\; 2^{O(d)} \frac{n^{d/4}}{d^{d/4}}.
$$

Proof. Let M_f be the SoS-symmetric matrix representation of f. Let $I^* = (i_1, \ldots, i_{d/2}) \in [n]^{d/2}$ be the multi-index of any row of M_f with maximum row sum. Let S_I for $I \in [n]^{d/2}$, denote the sum of the row I of M_f. By Perron-Frobenius theorem, $\|\mathsf{M}_f\| \leq S_{I^*}$. Thus $\Lambda(f) \leq S_{I^*}$.

We next proceed to bound $\|f\|_2$ from below. To this end, let $x^* := y^*/\|y^*\|$ where,

$$y^* := \frac{\mathbf{1}}{\sqrt{n}} + \frac{1}{\sqrt{d/2}} \sum_{i \in I^*} e_i$$

Since f is multilinear, I^* has all distinct elements, and so the second term in the definition of y^* is of unit length. Thus $\|y^*\| = \Theta(1)$, which implies that $\|f\|_2 \geq f(x^*) \geq f(y^*)/2^{O(d)}$. Now we have,

$$\begin{aligned}
f(y^*) &= ((y^*)^{\otimes d/2})^T \mathsf{M}_f\, (y^*)^{\otimes d/2} \\
&\geq \sum_{I \in \mathcal{O}(I^*)} \frac{1}{(nd)^{d/4}}\, e_{I(1)}^T \otimes \cdots \otimes e_{I(d/2)}^T\, \mathsf{M}_f\, \mathbf{1}^{\otimes d/2} \quad \text{(by non-negativity of } \mathsf{M}_f\text{)} \\
&= \sum_{I \in \mathcal{O}(I^*)} \frac{1}{(nd)^{d/4}}\, e_I^T\, \mathsf{M}_f\, \mathbf{1}\ (\in \mathbb{R}^{[n]^{d/2}}) \\
&= \sum_{I \in \mathcal{O}(I^*)} \frac{S_I}{(nd)^{d/4}} \\
&= \sum_{I \in \mathcal{O}(I^*)} \frac{S_{I^*}}{(nd)^{d/4}} \quad \text{(by SoS-symmetry of } \mathsf{M}_f\text{)} \\
&= \frac{(d/2)!\, S_{I^*}}{(nd)^{d/4}} \quad (|\mathcal{O}(I^*)| = (d/2)! \text{ by multilinearity of } f) \\
&\geq \frac{d^{d/4} S_{I^*}}{n^{d/4}\, 2^{O(d)}} \geq \frac{d^{d/4} \Lambda(f)}{n^{d/4}\, 2^{O(d)}}.
\end{aligned}$$

This completes the proof. ∎

Theorem 8.7.15. *Let $f(x)$ be a homogeneous n-variate degree-d polynomial with non-negative coefficients. Then for any even q such that d divides q,*

$$\frac{(\Lambda(f^{q/d}))^{d/q}}{\|f\|_2} \leq 2^{O(d)}\, \frac{n^{d/4}}{q^{d/4}}.$$

Proof. Applying Theorem 8.7.4 to $f^{q/d}$ and combining this with Theorem 8.7.14 yields the claim. ∎

8.7.3 $(n/q)^{d/2}$-Approximation for General Polynomials

Theorem 8.7.16. *Consider any homogeneous multilinear n-variate degree-d (for even d) polynomial $f(x)$. We have,*

$$\frac{\Lambda(f)}{\|f\|_2} \leq 2^{O(d)}\, \frac{n^{d/2}}{d^{d/2}}.$$

Proof. Let M_f be the SoS-symmetric matrix representation of f, i.e.

$$\mathsf{M}_f[I, J] = \frac{f_{\alpha(I)+\alpha(J)}}{|\mathcal{O}(\alpha(I) + \alpha(J))|}.$$

By the Gershgorin circle theorem, we can bound $\|\mathsf{M}_f\|_2$, and hence $\Lambda(f)$ by $n^{d/2} \cdot (\max_\beta |f_\beta| / d!)$. Here, we use the multilinearity of f. On the other hand for a multilinear polynomial, using $x = \beta/\sqrt{|\beta|}$ (where $|\beta| = d$ by multilinearity), gives $\|f\|_2 \geq d^{-d/2} \cdot |f_\beta|$. Thus, we easily get

$$\Lambda(f) \;\leq\; \frac{d^{d/2}}{d!} \cdot n^{d/2} \cdot \|f\|_2 \;=\; 2^{O(d)}\, \frac{n^{d/2}}{d^{d/2}}.$$

■

Theorem 8.7.17. *Let $f(x)$ be a homogeneous n-variate degree-d polynomial, and assume that $2d$ divides q. Then*

$$\frac{(\Lambda(f^{q/d}))^{d/q}}{\|f\|_2} \;\leq\; 2^{O(d)}\, \frac{n^{d/2}}{q^{d/2}}.$$

Proof. Applying Theorem 8.7.13 to $f^{q/d}$ and combining this with Theorem 8.7.16 yields the claim. ■

8.7.4 $\sqrt{m/q}$-Approximation for m-sparse polynomials

Lemma 8.7.18. *Consider any homogeneous multilinear n-variate degree-d (for even d) polynomial $f(x)$ with m non-zero coefficients. We have,*

$$\frac{\Lambda(f)}{\|f\|_2} \;\leq\; 2^{O(d)}\sqrt{m}.$$

Proof. Let M_f be the SoS-symmetric matrix representation of f, i.e.

$$\mathsf{M}_f[I, J] = \frac{f_{\alpha(I)+\alpha(J)}}{|\mathcal{O}(\alpha(I) + \alpha(J))|}.$$

Now $\Lambda(f) \leq \|\mathsf{M}_f\| \leq \|\mathsf{M}_f\|_F$. Thus we have,

$$\begin{aligned}
\|\mathsf{M}_f\|_F^2 &= \sum_{I,J \in [n]^{d/2}} \mathsf{M}_f[I,J]^2 \\
&= \sum_{\beta \in \{0,1\}^n_d} \frac{f_\beta^2}{|\mathcal{O}(\beta)|} \\
&= \sum_{\beta \in \{0,1\}^n_d} \frac{f_\beta^2}{d!} \\
&\leq \frac{m}{d!} \cdot \max_\beta |f_\beta|
\end{aligned}$$

On the other hand, since f is multilinear, using $x = \beta/\sqrt{|\beta|}$ (where $|\beta| = d$ by multilinearity), implies $\|f\|_2 \geq d^{-d/2} \cdot |f_\beta|$ for any β. This implies the claim. ∎

Theorem 8.7.19. *Let $f(x)$ be a homogeneous n-variate degree-d polynomial with m non-zero coefficients, and assume that $2d$ divides q. Then*

$$\frac{(\Lambda(f^{q/d}))^{d/q}}{\|f\|_2} \leq 2^{O(d)}\sqrt{m/q}.$$

Proof. Combining Theorem 8.7.13 and Lemma 8.7.18, yields that for any degree-q homogeneous polynomial g with sparsity $\overline{m}$, we have

$$\frac{(\Lambda(g))}{\|g\|_2} \leq 2^{O(q)}\sqrt{\overline{m}}.$$

Lastly, taking $g = f^{q/d}$ and observing that the sparsity of g is at most $\left(\!\binom{m}{q/d}\!\right)$ implies the claim. ∎

8.8 Weak Decoupling/Approximating 2-norms via Folding

8.8.1 Preliminaries

Recall that we call a folded polynomial multilinear if all its monomials are multilinear. In particular, there's no restriction on the folds of the polynomial.

Lemma 8.8.1 (Folded Analogue of Lemma 8.7.1)**.**
Let $(\mathbb{R}_{d_2}[x])_{d_1}[x] \ni f(x) := \sum_{\beta \in \mathbb{N}^n_{d_1}} \overline{f}_\beta(x) \cdot x^\beta$ be a (d_1, d_2)-folded polynomial. f can be written as

$$\sum_{\alpha \in \mathbb{N}^n_{\leq d_1/2}} F_{2\alpha}(x) \cdot x^{2\alpha}$$

where for any $\alpha \in \mathbb{N}^n_{\leq d_1/2}$, $F_{2\alpha}(x)$ is a multilinear $(d_1 - 2|\alpha|, d_2)$-folded polynomial.

Proof. Simply consider the folded polynomial

$$F_{2\alpha}(x) = \sum_{\gamma \in \{0,1\}^n_{d_1 - 2|\alpha|}} \overline{(F_{2\alpha})}_\gamma \cdot x^\gamma$$

where $\overline{(F_{2\alpha})}_\gamma = \overline{f}_{2\alpha+\gamma}$. ∎

8.8.2 Reduction to Multilinear Folded Polynomials

Here we will prove a generalized version of Lemma 8.7.2, which is a generalization in two ways; firstly it allows for folds instead of just coefficients, and secondly it allows a more general set of constraints than just the hypersphere since we will need to add some additional non-negativity constraints for the case of non-negative coefficient polynomials (so that $\Lambda_C()$ satisfies monotonicity over NNC polynomials which will come in handy later).

Recall that $\Lambda_C()$ is defined in Section 7.4 and that $\|f\|_2$ and $\Lambda_C(f)$ for a folded polynomial f, are applied to the unfolding of f.

Relating $\Lambda_C(f)$ to $\Lambda_C(F_{2\alpha})$

Lemma 8.8.2 (Folded Analogue of Lemma 8.7.2)**.**
Let C be a system of polynomial constraints of the form $\{\|x\|_2^2 = 1\} \cup C'$ where C' is a moment non-negativity constraint set. Let $f \in (\mathbb{R}_{d_2}[x])_{d_1}[x]$ be a (d_1, d_2)-folded polynomial. We have,

$$\Lambda_C(f) \;\le\; \max_{\alpha \in \mathbb{N}^n_{\le d_1/2}} \frac{\Lambda_C(F_{2\alpha})}{|\mathcal{O}(\alpha)|}\,(1 + d_1/2)$$

Proof. Consider any degree-$(d_1 + d_2)$ pseudo-expectation operator $\widetilde{\mathbf{E}}_C$. We have,

$$\begin{aligned}
\widetilde{\mathbf{E}}_C[f] &= \sum_{\alpha \in \mathbb{N}^n_{\le d_1/2}} \widetilde{\mathbf{E}}_C\left[F_{2\alpha}(x) \cdot x^{2\alpha}\right] && \text{(by Lemma 8.8.1)} \\
&\le \sum_{\alpha \in \mathbb{N}^n_{\le d_1/2}} \widetilde{\mathbf{E}}_C\left[x^{2\alpha}\right] \cdot \Lambda_C(F_{2\alpha}) && \text{(by Lemma 7.4.2)} \\
&= \sum_{0 \le t \le \frac{d_1}{2}} \sum_{\alpha \in \mathbb{N}^n_t} \widetilde{\mathbf{E}}_C\left[x^{2\alpha}\right] \cdot \Lambda_C(F_{2\alpha}) \\
&= \sum_{0 \le t \le \frac{d_1}{2}} \sum_{\alpha \in \mathbb{N}^n_t} \widetilde{\mathbf{E}}_C\left[|\mathcal{O}(\alpha)| x^{2\alpha}\right] \cdot \frac{\Lambda_C(F_{2\alpha})}{|\mathcal{O}(\alpha)|} \\
&\le \sum_{0 \le t \le \frac{d_1}{2}} \sum_{\alpha \in \mathbb{N}^n_t} \widetilde{\mathbf{E}}_C\left[|\mathcal{O}(\alpha)| x^{2\alpha}\right] \cdot \max_{\beta \in \mathbb{N}^n_{\le d_1/2}} \frac{\Lambda_C(F_{2\beta})}{|\mathcal{O}(\beta)|} && \left(\widetilde{\mathbf{E}}_C\left[x^{2\alpha}\right] \ge 0\right) \\
&= \sum_{0 \le t \le \frac{d_1}{2}} \widetilde{\mathbf{E}}_C\left[\sum_{\alpha \in \mathbb{N}^n_t} |\mathcal{O}(\alpha)| x^{2\alpha}\right] \cdot \max_{\beta \in \mathbb{N}^n_{\le d_1/2}} \frac{\Lambda_C(F_{2\beta})}{|\mathcal{O}(\beta)|} \\
&= \sum_{0 \le t \le \frac{d_1}{2}} \widetilde{\mathbf{E}}_C\left[\|x\|_2^{2t}\right] \cdot \max_{\beta \in \mathbb{N}^n_{\le d_1/2}} \frac{\Lambda_C(F_{2\beta})}{|\mathcal{O}(\beta)|} \\
&= \sum_{0 \le t \le \frac{d_1}{2}} \max_{\beta \in \mathbb{N}^n_{\le d_1/2}} \frac{\Lambda_C(F_{2\beta})}{|\mathcal{O}(\beta)|} \\
&= \max_{\beta \in \mathbb{N}^n_{\le d_1/2}} \frac{\Lambda_C(F_{2\beta})}{|\mathcal{O}(\beta)|}(1 + d_1/2)
\end{aligned}$$

■

8.8.3 Relating Evaluations of f to Evaluations of $F_{2\alpha}$

Here we would like to generalize Lemma 8.7.3 and Lemma 8.7.11 to allow folds, however for technical reasons related to decoupling of the domain of the folds from the domain of the monomials of a folded polynomial, we instead generalize claims implicit in the proofs of Lemma 8.7.3 and Lemma 8.7.11.

Let $f \in (\mathbb{R}_{d_2}[x])_{d_1}[x]$ be a (d_1, d_2)-folded polynomial. Recall that an evaluation of a folded polynomial treats the folds as coefficients and only substitutes values in the monomials of the folded polynomial. Thus for any fixed $y \in \mathbb{R}^n$, $f(y)$ (sometimes denoted by $(f(y))(x)$ for contextual clarity) is a degree-d_2 polynomial in x, i.e. $f(y) \in \mathbb{R}_{d_2}[x]$.

Lemma 8.8.3 (Folded Analogue of Lemma 8.7.3)**.**
Let $f \in \left(\mathbb{R}^+_{d_2}[x]\right)_{d_1}[x]$ be a (d_1, d_2)-folded polynomial whose folds have non-negative coefficients. Then for any $\alpha \in \mathbb{N}^n_{\le d_1/2}$ and any $y \ge 0$,

$$\left(f\left(y + \frac{\sqrt{\alpha}}{\sqrt{|\alpha|}}\right)\right)(x) \;\ge\; \frac{(F_{2\alpha}(y))(x)}{|\mathcal{O}(\alpha)|} \cdot 2^{-O(d_1)}$$

where the ordering is coefficient-wise.

Proof. Identical to the proof of Lemma 8.7.3. ∎

Lemma 8.8.4 (Folded Analogue of Lemma 8.7.11)**.**
Let $f \in (\mathbb{R}_{d_2}[x])_{d_1}[x]$ be a (d_1, d_2)-folded polynomial. Consider any $\alpha \in \mathbb{N}^n_{\le d_1/2}$ and any y, and let

$$z \quad := \quad \Xi \cdot y \circ \frac{1}{2\alpha + \mathbf{1}} \circ b \;+\; \frac{\sqrt{\alpha} \circ \zeta}{\sqrt{|\alpha|}}$$

where Ξ is an independent and uniformly randomly chosen $(d_1 - 2|\alpha| + 1)$-th root of unity, and for any $i \in [n]$, ζ_i is an independent and uniformly randomly chosen $(2\alpha_i + 1)$-th root of unity, and b_i is an independent Bernoulli(p) *random variable (p is an arbitrary parameter in $[0, 1]$). Then*

$$\mathbb{E}\left[(f(z))(x) \cdot \Xi \cdot \prod_{i \in [n]} \zeta_i\right] \;=\; p^{d_1 - 2|\alpha|} \cdot \frac{(F_{2\alpha}(y))(x)}{|\mathcal{O}(\alpha)|} \cdot 2^{-O(d_1)} \;+\; r(p)$$

where $r(p)$ is a univariate polynomial in p with degree less than $d_1 - 2|\alpha|$ (and whose coefficients are in $\mathbb{R}_{d_2}[x]$).

Proof. This follows by going through the proof of Lemma 8.7.11 for every fixed x. ∎

8.8.4 Bounding $\Lambda_C()$ of Multilinear Folded Polynomials

Here we bound $\Lambda_C()$ of a multilinear folded polynomial in terms of properties of the polynomial that are inspired by treating the folds as coefficients and generalizing the coefficient-based approximations for regular (non-folded) polynomials from Theorem 8.7.16 and Theorem 8.7.14.

General Folds: Bounding $\Lambda()$ in terms of $\Lambda()$ of the "worst" fold

Here we will give a folded analogue of the proof of Theorem 8.7.16 wherein we used Gershgorin-Circle theorem to bound SOS value in terms of the max-magnitude-coefficient.

Lemma 8.8.5 (Folded Analogue of Gershgorin Circle Bound on Spectral Radius)**.** *For even d_1, d_2, let $d = d_1 + d_2$, let $f \in \left(\mathbb{R}_{d_2}[x]\right)_{d_1}[x]$ be a multilinear (d_1, d_2)-folded polynomial. We have,*

$$\Lambda(f) \;\leq\; 2^{O(d)}\, \frac{n^{d_1/2}}{d_1^{d_1}} \max_{\gamma \in \{0,1\}^n_{d_1}} \|\overline{f}_\gamma\|_{sp}.$$

Proof. Since $\Lambda(f) \leq \|f\|_{sp}$, it is sufficient to bound $\|f\|_{sp}$.

Let $M_{\overline{f}_\gamma}$ be the matrix representation of $\overline{f}_\gamma$ realizing $\|\overline{f}_\gamma\|_{sp}$. Let M_f be an $[n]^{d_1/2} \times [n]^{d_1/2}$ block matrix with $[n]^{d_2/2} \times [n]^{d_2/2}$ size blocks, where for any $I, J \in [n]^{d_1/2}$ the block of M_f at index (I, J) is defined to be $\frac{1}{d_1!} \cdot M_{\overline{f}_{\alpha(I)+\alpha(J)}}$. Clearly M_f (interpreted as an $[n]^{d/2} \times [n]^{d/2}$) is a matrix representation of the unfolding of f since f is a multilinear folded polynomial. Lastly, applying Block-Gershgorin circle theorem to M_f and upper bounding the sum of spectral norms over a block row by $n^{d_1/2}$ times the max term implies the claim. ∎

Non-Negative Coefficient Folds: Relating SoS Value to the SoS Value of the $d_1/2$-collapse

Observe that in the case of a multilinear degree-d polynomial, the $d/2$-collapse corresponds (up to scaling) to the sum of a row of the SOS symmetric matrix representation of the polynomial. We will next develop a folded analogue of the proof of Theorem 8.7.14 wherein we employed Perron-Frobenius theorem to bound SOS value in terms of the $d/2$-collapse.

The proof here however, is quite a bit more subtle than in the general case above. This is because one can apply the block-matrix analogue of Gershgorin theorem (due to Feingold et al. [FV+62]) to a matrix representation of the folded polynomial (whose spectral norm is an upper bound on $\Lambda()$) in the general case. Loosely speaking, this corresponds to bounding $\Lambda(f)$ in terms of

$$\max_{\gamma \in \{0,1\}^n_k} \sum_{\theta \in \{0,1\}^n_k} \Lambda\left(\overline{f}_{\gamma+\theta}\right)$$

where $k = d_1/2$. This however is not enough in the nnc case as in order to win the $1/2$ in the exponent, one needs to relate $\Lambda_C(f)$ to

$$\max_{\gamma \in \{0,1\}^n_k} \Lambda\left(\sum_{\theta \in \{0,1\}^n_k} \overline{f}_{\gamma+\theta}\right).$$

This however, cannot go through Block-Gershgorin since it is **not** true that the spectral norm of a non-negative block matrix is upper bounded by the max over rows of the spectral norm of the sum of blocks in that row. It instead, can only be upper bounded by the max over rows of the sum of spectral norms of the blocks in that row.

To get around this issue, we skip the intermediate step of bounding $\Lambda_C(f)$ by the spectral norm of a matrix and instead prove the desired relation directly through the use of pseudoexpectation operators. This involved first finding a pseudo-expectation based proof of Gershgorin/Perron-Frobenius bound on spectral radius that generalizes to folded polynomials in the right way.

Lemma 8.8.6 (Folded analogue of Perron-Frobenius Bound on Spectral Radius)**.** *For even $d_1 = 2k$, let $f \in \left(\mathbb{R}^+_{d_2}[x]\right)_{d_1}[x]$ be a multilinear (d_1, d_2)-folded polynomial whose folds have non-negative coefficients. Let C be the system of polynomial constraints given by $\{\|x\|_2^2 = 1; \forall \beta \in \mathbb{N}^n_{d_2}, x^\beta \geq 0\}$. We have,*

$$\Lambda_C(f) \;\leq\; \max_{\gamma \in \{0,1\}^n_k} \Lambda_C(\bar{g}_\gamma) \cdot \frac{1}{k!}$$

where

$$\bar{g}_\gamma(x) := \overline{\mathsf{C}_k(f)}_\gamma = \sum_{\substack{\theta \leq \mathbf{1} - \gamma \\ \theta \in \mathbb{N}^n_k}} \bar{f}_{\gamma+\theta}(x).$$

Proof. Consider any pseudo-expectation operator $\widetilde{\mathbf{E}}_C$ of degree at least $d_1 + d_2$. Note that since $\widetilde{\mathbf{E}}_C$ satisfies $\{\forall \beta \in \mathbb{N}^n_{d_2}, x^\beta \geq 0\}$, by linearity $\widetilde{\mathbf{E}}_C$ must also satisfy $\{h \geq 0\}$ for any $h \in \mathbb{R}^+_{d_2}[x]$ - a fact we will use shortly.

Since f is a multilinear folded polynomial, $\bar{f}_\alpha$ is only defined when $0 \leq \alpha \leq \mathbf{1}$. If $\alpha \not\leq \mathbf{1}$, we define $\bar{f}_\alpha := 0$ We have,

$$\begin{aligned}
\widetilde{\mathbf{E}}_C[f] &= \sum_{\alpha \in \{0,1\}^n_{d_1}} \widetilde{\mathbf{E}}_C\left[\bar{f}_\alpha \cdot x^\alpha\right] && (f \text{ is a multilinear folded polynomial}) \\
&= \sum_{I \in [n]^k} \sum_{J \in [n]^k} \widetilde{\mathbf{E}}_C\left[\bar{f}_{\alpha(I)+\alpha(J)} \cdot x^I x^J\right] \cdot \frac{1}{d_1!} && (\text{by multilinearity}) \\
&\leq \sum_{I \in [n]^k} \sum_{J \in [n]^k} \widetilde{\mathbf{E}}_C\left[\bar{f}_{\alpha(I)+\alpha(J)} \cdot \frac{(x^I)^2 + (x^J)^2}{2}\right] \cdot \frac{1}{d_1!} && (\widetilde{\mathbf{E}}_C \text{ satisfies } \bar{f}_\alpha \geq 0) \\
&= \sum_{I \in [n]^k} \sum_{J \in [n]^k} \widetilde{\mathbf{E}}_C\left[\bar{f}_{\alpha(I)+\alpha(J)} \cdot (x^I)^2\right] \cdot \frac{1}{d_1!} \\
&= \sum_{I \in [n]^k} \widetilde{\mathbf{E}}_C\left[(x^I)^2 \cdot \sum_{J \in [n]^k} \bar{f}_{\alpha(I)+\alpha(J)}\right] \cdot \frac{1}{d_1!}
\end{aligned}$$

$$= \sum_{I \in [n]^k} \widetilde{\mathbf{E}}_C \left[(x^I)^2 \cdot \sum_{\substack{\theta \leq \mathbf{1} - \alpha(I) \\ \theta \in \mathbb{N}_k^n}} \overline{f}_{\alpha(I)+\theta} \right] \cdot \frac{k!}{d_1!} \qquad \text{(by multilinearity)}$$

$$= \sum_{I \in [n]^k} \widetilde{\mathbf{E}}_C \left[(x^I)^2 \cdot g_{\alpha(I)} \right] \cdot \frac{k!}{d_1!}$$

$$\leq \sum_{I \in [n]^k} \widetilde{\mathbf{E}}_C \left[(x^I)^2 \right] \cdot \Lambda_C \left(\overline{g}_{\alpha(I)} \right) \cdot \frac{1}{k!} \qquad \text{(by Lemma 7.4.2)}$$

$$\leq \sum_{I \in [n]^k} \widetilde{\mathbf{E}}_C \left[(x^I)^2 \right] \cdot \max_{\gamma \in \{0,1\}_k^n} \Lambda_C(\overline{g}_\gamma) \cdot \frac{1}{k!} \qquad \left(\widetilde{\mathbf{E}}_C \left[(x^I)^2 \right] \geq 0\right)$$

$$= \widetilde{\mathbf{E}}_C \left[\|x\|_2^{d_1} \right] \cdot \max_{\gamma \in \{0,1\}_k^n} \Lambda_C(\overline{g}_\gamma) \cdot \frac{1}{k!}$$

$$= \max_{\gamma \in \{0,1\}_k^n} \Lambda_C(\overline{g}_\gamma) \cdot \frac{1}{k!}$$

■

We are finally equipped to prove the main results of this section.

8.8.5 $(n/q)^{d/4-1/2}$-Approximation for Non-negative Coefficient Polynomials

Theorem 8.8.7. *Consider any $f \in \mathbb{R}_d^+[x]$ for $d \geq 2$, and any q divisible by $2d$. Let C be the system of polynomial constraints given by $\{\|x\|_2^2 = 1; \forall \beta \in \mathbb{N}_{2q/d}^n, x^\beta \geq 0\}$. Then we have,*

$$\frac{\Lambda_C \left(f^{q/d} \right)^{d/q}}{\|f\|_2} \leq 2^{O(d)} \frac{n^{d/4-1/2}}{q^{d/4-1/2}}.$$

Proof. Let h be any $(d-2,2)$-folded polynomial whose unfolding yields f and whose folds have non-negative coefficients and let s be the $(\overline{q}, 2q/d)$-folded polynomial given by $h^{q/d}$ where $\overline{q} := (d-2)q/d$. Finally, consider any $\alpha \in \mathbb{N}_{\leq \overline{q}/2}^n$ and let $S_{2\alpha}$ be the multilinear component of s as defined in Lemma 8.8.1. We will establish that for any $\gamma \in \{0,1\}_k^n$ (where $k := \overline{q}/2 - |\alpha|$),

$$\|f\|_2^{q/d} \geq \frac{2^{-O(q)} \cdot \Lambda_C \left(\overline{C_{\overline{q}/2-|\alpha|}(S_{2\alpha})_\gamma} \right)}{(\overline{q}/2 - |\alpha|)^{\overline{q}/4-|\alpha|/2} \cdot |\mathcal{O}(\alpha)| \cdot n^{\overline{q}/4-|\alpha|/2}} \tag{8.7}$$

which on combining with the application of Lemma 8.8.2 to s and its composition with Lemma 8.8.6, yields the claim. To elaborate, we apply Lemma 8.8.2 to s with $d_1 = \overline{q}, d_2 = 2q/d$ and then for every $\alpha \in \mathbb{N}_{\leq \overline{q}/2}^n$ we apply Lemma 8.8.6 with $d_1 = \overline{q} - 2|\alpha|, d_2 = 2q/d$, to get

$$\Lambda_C \left(f^{q/d} \right) = \Lambda_C(s) \leq 2^{O(q)} \cdot \max_{\alpha \in \mathbb{N}_{\leq \overline{q}/2}^n} \max_{\gamma \in \{0,1\}_{\overline{q}/2-|\alpha|}^n} \frac{\Lambda_C \left(\overline{C_{\overline{q}/2-|\alpha|}(S_{2\alpha})_\gamma} \right)}{(\overline{q}/2 - |\alpha|)! \cdot |\mathcal{O}(\alpha)|}$$

which on combining with Eq. (8.7) yields the claim.

It remains to establish Eq. (8.7). So fix any α, γ satisfying the above conditions. Let $t := |\alpha|$ and let $k := \bar{q}/2 - |\alpha|$. Clearly $\|f\|_2 \geq f(y/\|y\|_2)$ where $y := a + z$, and

$$z := \frac{\mathbf{1}}{\sqrt{n}} + \frac{\gamma}{\sqrt{k}} + \frac{\sqrt{\alpha}}{\sqrt{t}}$$

and a is the unit vector that maximizes the quadratic polynomial

$$(h(z))(x).$$

Since $\|y\|_2 = O(1)$, $\|f\|_2 \geq f(y)/2^{O(d)}$. Now clearly by non-negativity we have,

$$f(y) \geq (h(z))(a) = \|h(z)\|_2$$

Thus we have,

$$\begin{aligned}
\|f\|_2^{q/d} &\geq \|(h(z))(x)\|_2^{q/d} \cdot 2^{-O(q)} \\
&= \|h(z)^{q/d}(x)\|_2 \cdot 2^{-O(q)} \\
&= \Lambda_C\Big(h(z)^{q/d}(x)\Big) \cdot 2^{-O(q)} && \text{(SOS exact on powered quadratics)} \\
&= \Lambda_C(s(z)(x)) \cdot 2^{-O(q)} \\
&\geq \Lambda_C\Big(S_{2\alpha}(\mathbf{1}/\sqrt{n} + \gamma/\sqrt{k})(x)\Big) \cdot \frac{2^{-O(q)}}{|\mathcal{O}(\alpha)|} && \text{(by Lemma 7.4.3 and Lemma 8.8.3)} \\
&\geq \frac{\Lambda_C\Big(\overline{\mathsf{C}_k(S_{2\alpha})_\gamma}\Big)}{k^{k/2} \cdot n^{k/2}} \cdot \frac{2^{-O(q)}}{|\mathcal{O}(\alpha)|} && \text{(by Lemma 7.4.3, and} \\
& && S_{2\alpha}(\frac{\mathbf{1}}{\sqrt{n}} + \frac{\gamma}{\sqrt{k}}) \geq \overline{\mathsf{C}_k(S_{2\alpha})_\gamma} \text{ coefficient-wise)}
\end{aligned}$$

which completes the proof since we've established Eq. (8.7). ■

8.8.6 $(n/q)^{d/2-1}$-Approximation for General Polynomials

Theorem 8.8.8. *Consider any $f \in \mathbb{R}_d^+[x]$ for $d \geq 2$, and any q divisible by $2d$. Then we have,*

$$\frac{\Lambda\big(f^{q/d}\big)^{d/q}}{\|f\|_2} \;\leq\; 2^{O(d)}\;\frac{n^{d/2-1}}{q^{d/2-1}}.$$

Proof. Let h be the unique $(d-2, 2)$-folded polynomial whose unfolding yields f and such that for any $\beta \in \mathbb{N}_{d-2}^n$, the fold $\bar{h}_\beta$ of h is equal up to scaling, to the quadratic form of the corresponding $(n \times n)$ block of the SOS-symmetric matrix representation M_f of f. That is, for any $I, J \in [n]^{d/2-1}$, s.t. $\alpha(I) + \alpha(J) = \beta$,

$$\bar{h}_\beta(x) = \frac{x^T \mathsf{M}_f[I, J] x}{|\mathcal{O}(\beta)|}.$$

Let s be the $(\overline{q}, 2q/d)$-folded polynomial given by $h^{q/d}$ where $\overline{q} := (d-2)q/d$. Consider any $\alpha \in \mathbb{N}^n_{\leq \overline{q}/2}$ and $\gamma \in \{0,1\}^n_{\overline{q}-2|\alpha|}$, and let $S_{2\alpha}$ be the multilinear component of s as defined in Lemma 8.8.1. Below we will show,

$$\|f\|_2^{q/d} \geq \frac{2^{-O(q)} \cdot \|\overline{(S_{2\alpha})}_\gamma\|_{sp}}{(\overline{q}-2|\alpha|)^{\overline{q}/2-|\alpha|} \cdot |\mathcal{O}(\alpha)|} \tag{8.8}$$

which would complete the proof after applying Lemma 8.8.2 to s and composing the result with Lemma 8.8.5. To elaborate, we apply Lemma 8.8.2 to s with $d_1 = \overline{q}, d_2 = 2q/d$ and then for every $\alpha \in \mathbb{N}^n_{\leq \overline{q}/2}$ we apply Lemma 8.8.5 with $d_1 = \overline{q} - 2|\alpha|, d_2 = 2q/d$, to get

$$\Lambda\left(f^{q/d}\right) = \Lambda(s) \leq 2^{O(q)} \cdot \max_{\alpha \in \mathbb{N}^n_{\leq \overline{q}/2}} \max_{\gamma \in \{0,1\}^n_{\overline{q}-2|\alpha|}} \frac{\|\overline{(S_{2\alpha})}_\gamma\|_{sp}}{(\overline{q}-2|\alpha|)^{\overline{q}-2|\alpha|} \cdot |\mathcal{O}(\alpha)|}$$

which on combining with Eq. (8.8) yields the claim.

Fix any α, γ satisfying the above conditions. Let $k := \overline{q} - 2\alpha$. Let $t := |\alpha|$, and let

$$z := \Xi \cdot \frac{1}{\sqrt{k}} \cdot \gamma \circ \frac{1}{2\alpha + \mathbf{1}} \circ b \;\; + \;\; \frac{\sqrt{\alpha} \circ \zeta}{\sqrt{t}}$$

Ξ is an independent and uniformly randomly chosen $(k+1)$-th root of unity, and for any $i \in [n]$, ζ_i is an independent and uniformly randomly chosen $(2\alpha_i + 1)$-th root of unity, and for any $i \in [n]$, b_i is an independent Bernoulli(p) random variable (p is a parameter that will be set later). By *Lemma* 8.7.7 and definition of h, we see that for any y, $\|f\|_2^c \geq \|(h(y))(x)\|_2^c$. Thus we have,

$$\begin{aligned}
\|f\|_2^{q/d} &= \|f^{q/d}\|_2 \\
&\geq \|f^{q/d}\|_2^c \cdot 2^{-O(q)} && \text{(by Lemma 8.7.8)} \\
&\geq \max_{p\in[0,1]} \mathbb{E}\left[\|h(z)^{q/d}(x)\|_2\right] \cdot 2^{-O(q)} && \text{(by Lemma 8.7.7)} \\
&= \max_{p\in[0,1]} \mathbb{E}\left[\|h(z)^{q/d}(x)\|_{sp}\right] \cdot 2^{-O(q)} && \text{(SOS exact on powered quadratics)} \\
&= \max_{p\in[0,1]} \mathbb{E}\left[\|h(z)^{q/d}(x) \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\|_{sp}\right] \cdot 2^{-O(q)} \\
&\geq \max_{p\in[0,1]} \left\|\mathbb{E}\left[h(z)^{q/d}(x) \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]\right\|_{sp} \cdot 2^{-O(q)} \\
&= \max_{p\in[0,1]} \left\|\mathbb{E}\left[(s(z))(x) \cdot \Xi \cdot \prod_{i\in[n]} \zeta_i\right]\right\|_{sp} \cdot 2^{-O(q)} \\
&= \max_{p\in[0,1]} \left\| p^k \cdot \frac{(S_{2\alpha}(\gamma/\sqrt{k}))(x)}{|\mathcal{O}(\alpha)|} + r(p)\right\|_{sp} \cdot 2^{-O(q)} && \text{(by Lemma 8.8.4, } \deg(r) < k)
\end{aligned}$$

$$= \max_{p\in[0,1]} \left\| p^k \cdot \frac{\overline{(S_{2\alpha})}_\gamma(x)}{k^{k/2}\cdot|\mathcal{O}(\alpha)|} + r(p)\right\|_{sp} \cdot 2^{-O(q)}$$

$$\geq \frac{\|\overline{(S_{2\alpha})}_\gamma\|_{sp}}{k^{k/2}\cdot|\mathcal{O}(\alpha)|}\cdot 2^{-O(q+k)} \qquad \text{(Chebyshev Inequality - Lemma 8.7.5)}$$

where the last inequality follows by the following argument: one would like to show that there always exists $p \in [0,1]$ such that $\|p^k \cdot h_k(x) + \dots p^0 \cdot h_0(x))\|_{sp} \geq \|h_k(x)\|_{sp} \cdot 2^{-O(k)}$. So let p be such that $|p^k \cdot u^T M_k v + \dots p^0 \cdot u^T M_0 v| \geq |u^T M_k v| \cdot 2^{-O(k)}$ (such a p exists by Chebyshev inequality) where M_k is the matrix representation of $h_k(x)$ realizing $\|h_k\|_{sp}$ and u, v are the maximum singular vectors of M_k. $M_{k-1}, \dots, M_0$ are arbitrary matrix representations of $h_{k-1}, \dots h_0$ respectively. But $p^k \cdot M_k + \dots p^0 \cdot M_0$ is a matrix representation of $p^k \cdot h_k + \dots p^0 \cdot h_0$. Thus $\|p^k \cdot h_k + \dots p^0 \cdot h_0\|_{sp} \geq |u^T M_k v|/2^{-O(k)} = \|h_k\|_{sp} \cdot 2^{-O(q)}$.

This completes the proof as we've established Eq. (8.8). ■

8.8.7 Algorithms

It is straightforward to extract algorithms from the proofs of Theorem 8.8.7 and Theorem 8.8.8.

Non-negative coefficient polynomials

Let f be a degree-d polynomial with non-negative coefficients and let h be a $(d-2,2)$-folded polynomial whose unfolding yields f. Consider any q divisible by $2d$ and let $\overline{q} := (d-2)q/d$. Pick and return the best vector from the set

$$\left\{ \frac{\mathbf{1}}{\sqrt{n}} + \frac{\sqrt{\alpha}}{\sqrt{|\alpha|}} + \frac{\gamma}{\sqrt{|\gamma|}} + \arg\max \left\| h\left(\frac{\mathbf{1}}{\sqrt{n}} + \frac{\sqrt{\alpha}}{\sqrt{|\alpha|}} + \frac{\gamma}{\sqrt{|\gamma|}}\right)(x)\right\|_2 \;\middle|\; \alpha \in \mathbb{N}^n_{\leq \overline{q}/2}, \gamma \in \mathbb{N}^n_{\overline{q}/2-|\alpha|} \right\}$$

General Polynomials

Let f be a degree-d polynomial and let h be the unique $(d-2,2)$-folded polynomial whose unfolding yields f and such that for any $\beta \in \mathbb{N}^n_{d-2}$, the fold $\overline{h}_\beta$ of h is equal up to scaling, to the quadratic form of the corresponding $(n \times n)$ block of the SOS-symmetric matrix representation M_f of f. That is, for any $I, J \in [n]^{d/2-1}$, s.t. $\alpha(I) + \alpha(J) = \beta$,

$$\overline{h}_\beta(x) = \frac{x^T \mathsf{M}_f[I,J] x}{|\mathcal{O}(\beta)|}.$$

Consider any q divisible by $2d$ and let $\overline{q} := (d-2)q/d$. Let the set S be defined by,

$$S := \left\{ \Xi \cdot \frac{1}{\sqrt{|\gamma|}} \cdot \gamma \circ \frac{1}{2\alpha + \mathbf{1}} \circ b + \frac{\sqrt{\alpha}\circ\zeta}{\sqrt{|\alpha|}} \;\middle|\; \begin{array}{l} \Xi \in \Omega_{k+1},\ \zeta_i \in \Omega_{2\alpha_i+1},\ b \in \{0,1\}^n, \\ \alpha \in \mathbb{N}^n_{\leq \overline{q}/2},\ \gamma \in \{0,1\}^n_{\overline{q}-2|\alpha|} \end{array} \right\}$$

where Ω_p denotes the set of p-th roots of unity. Pick and return the best vector from the set

$$\left\{ c_1 \cdot y + c_2 \cdot \arg\max \|(h(y))(x)\|_2 \;\middle|\; y \in S,\; c_1 \in [-(d-2),(d-2)],\; c_2 \in [-2,2] \right\}$$

Note that one need only search through all roots of unity vectors ζ supported on $\operatorname{Supp}\gamma$ and all $\{0,1\}$-vectors b supported on $\operatorname{Supp}\alpha$. Lemma 8.7.7 can trivially be made constructive in time $2^{O(q)}$. Lastly, to go from complexes to reals, Lemma 8.7.8 can trivially be made constructive using $2^{O(d)}$ time. Thus the algorithm runs in time $n^{O(q)}$.

8.9 Constant Level Lower Bounds for Polynomials with Nonnegative Coefficients

Let $G = (V, E)$ be a random graph drawn from the distribution $G_{n,p}$ for $p \geq n^{-1/3}$. Let $\mathcal{C} \subseteq \binom{V}{4}$ be the set of 4-cliques in G. The polynomial f is defined as

$$f(x_1, \ldots, x_n) := \sum_{\{i_1,i_2,i_3,i_4\} \in \mathcal{C}} x_{i_1} x_{i_2} x_{i_3} x_{i_4}.$$

Clearly, f is multilinear and every coefficient of f is nonnegative. In this section, we prove the following two lemmas that establish a polynomial gap between $\|f\|_2$ and $\Lambda(f)$.

Lemma 8.9.1 (Soundness)**.** *With probability at least $1 - \frac{1}{n}$ over the choice of the graph G, we have $\|f\|_2 \leq n^2 p^6 \cdot (\log n)^{O(1)}$.*

Lemma 8.9.2 (Completeness)**.** *With probability at least $1 - \frac{1}{n}$ over the choice of the graph G, we have*

$$\Lambda(f) \;\geq\; \Omega\Big(\frac{n^{1/2} \cdot p}{\log^2 n}\Big)$$

when $p \in [n^{-1/3}, n^{-1/4}]$.

Note that the gap between the two quantities if $\tilde{\Omega}(n^{1/6})$ when $p = n^{-1/3}$, which is the choice we make.

8.9.1 Upper Bound on $\|f\|_2$

Reduction to counting shattered cliques

We say that an ordered 4-clique $(i_1, \ldots, i_4)$ is *shattered* by 4 disjoint sets $Z_1, \ldots, Z_4$ if for each $k \in [4]$, $i_k \in Z_k$. Let $Y_{j_1}, \ldots, Y_{j_4}$ be the sets containing the coordinates $i_1, \ldots, i_4$. Let $\mathcal{C}_G$ denote the set of (ordered) 4-cliques in G, and let $\mathcal{C}_G(Z_1, Z_2, Z_3, Z_4)$ denote the set of cliques shattered by $Z_1, \ldots, Z_4$.

We reduce the problem of bounding $\|f\|_2$, to counting shattered 4-cliques.

Claim 8.9.3. *There exist disjoint sets $Z_1, \ldots, Z_4 \subseteq [n]$ such that*

$$|\mathcal{C}_G(Z_1, Z_2, Z_3, Z_4)| \;\geq\; \left(\prod_{k=1}^{4} |Z_k|\right)^{1/2} \cdot O\left(\frac{\|f\|_2}{(\log n)^4}\right).$$

Proof. Let $x^* \in \mathbb{S}^{n-1}$ be the vector that maximizes f. Without loss of generality, assume that every coordinate of x^* is nonnegative. Let y^* be another unit vector defined as

$$y^* \;:=\; \frac{(x^* + \mathbf{1}/\sqrt{n})}{\|x^* + \mathbf{1}/\sqrt{n}\|_2}.$$

Since both x^* and $\frac{\mathbf{1}}{\sqrt{n}}$ are unit vectors, the denominator is at most 2. This implies that $f(y^*) \geq \frac{f(x^*)}{2^4}$, and each coordinate of y^* is at least $\frac{1}{2\sqrt{n}}$. For $1 \leq j \leq \log_2 n$, let Y_j be the set

$$Y_j \;:=\; \left\{ i \in [n] \;\mid\; 2^{-j} < y_i^* \leq 2^{-(j-1)} \right\}.$$

The sets $Y_1, \ldots, Y_{\log_2 n}$ partition $[n]$. Since $1 = \sum_{i \in [n]} y_i^2 > |Y_j| \cdot 2^{-2j}$, we have for each j, $|Y_j| \leq 2^{2j}$. Let Z_1, Z_2, Z_3, and Z_4 be pairwise disjoint random subsets of $[n]$ chosen as follows:

- Randomly partition each Y_j to $Y_{j,1}, \ldots, Y_{j,4}$ where each element of Y_j is put into exactly one of $Y_{j,1}, \ldots, Y_{j,4}$ uniformly and independently.
- Sample $r_1, \ldots, r_4$ independently and randomly from $\{1, \ldots, \log_2 n\}$.
- For $k = 1, \ldots, 4$, take $Z_k := Y_{r_k,k}$

We use $\mathcal{P}$ to denote random partitions $\left\{ (Y_{j,1}, \ldots, Y_{j,4}) \right\}_{j \in [\log n]}$ and r to denote the random choices $r_1, \ldots, r_4$. Note that the events $i_k \in Z_k$ are independent for different k, and that $Z_1, \ldots, Z_4$ are independent given $\mathcal{P}$. Thus, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathcal{P},r}\left[\frac{\mathbb{1}\left[(i_1, i_2, i_3, i_4) \text{ is shattered}\right]}{\sqrt{|Z_1|\,|Z_2|\,|Z_3|\,|Z_4|}}\right] &= \mathop{\mathbb{E}}_{\mathcal{P}}\left[\prod_{k=1}^{4} \mathop{\mathbb{E}}_{r_k}\left[\frac{\mathbb{1}\left[i_k \in Z_k\right]}{\sqrt{|Z_k|}}\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathcal{P}}\left[\prod_{k=1}^{4} \mathop{\mathbb{E}}_{r_k}\left[\frac{\mathbb{1}\left[r_k = j_k\right] \cdot \mathbb{1}\left[i_k \in Y_{j_k,k}\right]}{\sqrt{|Y_{j_k,k}|}}\right]\right] \\
&\geq \mathop{\mathbb{E}}_{\mathcal{P}}\left[\prod_{k=1}^{4} \mathop{\mathbb{E}}_{r_k}\left[\frac{\mathbb{1}\left[r_k = j_k\right] \cdot \mathbb{1}\left[i_k \in Y_{j_k,k}\right]}{\sqrt{|Y_{j_k}|}}\right]\right] \\
&= \mathop{\mathbb{E}}_{\mathcal{P}}\left[\prod_{k=1}^{4} \left(\frac{1}{\log n} \cdot \frac{\mathbb{1}\left[i_k \in Y_{j_k,k}\right]}{\sqrt{|Y_{j_k}|}}\right)\right]
\end{aligned}$$

$$
\begin{aligned}
&= \mathop{\mathbb{E}}_{\mathcal{P}}\left[\prod_{k=1}^{4}\left(\frac{1}{\log n}\cdot\frac{\mathbb{1}\left[i_k \in Y_{j_k,k}\right]}{\sqrt{\left|Y_{j_k}\right|}}\right)\right] \\
&= \frac{1}{(4\log n)^4}\cdot\frac{1}{\sqrt{\left|Y_{j_1}\right|\left|Y_{j_2}\right|\left|Y_{j_3}\right|\left|Y_{j_4}\right|}} \\
&\geq \frac{1}{(4\log n)^4}\cdot 2^{j_1+j_2+j_3+j_4} \\
&\geq \frac{1}{(8\log n)^4}\cdot y^*_{i_1}y^*_{i_2}y^*_{i_3}y^*_{i_4}\,.
\end{aligned}
$$

Then, by linearity of expectation,

$$
\begin{aligned}
\mathop{\mathbb{E}}_{\mathcal{P},r}\left[\frac{|\mathcal{C}_G(Z_1,Z_2,Z_3,Z_4)|}{\sqrt{|Z_1|\,|Z_2|\,|Z_3|\,|Z_4|}}\right] &\geq \frac{1}{(8\log n)^4}\cdot\sum_{(i_1,\ldots,i_4)\in\mathcal{C}_G} y^*_{i_1}y^*_{i_2}y^*_{i_3}y^*_{i_4} \\
&= \frac{4!}{(8\log n)^4}\cdot f\left(y^*\right) \\
&\geq \frac{4!}{(16\log n)^4}\cdot f\left(x^*\right) \;=\; \frac{4!}{(16\log n)^4}\cdot\|f\|_2\,,
\end{aligned}
$$

which proves the claim. ■

We will show that with high probability, G satisfies the property that every four disjoint sets $Z_1,\ldots,Z_4 \subseteq V$ shatter at most $O\left(\sqrt{|Z_1||Z_2||Z_3||Z_4|}\cdot n^2p^6\cdot(\log n)^{O(1)}\right)$ cliques, proving Lemma 8.9.1.

Counting edges and triangles

For a vertex $i \in [n]$, we use $\mathsf{N}(i)$ to denote the set of vertices in the graph G. For ease of notation, we use $a \lesssim b$ to denote $a \leq b\cdot(\log n)^{O(1)}$. We first collect some simple consequences of Chernoff bounds.

Claim 8.9.4. *Let $G \sim G_{n,p}$ with $p \geq n^{-1/3}$. Then, with probability $1-\frac{1}{n}$, we have*

- *For all distinct $i_1, i_2 \in [n]$, $|\mathsf{N}(i_1)\cap\mathsf{N}(i_2)| \lesssim np^2$.*
- *For all distinct $i_1, i_2, i_3 \in [n]$, $|\mathsf{N}(i_1)\cap\mathsf{N}(i_2)\cap\mathsf{N}(i_3)| \lesssim np^3$.*
- *For all sets $S_1, S_2 \subseteq [n]$, $|E(S_1,S_2)| \lesssim \max\{|S_1|\,|S_2|\,p, |S_1|+|S_2|\}$.*

We also need the following bound on the number of triangles shattered by three disjoint sets S_1, S_2 and S_3, denoted by $\Delta_G(S_1,S_2,S_3)$. As for 4-cliques, a triangle is said to be shattered if it has exactly one vertex in each the sets.

Claim 8.9.5. *Let $G \sim G_{n,p}$ with $p \geq n^{-1/3}$. Then, with probability $1 - \frac{1}{n}$, for all disjoint sets $S_1, S_2, S_3 \subseteq [n]$*

$$|\Delta_G(S_1, S_2, S_3)| \lesssim |S_3| + |E(S_1, S_2)| \cdot \left(np^3 \cdot |S_3|\right)^{1/2} .$$

Proof. With probability at least $1 - \frac{1}{n}$, G satisfies the conclusion of Claim 8.9.4. Fix such a G, and consider arbitrary subsets $S_1, S_2, S_3 \subseteq V$. Consider the bipartite graph H where the left side vertices correspond to edges in $E(S_1, S_2)$, the right side vertices correspond to vertices in S_3, and there is an edge from $(i_1, i_2) \in E(S_1, S_2)$ to $i_3 \in S_3$ when both $(i_1, i_3), (i_2, i_3) \in E$. Clearly, $|\Delta_G(S_1, S_2, S_3)|$ is equal to the number of edges in H.

Consider two different edges $(i_1, i_2), (i_1', i_2') \in E(S_1, S_2)$. These two edges are incident on at least 3 distinct vertices, say $\{i_1, i_2, i_1'\}$. Hence, the number of vertices $i_3 \in [n]$ that are adjacent to all $\{i_1, i_2, i_1', i_2'\}$ in G is at most $|\mathsf{N}(i_1) \cap \mathsf{N}(i_2) \cap \mathsf{N}(i_1')| \lesssim np^3$. This gives that the number of pairs triangles sharing a common vertex in S_3 is at most $|E(S_1, S_2)|^2 \cdot np^3 (\log n)^{O(1)}$.

Let $d_H(i_3)$ denote the degree of a vertex i_3 in H, and let Δ denote the number of shattered triangles. Counting the above pairs of triangles using the degrees gives

$$\sum_{i_3 \in S_3} \binom{d_H(i_3)}{2} \lesssim |E(S_1, S_2)|^2 \cdot np^3 .$$

An application of Cauchy-Schwarz gives

$$\Delta^2 - \Delta \cdot |S_3| \lesssim |S_3| \cdot |E(S_1, S_2)|^2 \cdot np^3 ,$$

which proves the claim. ■

Bounding 4-clique Density

Let $G \sim G_{n,p}$ be a graph satisfying the conclusions of Claims Claim 8.9.4 and Claim 8.9.5. Let $S_1, \ldots, S_4 \subseteq [n]$ be disjoint sets with sizes $n_1 \leq n_2 \leq n_3 \leq n_4$. We consider two cases:

- **Case 1**: $|E(S_1, S_2)| \lesssim n_1 n_2 p$
  Note that each edge (i_1, i_2) can only participate in at most $|\mathsf{N}(i_1) \cap \mathsf{N}(i_2)|$ triangles, and each triangle (i_1, i_2, i_3) can only be extended to at most $|\mathsf{N}(i_1) \cap \mathsf{N}(i_2) \cap \mathsf{N}(i_3)|$ 4-cliques. Thus, Claim Claim 8.9.4 gives

$$|\mathcal{C}_G(S_1, S_2, S_3, S_4)| \lesssim n_1 n_2 p \cdot np^2 \cdot np^3 \lesssim (n_1 n_2 n_3 n_4)^{1/2} \cdot n^2 p^6 .$$

- **Case 2**: $|E(S_1, S_2)| \lesssim n_1 + n_2$
  Claim Claim 8.9.5 gives

$$|\Delta_G(S_1, S_2, S_3)| \lesssim n_3 + (n_1 + n_2) \cdot \left(n_3 \cdot np^3\right)^{1/2} ,$$

which together with Claim Claim 8.9.4 implies

$$|\mathcal{C}_G(S_1, S_2, S_3, S_4)| \lesssim n_3 \cdot np^3 + (n_1 + n_2) \cdot n_3^{1/2} \cdot \left(np^3\right)^{3/2} .$$

Considering the first term, we note that

$$n_3 \cdot np^3 \leq (n_3 n_4)^{1/2} \cdot n^2 p^6 \leq (n_1 n_2 n_3 n_4)^{1/2} \cdot n^2 p^6 ,$$

since $n_3 \leq n_4$ and $np^3 \geq 1$. Similarly, for the second term, we have

$$(n_1 + n_2) \cdot n_3^{1/2} \cdot \left(np^3\right)^{3/2} \leq 2(n_2 n_3 n_4)^{1/2} \cdot \left(np^3\right)^{3/2} \leq 2 \cdot (n_1 n_2 n_3 n_4)^{1/2} \cdot n^2 p^6 .$$

Combined with Claim Claim 8.9.3, this completes the proof of Lemma 8.9.1.

8.9.2 Lower Bound on $\Lambda(f)$

Recall that given a random graph $G = ([n], E)$ drawn from the distribution $G_{n,p}$, the polynomial f is defined as

$$f(x_1, \ldots, x_n) := \sum_{\{i_1,i_2,i_3,i_4\} \in \mathcal{C}} x_{i_1} x_{i_2} x_{i_3} x_{i_4},$$

where $\mathcal{C} \subseteq \binom{[n]}{4}$ is the set of 4-cliques in G. Let $\mathsf{A} \in \mathbb{R}^{[n]^2 \times [n]^2}$ be the natural matrix representation of $24f$ (corresponding to ordered copies of cliques) with

$$\mathsf{A}[(i_1, i_2), (i_3, i_4)] = \begin{cases} 1 & \text{if } \{i_1, \ldots, i_4\} \in \mathcal{C} \\ 0 & \text{otherwise} \end{cases}$$

Let $E' \subseteq [n]^2$ be the set of ordered edges i.e., $(i_1, i_2) \in E'$ if and only if $\{i_1, i_2\} \in E$. Note that $|E'| = 2m$ where m is the number of edges in G. All nonzero entries of A are contained in the principal submatrix $\mathsf{A}_{E'}$, formed by the rows and columns indexed by E'.

A simple lower bound on $\|f\|_{sp}$

We first give a simple proof that $\|f\|_{sp} \geq \sqrt{n^2 p^5}$ with high probability.

Lemma 8.9.6. *$\|f\|_{sp} \geq \Omega(\sqrt{n^2 p^5}) = \Omega(n^{1/6})$ with high probability.*

Proof. Consider any matrix representation M of $24f$ and its principal submatrix $\mathsf{M}_{E'}$. It is easy to observe that the Frobenius norm of $M_{E'}$ satisfies $\|M_{E'}\|_F^2 \geq 24\,|\mathcal{C}|$, minimized when $M = \mathsf{A}$. Since $\|M_{E'}\|_F^2 \leq |E'| \cdot \|A_{E'}\|_2^2$, we have that with high probability,

$$\|A\|_2 \geq \|A_{E'}\|_2 \geq \sqrt{\frac{24\,|\mathcal{C}|}{2|E|}} = \Omega\left(\frac{\sqrt{n^4 p^6}}{\sqrt{n^2 p}}\right) = \Omega\left(\sqrt{n^2 p^5}\right).$$

■

Lower bound for the stronger relaxation computing $\Lambda(f)$

We now prove Lemma 8.9.2, which says that $\Lambda(f) \geq \frac{n^{1/6}}{\log^2 n}$ with high probability. In order to show a lower bound, we look at the dual SDP for computing $\Lambda(f)$, which is a maximization problem over positive semidefinite, SoS-symmetric matrices M with $\mathsf{Tr}(\mathsf{M}) = 1$. We exhibit such a matrix $\mathsf{M} \in \mathbb{R}^{[n]^2 \times [n]^2}$ for which the value of the objective $\langle \mathsf{A}, \mathsf{M} \rangle$ is large.

For large $\langle \mathsf{A}, \mathsf{M} \rangle$, one natural attempt is to take M to be A and modify it to satisfy other conditions. Note that A is already SoS-symmetric. However, $\mathsf{Tr}(\mathsf{A}) = 0$, which implies that the minimum eigenvalue is negative.

Let $\lambda_{\min}$ be the minimum eigenvalue of A, which is also the minimum eigenvalue of $\mathsf{A}_{E'}$. Let $I_{E'} \in \mathbb{R}^{[n]^2 \times [n]^2}$ be such that $I[(i_1, i_2), (i_1, i_2)] = 1$ if $(i_1, i_2) \in E'$ and all other entries are 0. Note that $I_{E'}$ is a diagonal matrix with $\mathsf{Tr}(I_{E'}) = 2m$. Adding $-\lambda_{\min} \cdot I_{E'}$ to A makes it positive semidefinite, so setting

$$\mathsf{M} \;=\; \frac{\mathsf{A} - \lambda_{\min} I_{E'}}{\mathsf{Tr}(\mathsf{A} - \lambda_{\min} I_{E'})} \;=\; \frac{\mathsf{A} - \lambda_{\min} I_{E'}}{-2m\lambda_{\min}} \;=\; \frac{\mathsf{A} + |\lambda_{\min}| \cdot I_{E'}}{2m \cdot |\lambda_{\min}|} \tag{8.9}$$

makes sure that M is positive semidefinite, $\mathsf{Tr}(\mathsf{M}) = 1$, and $\langle \mathsf{A}, \mathsf{M} \rangle = \frac{12|\mathcal{C}|}{m \cdot |\lambda_{\min}|}$ (each 4-clique in $\mathcal{C}$ contributes 24). Since $|\mathcal{C}| = \Theta(n^4 p^6)$ and $m = \Theta(n^2 p)$ with high probability, if $|\lambda_{\min}| = O(np^{5/2})$, $\langle \mathsf{A}, \mathsf{M} \rangle = \Theta\left(n^2 p^{5/2}\right)$, which is $\Omega(n^{1/6})$ when $p = \Omega(n^{-1/3})$.

The M defined in Eq. (8.9) does not directly work since it is not SoS-symmetric. However, the following claim proves that this issue can be fixed by losing a factor 2 in $\langle \mathsf{A}, \mathsf{M} \rangle$.

Claim 8.9.7. *There exists* M *such that it is SoS-symmetric, positive semidefinite with* $\mathsf{Tr}(\mathsf{M}) = 1$, *and* $\langle \mathsf{A}, \mathsf{M} \rangle \geq \frac{6|\mathcal{C}|}{m \cdot |\lambda_{\min}|}$.

Proof. Let $Q_{E'} \in \mathbb{R}^{[n]^2 \times [n]^2}$ be the matrix such that

- For $(i_1, i_2) \in E'$, $Q_{E'}[(i_1, i_1), (i_2, i_2)] = Q_{E'}[(i_2, i_2), (i_1, i_1)] = 1$.
- For $i \in [n]$, $Q_{E'}[(i, i), (i, i)] = \deg_G(i)$, where $\deg_G(i)$ denotes the degree of i in G.
- All other entries are 0.

We claim that $I_{E'} + Q_{E'}$ is SoS-symmetric: $(I_{E'} + Q_{E'})[(i_1, i_2), (i_3, i_4)]$ has a nonzero entry if and only if $i_1 = i_2 = i_3 = i_4$ or two different numbers j_1, j_2 appear exactly twice and $(j_1, j_2) \in E'$ (in this case $(I_{E'} + Q_{E'})[(i_1, i_2), (i_3, i_4)] = 1$). Since A is SoS-symmetric, so $\mathsf{A} + |\lambda_{\min}| \cdot (I_{E'} + Q_{E'})$ is also SoS-symmetric.

It is easy to see that $Q_{E'}$ is diagonally dominant, and hence positive semidefinite. Since we already argued that $\mathsf{A} + |\lambda_{\min}| \cdot I_{E'}$ is positive semidefinite, $\mathsf{A} + |\lambda_{\min}| \cdot (I_{E'} + Q_{E'})$ is also positive semidefinite. Also, $\mathsf{Tr}(Q_{E'}) = \sum_{i \in [n]} \deg_G(i) = 2m$. Thus, we take

$$\mathsf{M} \;=\; \frac{\mathsf{A} + |\lambda_{\min}| \cdot (I_{E'} + Q_{E'})}{\mathsf{Tr}(\mathsf{A} + |\lambda_{\min}| \cdot (I_{E'} + Q_{E'}))} \;=\; \frac{\mathsf{A} + |\lambda_{\min}| \cdot I_{E'}}{4m \cdot |\lambda_{\min}|}.$$

By the above arguments, we have that M that is PSD, SoS-symmetric with $\mathrm{Tr}(\mathsf{M}) = 1$, and

$$\langle \mathsf{A}, \mathsf{M} \rangle = \frac{6|\mathcal{C}|}{m \cdot |\lambda_{\min}|}$$

as desired. ■

It only remains to bound $\lambda_{\min}$, which is the minimum eigenvalue of A and $\mathsf{A}_{E'}$. For p in the range $[n^{-1/3}, n^{-1/4}]$, we will show a bound of $\tilde{O}(n^{3/2}p^4)$ below, which when combined with the above claim, completes the proof of Lemma 8.9.2.

Bounding the smallest eigenvalue via the trace method

Our estimate $|\lambda_{\min}| = O(np^{5/2})$ is based on the following observation: $\mathsf{A}_{E'}$ is a $2m \times 2m$ random matrix where each row and column is expected to have $\Theta(n^2p^5)$ ones (the expected number of 4-cliques an edge participates in). An adjacency matrix of a random graph with average degree d has a minimum eigenvalue $-\Theta(\sqrt{d})$, hence the estimate $|\lambda_{\min}| = O(np^{5/2})$. Even though $\mathsf{A}_{E'}$ is not sampled from a typical random graph model (and even E' is a random variable), we will be able to prove the following weaker estimate, which suffices for our purposes.

Lemma 8.9.8. *With high probability over the choice of the graph G, we have*

$$|\lambda_{\min}| \;=\; \begin{cases} \tilde{O}\left(n^{3/2} \cdot p^4\right) & \text{for } p \in \left[n^{-1/3}, n^{-1/4}\right] \\ \tilde{O}\left(n^{5/3} \cdot p^{14/3}\right) & \text{for } p \in \left[n^{-1/4}, 1/2\right] \end{cases}$$

Proof. Instead of $\mathsf{A}_{E'}$, we directly study A to bound $\lambda_{\min}$. For simplicity, we consider the following matrix $\widehat{\mathsf{A}}$, where each row and column is indexed by an unordered pair $\{i,j\} \in \binom{[n]}{2}$, and $\widehat{\mathsf{A}}[\{i_1,i_2\},\{i_3,i_4\}] = 1$ if and only if i_1, i_2, i_3, i_4 form a 4-clique. A has only zero entries in the rows or columns indexed by (i,i) for all $i \in [n]$, and for two pairs $i_1 \neq i_2$ and $i_3 \neq i_4$, we have

$$\begin{aligned} \widehat{\mathsf{A}}[\{i_1,i_2\},\{i_3,i_4\}] \;:=\; & \frac{1}{4} \cdot \{\mathsf{A}[(i_1,i_2),(i_3,i_4)] + \mathsf{A}[(i_1,i_2),(i_4,i_3)]\} \\ & + \frac{1}{4} \cdot \{\mathsf{A}[(i_2,i_1),(i_3,i_4)] + \mathsf{A}[(i_2,i_1),(i_4,i_3)]\} \,. \end{aligned}$$

Therefore, $|\lambda_{\min}(\mathsf{A})| \leq 4 \cdot \left|\lambda_{\min}\left(\widehat{\mathsf{A}}\right)\right|$ and it suffices to bound the minimum eigenvalue of $\widehat{\mathsf{A}}$. We consider the matrix $\widehat{\mathsf{N}}_E := \widehat{\mathsf{A}} - p^4 \cdot \widehat{\mathsf{J}}_E$, where $\widehat{\mathsf{J}}_E \in \mathbb{R}^{\binom{[n]}{2} \times \binom{[n]}{2}}$ is such that

$$\widehat{\mathsf{J}}_E[\{i_1,i_2\},\{i_3,i_4\}] \;=\; \begin{cases} 1 & \text{if } \{i_1,i_2\},\{i_3,i_4\} \in E \\ 0 & \text{otherwise} \end{cases} .$$

Since $\widehat{\mathsf{J}}_E$ is a rank-1 matrix with a positive eigenvalue, the minimum eigenvalues of $\widehat{\mathsf{A}}$ and $\widehat{\mathsf{N}}_E$ are the same. In summary, $\widehat{\mathsf{N}}_E$ is the following matrix.

$$\widehat{\mathsf{N}}_E[\{i_1,i_2\},\{i_3,i_4\}] \;=\; \begin{cases} 1-p^4 & \text{if } \{i_1,i_2,i_3,i_4\}\in\mathcal{C} \\ -p^4 & \text{if } \{i_1,i_2,i_3,i_4\}\notin\mathcal{C} \text{ but } \{i_1,i_2\},\{i_3,i_4\}\in E \\ 0 & \text{otherwise} \end{cases}$$

We use the trace method to bound $\|\widehat{\mathsf{N}}_E\|_2$, based on the observation that for every even $r\in\mathbb{N}$, $\|\widehat{\mathsf{N}}_E\|_2 \le \left(\mathsf{Tr}\left((\widehat{\mathsf{N}}_E)^r\right)\right)^{1/r}$. Fix an even $r\in\mathbb{N}$. The expected value of the trace can be represented as

$$\mathbb{E}\left[\mathsf{Tr}\left((\widehat{\mathsf{N}}_E)^r\right)\right] \;=\; \mathbb{E}\left[\sum_{I^1,\dots,I^r\in\binom{[n]}{2}}\prod_{k=1}^{r}\widehat{\mathsf{N}}_E[I^k,I^{k+1}]\right] \;=\; \sum_{I^1,\dots,I^r\in\binom{[n]}{2}}\mathbb{E}\left[\prod_{k=1}^{r}\widehat{\mathsf{N}}_E[I^k,I^{k+1}]\right]$$

where each $I^j = \left\{i_1^j, i_2^j\right\} \in \binom{[n]}{2}$ is an edge of the complete graph on n vertices (call it a *potential edge*) and $I^{r+1} := I^1$.

Fix r potential edges $I^1,\dots,I^r$, let $t := \prod_{k=1}^{r}\widehat{\mathsf{N}}_E[I^k,I^{k+1}]$, and consider $\mathbb{E}[t]$. Let $E_0 := \{I^1,\dots,I^r\}$ be the set of distinct edges represented by $I^1,\dots,I^r$. Note that the expected value is 0 if one of I^j does not become an edge. Therefore, $\mathbb{E}[t] = p^{|E_0|}\cdot\mathbb{E}[t \mid E_0\subseteq E]$.

Let $D\subseteq[r]$ be the set of $j\in[r]$ such that all four vertices in I^j and I^{j+1} are distinct i.e.,

$$D \;:=\; \left\{j\in[r] \;\middle|\; \left|\left\{i_1^j, i_2^j, i_1^{j+1}, i_2^{j+1}\right\}\right| = 4\right\}.$$

For $j\in[r]\setminus D$, $\{i_1^j, i_2^j, i_1^{j+1}, i_2^{j+1}\}$ cannot form a 4-clique, so given that $I^j, I^{j+1}\in E$, we have $\widehat{\mathsf{N}}_E[I^j,I^{j+1}] = -p^4$. For $j\in D$, let $E_j := \left\{\left\{i_1^j,i_1^{j+1}\right\},\left\{i_1^j,i_2^{j+1}\right\},\left\{i_2^j,i_1^{j+1}\right\},\left\{i_2^j,i_2^{j+1}\right\}\right\}\setminus E_0$ be the set of edges in the 4-clique created by $\left\{i_1^j, i_2^j, i_1^{j+1}, i_2^{j+1}\right\}$ except ones in E_0. Then

$$\mathbb{E}[t] \;=\; p^{|E_0|}\cdot\mathbb{E}[t|E_0\subseteq E] \;=\; p^{|E_0|}\cdot(-p^4)^{r-|D|}\cdot\mathbb{E}\left[\prod_{k\in D}\widehat{\mathsf{N}}_E[I^k,I^{k+1}] \;\middle|\; E_0\subseteq E\right].$$

Suppose there exists $j\in D$ such that $\left|E_j\right| = 4$ and $E_j\cap\left(\cup_{j'\in D\setminus\{j\}}E_{j'}\right) = \emptyset$. Then, given that $E_0\subseteq E$, $\widehat{\mathsf{N}}_E[I^j,I^{j+1}]$ is independent of all $\left\{\widehat{\mathsf{N}}_E[I^k,I^{k+1}]\right\}_{k\in D\setminus\{j\}}$, and

$$\mathbb{E}\left[\widehat{\mathsf{N}}_E[I^j,I^{j+1}]|E_0\subseteq E\right] \;=\; p^4(1-p^4)+(1-p^4)(-p^4) \;=\; 0\,.$$

Therefore, $\mathbb{E}[t]=0$ unless for all $j\in D$, either $|E_j|\le 3$ or there exists $j'\in D\setminus\{j\}$ with $E_j\cap E_{j'}\neq\emptyset$.

Let $E_D := \bigcup_{j\in D}E_j$. Note that E_0 and E_D completely determines t. $\mathbb{E}[t]$ can be written as

$$p^{|E_0|}\cdot(-p^4)^{r-|D|}\cdot\mathbb{E}\left[\prod_{k\in D}\widehat{\mathsf{N}}_E[I^k,I^{k+1}] \;\middle|\; E_0\subseteq E\right]$$

$$= p^{|E_0|} \cdot (-p^4)^{r-|D|} \cdot \sum_{F \subseteq E_D} \left(p^{|F|}(1-p)^{|E_D|-|F|} \cdot \mathbb{E}\left[\prod_{k \in D} \widehat{\mathsf{N}}_E[I^k, I^{k+1}] \mid E_0 \subseteq E, E_D \cap E = F \right] \right)$$
$$= p^{|E_0|} \cdot (-p^4)^{r-|D|} \cdot \sum_{F \subseteq E_D} \left(p^{|F|}(1-p)^{|E_D|-|F|} \cdot (1-p^4)^{|D|-a(F)}(-p^4)^{a(F)} \right),$$

where $a(F)$ denotes the number of $j \in D$ with $E_j \not\subseteq F$. Since $E_D \subseteq F \cup \left(\bigcup_{j: E_j \not\subseteq F} E_j) \right)$ and $4a(F) + |F| \geq |E_D|$, we have

$$\begin{aligned} \mathbb{E}[t] &= p^{|E_0|} \cdot (-p^4)^{r-|D|} \cdot \sum_{F \subseteq E_D} \left(p^{|F|}(1-p)^{|E_D|-|F|} \cdot (1-p^4)^{|D|-a(F)}(-p^4)^{a(F)} \right) \\ &\leq p^{|E_0|} \cdot (p^4)^{r-|D|} \cdot 2^{|E_D|} \cdot p^{|E_D|} \\ &\leq 2^{4r} \cdot p^{4(r-D)+|E_0|+|E_D|}. \end{aligned}$$

We now count the number of terms which contribute to the sum. Fix a graph H with r labelled edges $I^1, \ldots, I^r$ (possibly repeated) and $q := q(H)$ vertices, without any isolated vertex (so $q \leq 2r$). There are at most $\binom{q}{2}^r \leq (2r)^{2r}$ such graphs. Then $I^1, \ldots, I^r$, as edges in $\binom{[n]}{2}$, are determined by a map $V_H \to [n]$. There are at most n^q such mappings.

Let $E_0 := E_0(H), D := D(H), E_j := E_j(H), E' := E'(H)$ be defined as before. Note that E_0 is set the edges of H. As observed before, the contribution from H is 0 if there exists $j \in D$ such that $|E_j| = 4$ and E_j is disjoint from $\{E_{j'}\}_{j' \in D \setminus \{j\}}$. Let $\mathcal{H}$ be the set of H that has nonzero contribution. Then,

$$\begin{aligned} \mathbb{E}\left[\mathsf{Tr}\left(\left(\widehat{\mathsf{N}}_E \right)^r \right) \right] &= \sum_{I^1, \ldots, I^r \in \binom{[n]}{2}} \mathbb{E}\left[\prod_{k=1}^r \widehat{\mathsf{N}}_E[I^k, I^{k+1}] \right] \\ &\leq \sum_{H \in \mathcal{H}} n^{q(H)} \cdot 2^{4r} \cdot p^{4(r-D(H))+|E_0(H)|+|E_D(H)|} \\ &\leq (2r)^{2r} \cdot \max_{H \in \mathcal{H}} \left(n^{q(H)} 2^{4r} \cdot p^{4(r-D(H))+|E_0(H)|+|E_D(H)|} \right) \\ &\leq (8r)^{2r} \cdot \max_{H \in \mathcal{H}} \left(n^{q(H)} p^{4(r-D(H))+|E_0(H)|+|E_D(H)|} \right) \end{aligned}$$

We will prove the following bound on the maximum contribution of any $H \in \mathcal{H}$.

Claim 8.9.9. *Let $\mathcal{H}$ be defined as above. Then, for all $H \in \mathcal{H}$, we have*

$$n^{q(H)} p^{4(r-D(H))+|E_0(H)|+|E_D(H)|} \leq n^2 \cdot B_p^r,$$

where

$$B_p = \begin{cases} n^{3/2} \cdot p^4 & \text{for } p \in \left[n^{-1/3}, n^{-1/4}\right] \\ n^{5/3} \cdot p^{14/3} & \text{for } p \in \left[n^{-1/4}, 1/2\right] \end{cases}.$$

Using the above claim, we can bound $\mathbb{E}[\mathsf{Tr}\left(\left(\widehat{\mathsf{N}}_E \right)^r \right)]$ as

$$\mathbb{E}\left[\mathsf{Tr}\left(\left(\widehat{\mathsf{N}}_E \right)^r \right) \right] \leq (8r)^{2r} \cdot \max_{H \in \mathcal{H}} \left(n^{q(H)} p^{4(r-D(H))+|E_0(H)|+|E_D(H)|} \right)$$

$$\leq \; (8r)^{2r} \cdot n^2 \cdot B_p^r \,,$$

where B_p is given by Claim Claim 8.9.9 for different ranges of p. By Markov's inequality, we get that with probability $1 - \frac{1}{n}$, we have $\mathrm{Tr}\left(\left(\widehat{\mathsf{N}}_E\right)^r\right) \leq (8r)^{2r} \cdot n^3 \cdot B_p^r$, which gives

$$\|\widehat{\mathsf{N}}_E\|_2 \;\leq\; (8r)^2 \cdot B_p \cdot n^{3/r} \,.$$

Choosing $r = \Theta(\log n)$ then proves the lemma. ■

It remains to prove Claim Claim 8.9.9.

Analyzing contributing subgraphs

Recall that graphs $H \in \mathcal{H}$ were constructed from edges $\{I^1, \ldots, I^r\}$, with edge I^j consisting of vertices $\{i_1^j, i_2^j\}$. Also, we define $q(H) = |V(H)|$. Moreover, we defined the following sets for graph H

$$\begin{aligned}
E_0(H) \;&:=\; \{I^1, \ldots, I^r\} \quad \text{(counting only distinct edges)} \\
D(H) \;&:=\; \left\{ j \in [r] \;\middle|\; \left|\left\{ i_1^j, i_2^j, i_1^{j+1}, i_2^{j+1} \right\}\right| = 4 \right\} \\
E_j(H) \;&:=\; \left\{ \{i_1^j, i_1^{j+1}\}, \{i_1^j, i_2^{j+1}\}, \{i_2^j, i_1^{j+1}\}, \{i_2^j, i_2^{j+1}\} \right\} \setminus E_0(H) \\
E_D(H) \;&:=\; \bigcup_{j \in D} E_j(H)
\end{aligned}$$

Moreover, the graph H is in $\mathcal{H}$ only if for every $j \in D$, either $\left|E_j(H)\right| \leq 3$ or there exists $j' \in D \setminus \{j\}$ such that $E_j(H) \cap E_{j'}(H) \neq \emptyset$. Claim Claim 8.9.9 then follows from the following combinatorial claim (taking $b = \log(1/p)/\log n$).

Claim 8.9.10. *Any graph $H \in \mathcal{H}$ satisfies, for all $b \in [0, 1/3]$*

$$q(H) \;\leq\; 2 + b \cdot (4(r - |D(H)|) + |E_0(H)| + |E_D(H)|) + c \cdot r \,,$$

where $c = 5/3 - 14b/3$ for $b \in [0, 1/4]$ and $c = 3/2 - 4b$ for $b \in [1/4, 1/3]$.

Proof. Fix a graph $H \in \mathcal{H}$. Let $j = 1, \ldots, r$, let $V_j := \left\{ i_1^j, i_2^j \right\}_{1 \leq j \leq r}$ (i.e., the set of vertices covered by $I^1, \ldots, I^j$). For each $j = 2, \ldots, r$, let $v_j := |V_j| - |V_{j-1}|$ and classify the index j to one of the following types.

- Type -1: $I^j \cap I^{j-1} \neq \emptyset$ (equivalently, $j - 1 \notin D$).
- Type k ($0 \leq k \leq 2$): I^j and I^{j-1} are disjoint, and $v_j = k$ (i.e., adding I^j introduces k new vertices).

Let T_k ($-1 \le k \le 2$) be the set of indices of Type k, and let $t_k := |T_k|$. The number of vertices q is bounded by

$$q \ \le \ 2 + 1 \cdot t_{-1} + 0 \cdot t_0 + 1 \cdot t_1 + 2 \cdot t_2 \ = \ 2 + t_{-1} + t_1 + 2t_2.$$

Let H_j be the graph with V_j as vertices and edges

$$E(H_j) \ = \ \left\{ I^1, \ldots, I^j \right\} \bigcup \left(\bigcup_{k \in D \cap [j-1]} E_k \right).$$

For $j = 2, \ldots, r$, let $e_j = \left|E(H_j)\right| - \left|E(H_{j-1})\right|$. For an index $j \in T_2$, adding two vertices i_1^j, i_2^j introduces at least 5 edges in H_j compared to H_{j-1} (i.e., six edges in the 4-clique on $\left\{ i_1^{j-1}, i_2^{j-1}, i_1^j, i_2^j \right\}$ except I^{j-1}), so $e_j \ge 5$. Similarly, we get $e_j \ge 3$ for each $j \in T_1$.

The lemma is proved via the following charging argument. For each index $j = 2, \ldots, r$, we get value b for each edge in $H_j \setminus H_{j-1}$ and get value c for the new index. If $j \in T_{-1}$, we get an additional value of $4b$. We give this value to vertices in $V_j \setminus V_{j-1}$. If we do not give more value than we get and each vertex in $V(H) \setminus V_1$ gets more than 1, this means

$$q - 2 \ \le \ b \cdot (|E_0| + |E_D| + 4(r - |D(H)|)) + c \cdot r,$$

proving the claim. For example, if j is an index of Type 1, it gets a value at least $3b + c$ and needs to give value 1, such a charging can be done if $3b + c \ge 1$. Similarly, a type 0 vertex does not need to give any value and has a surplus. We will choose parameters so that each j of types -1, 1 or 0 can distribute the value to vertices added in $V_j \setminus V_{j-1}$. However, if j is an index of Type 2, it needs to distribute the value it gets ($5b + c$) to two vertices, and we will allow it to be "compensated" by vertices of other types, which may have a surplus.

Consider an index $j \in T_2$. The fact that $j \in T_2$ guarantees that earlier edges $I^1, \ldots, I^{j-1}$ are all vertex disjoint from I^j. If later edges $I^{j+1}, \ldots, I^r$ are all vertex disjoint from I^j, then $|E_{j-1}| = 4$ and E_{j-1} is disjoint from $\{E_{j'}\}_{j' \in D \setminus \{j-1\}}$, and this means that $H \notin \mathcal{H}$. Thus, there exists $j' > j$ such that $I^{j'}$ and I^j share an vertex. Take the smallest $j' > j$, and say that j' *compensates* j. Note that $j' \notin T_2$.

We will allow a type 1 index to compensate at most one type 2 index, and a type -1 or 0 index to compensate at most two type 2 indices. We consider below the constraints implied by each kind of compensation.

1. *One Type* 1 *index j' compensates one Type* 2 *index j*

   $v_{j'} + v_j = 3$ and $e_{j'} + e_j \ge 8$ (5 from e_j and 3 from $e_{j'}$). This is possible if $8b + 2c \ge 3$.

2. *One Type* 0 *index j' compensates one Type* 2 *index j*

   $v_{j'} + v_j = 2$ and $e_{j'} + e_j \ge 5$ (5 from e_j). This is possible if $5b + 2c \ge 2$.

3. *One Type* 0 *index j' compensates two Type* 2 *indices j_1 and j_2 (say $j_1 < j_2$).*

   There are two cases.

(a) $e_{j'} + e_{j_1} + e_{j_2} \geq 11$: $v_{j'} + v_{j_1} + v_{j_2} = 4$. This is possible if $11b + 3c \geq 4$.

(b) $e_{j'} + e_{j_1} + e_{j_2} = 10$: since $e_{j_1}, e_{j_2} \geq 5$, this means that $e_{j'} = 0$.

First, we note that since j_1 is a type 2 index and j' is the smallest index j such that $I^{j_1} \cap I_j \neq \emptyset$, in the graph $H_{j'-1}$, vertices in I^{j_1} only have edges to vertices in I^{j_1-1} and I^{j_1+1}. Similarly, vertices in I^{j_2} only have edges to vertices in I^{j_2-1} and I^{j_2+1}.

Since $I^{j'}$ shares one vertex each with I^{j_1} and I^{j_2}, let $I^{j'} = \{i_1^{j'}, i_2^{j'}\}$ with $i_1^{j'} \in I^{j_1}$ and $i_2^{j'} \in I^{j_2}$. Since $e_{j'} = 0$ means that $I^{j'} = \{i_1^{j'}, i_2^{j'}\}$ was in $H_{j'-1}$. However, this is an edge between vertices in I^{j_1} and I^{j_2}. By the above argument, this is only possible if $j_2 = j_1 + 1$. Also, since j' is type 0 and $I^{j'}$ shares a vertex with I^{j_2}, we must have $j' > j_2 + 1$ (otherwise j' would be type -1).

Consider $I^{j'-1}$, which are vertex disjoint from both I^{j_1} and I^{j_2}. If $I^{j'-1} \neq I^{j_1-1}$, at least one edge between $I^{j'-1}$ and I^{j} was not in $H_{j'-1}$, contradicting the assumption $e_{j'} = 0$. Therefore, $I^{j'-1} = I^{j_1-1}$. For the same reason, $I^{j'-1} = I^{j_2+1}$. Thus, in particular, we have $I^{j_2+1} = I^{j_1-1}$. Thus, $j_2 + 1$ is also a type 0 index. Moreover, it cannot compensate any previous index, since any such index would already be compensated by $j_1 - 1$.

In this case we consider that I^{j_2+1} and $I^{j'}$ jointly compensate j_1 and j_2. $v_{j'} + v_{j_2+1} + v_{j_1} + v_{j_2} = 4$ and $e_{j_2+1} + e_{j'} + e_{j_1} + e_{j_2} \geq 10$. Compensation is possible if $10b + 4c \geq 4$.

4. *One Type -1 index j' compensates one Type 2 index j.*

   $v_{j'} + v_j \leq 3$ and $e_{j'} + e_j \geq 5$ (5 from e_j). Compensation is possible if $5b + 4b + 2c \geq 2$.

5. *One Type -1 index j' compensates two Type 2 indices j_1 and j_2*

   We have $v_{j'} + v_{j_1} + v_{j_2} \leq 5$ and $e_{j'} + e_{j_1} + e_{j_2} \geq 10$. Compensation is possible if $10b + 4b + 3c \geq 5$.

Each index j of Type 2 is compensated by exactly one other index j'. We also require indices of types 1 and -1 which do not compensate any other index, to have value at least 1 (to account for the one vertex added). This is true if $3b + c \geq 1$ and $4b + c \geq 1$.

Aggregating the above conditions (and discarding the redundant ones), we take

$$c = \max\left\{\frac{3}{2} - 4b,\ 1 - \frac{5b}{2},\ \frac{4}{3} - \frac{11b}{3},\ \frac{5}{3} - \frac{14b}{3}\right\}$$

It is easy to check that the maximum is attained by $c = 5/3 - 14b/3$ when $b \in [0, 1/4]$ and $c = 3/2 - 4b$ when $b \in [1/4, 1/3]$. ■

8.10 Lifting $\|\cdot\|_{sp}$ lower bounds to higher levels

For a matrix $B \in \mathbb{R}^{[n]^{q/2}\times[n]^{q/2}}$, let B^S denote the matrix obtained by symmetrizing B, i.e. for any $I, J \in [n]^{q/2}$,

$$B^S[I,J] \;:=\; \frac{1}{|\mathcal{O}(\alpha(I)+\alpha(J))|}\cdot \sum_{\substack{\alpha(I')+\alpha(J')\\=\alpha(I)+\alpha(J)}} B[I',J']$$

Equivalently, B^S can be defined as follows:

$$B^S = \frac{1}{q!}\cdot\sum_{\pi\in\mathbb{S}_q} B^\pi$$

where for any $K \in [n]^q$, $B^\pi[K] := B[\pi(K)]$.

For a matrix $M \in \mathbb{R}^{[n]^2\times[n]^2}$ let $T \in \mathbb{R}^{[n]^4}$ denote the tensor given by, $T[i_1,i_2,i_3,i_4] = M[(i_1,i_2),(i_3,i_4)]$. Also for any non-negative integers x, y satisfying $x+y=4$, let $M_{x,y} \in \mathbb{R}^{[n]^x\times[n]^y}$ denote the matrix given by, $M[(i_1,\ldots,i_x),(j_1,\ldots j_y)] = T[i_1,\ldots,i_x,j_1,\ldots j_y]$. We will use the following result that we prove in Section 8.10.3.

Theorem 8.10.1 (Lifting "Stable" $\|\cdot\|_{sp}$ Lower Bounds)**.** *Let $M \in \mathbb{R}^{[n]^2\times[n]^2}$ be a degree-4-SOS-symmetric matrix satisfying*

$$\|M\|_{S_1},\ \|M_{3,1}\|_{S_1} \le 1.$$

Then for any q divisible by 4,

$$\|\left(M^{\otimes q/4}\right)^S\|_{S_1} = 2^{O(q)}$$

8.10.1 Gap between $\|\cdot\|_{sp}$ and $\|\cdot\|_2$ for Non-Neg. Coefficient Polynomials

Lemma 8.10.2. *Consider any homogeneous polynomial g of even degree-t and let $M_g \in \mathbb{R}^{[n]^{t/2}\times[n]^{t/2}}$ be its SoS-symmetric matrix representation. Then $\|g\|_{sp} \ge \|M_g\|_F^2/\|M_g\|_{S_1}$.*

Proof. We know by strong duality, that

$$\|g\|_{sp} = \max\left\{\langle X, M_g\rangle \;\middle|\; \|X\|_{S_1}=1,\ \ X \text{ is SoS-Symmetric},\ \ X\in\mathbb{R}^{[n]^{t/2}\times[n]^{t/2}}\right\}.$$

The claim follows by substituting $X := M_g/\|M_g\|_{S_1}$. ∎

Theorem 8.10.3. *For any q divisible by 4 and f as defined in Section 8.9, we have that w.h.p.*

$$\frac{\|f^{q/4}\|_{sp}}{\|f^{q/4}\|_2} \;\ge\; \frac{n^{q/24}}{(q\log n)^{O(q)}}.$$

Proof. Let f be the degree-4 homogeneous polynomial as defined in Section 8.9 and let $M = M_f$ be its SoS-symmetric matrix representation. Let $g := f^{q/4}$ and let M_g be its SoS-symmetric matrix representation. Thus $M_g = (M^{\otimes q/4})^S$ and it is easily verified that w.h.p., $\|M\|_F^2 \geq \widetilde{\Omega}(n^4 p^6) = \widetilde{O}(n^2)$ and also $\|M_g\|_F^2 \geq \widetilde{\Omega}((n^4p^6)^{q/4}/q^{O(q)}) = \widetilde{\Omega}(n^{q/2}/q^{O(q)})$.

It remains to estimate $\|M_g\|_{S_1}$ so that we may apply Lemma 8.10.2. Implicit, in the proof of Lemma 8.9.8, is that w.h.p. M has one eigenvalue of magnitude $O(n^2 p) = O(n^{5/3})$ and at most $O(n^2 p) = O(n^{5/3})$ eigenvalues of magnitude $\widetilde{O}(n^{3/2}p^4) = \widetilde{O}(n^{1/6})$. Thus $\|M\|_{S_1} = \widetilde{O}(n^{11/6})$ w.h.p. Now we have that $\|M_{1,3}\|_{S_1} \leq \sqrt{n} \cdot \|M_{1,3}\|_F = \sqrt{n} \cdot \|M\|_F = \widetilde{O}(n^{3/2})$ w.h.p. Thus on applying Theorem 8.10.1 to $M/\widetilde{O}(n^{11/6})$, we get that $\|M_g\|_{S_1}/\widetilde{O}(n^{11q/24}) \leq 2^{O(q)}$ w.h.p.

Thus, applying Lemma 8.10.2 yields the claim. ∎

8.10.2 Tetris Theorem

Let $M \in \mathbb{R}^{[n]^2 \times [n]^2}$ be a degree-4 SoS-Symmetric matrix, let $M_A := M_{3,1} \otimes M_{0,4} \otimes M_{3,1}$, let $M_B := M_{3,1} \otimes M_{1,3}$, let $M_C := M$ and let $M_D := \mathsf{Vec}(M)\,\mathsf{Vec}(M)^T = M_{0,4} \otimes M_{4,0}$. For any permutation $\pi \in \mathbb{S}_{q/2}$ let $\overline{\pi} \in \mathbb{S}_{n^{q/2}}$ denote the permutation that maps any $i \in [n]^{q/2}$ to $\pi(i)$. Also let $P_\pi \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ denote the row-permutation matrix induced by the permutation $\overline{\pi}$. Let $P := \sum_{\pi \in \mathbb{S}_{q/2}} P_\pi$. Let $\mathcal{R}(a,b,c,d) := (c!^{\,2}\, 2!^{\,2c})\, (b!\,(2a+b)!\,3!^{\,2a+2b})\, (d!\,(a+d)!\,4!^{\,a+2d})$. Define

$$\mathfrak{M}(a,b,c,d) := \frac{M_A^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d}}{\mathcal{R}(a,b,c,d)}$$

$$\overline{\mathfrak{M}}(a,b,c,d) := \frac{(M_A^T)^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d}}{\mathcal{R}(a,b,c,d)}$$

$$S^{\mathfrak{M}} := \left\{ P \cdot \mathfrak{M}(a,b,c,d) \cdot P^T \;\middle|\; 12a + 8b + 4c + 8d = q \right\} \cup \left\{ P \cdot \overline{\mathfrak{M}}(a,b,c,d) \cdot P^T \;\middle|\; 12a + 8b + 4c + 8d = q \right\}.$$

Theorem 8.10.4. *Let $M \in \mathbb{R}^{[n]^2 \times [n]^2}$ be a degree-4-SOS-symmetric matrix. Then*

$$(q/4)! \cdot 4!^{\,q/4} \cdot \sum_{\mathfrak{M} \in S^{\mathfrak{M}}} \mathfrak{M} \quad = \quad \sum_{\pi \in \mathbb{S}_q} \left(M^{\otimes q/4}\right)^{\pi} \quad = \quad q! \cdot \left(M^{\otimes q/4}\right)^S \qquad (8.10)$$

We shall prove this claim in Section 8.10.4 after first exploring its consequences.

8.10.3 Lifting Stable Degree-4 Lower Bounds

Theorem 8.10.5 (Lifting "Stable" $\|\cdot\|_{sp}$ Lower Bounds: Restatement of Theorem 8.10.1)**.** *Let $M \in \mathbb{R}^{[n]^2 \times [n]^2}$ be a degree-4-SOS-symmetric matrix satisfying*

$$\|M\|_{S_1},\ \|M_{3,1}\|_{S_1} \leq 1.$$

Then for any q divisible by 4*,*

$$\|\left(M^{\otimes q/4}\right)^S\|_{S_1} = 2^{O(q)}$$

Proof. Implicit in the proof of Theorem 8.10.4 is the following:

$$\begin{aligned} & q! \cdot \left(M^{\otimes q/4}\right)^S \\ = & \sum_{12a+8b+4c+8d=q} \frac{(q/4)! \cdot 4!^{q/4}}{\mathcal{R}(a,b,c,d)} \sum_{\sigma_1,\sigma_2 \in \mathbb{S}_{q/2}} \left(P_{\sigma_1}^T \left(M_A^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d}\right) P_{\sigma_2}\right) \\ + & \sum_{12a+8b+4c+8d=q} \frac{(q/4)! \cdot 4!^{q/4}}{\mathcal{R}(a,b,c,d)} \sum_{\sigma_1,\sigma_2 \in \mathbb{S}_{q/2}} \left(P_{\sigma_1}^T \left((M_A^T)^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d}\right) P_{\sigma_2}\right) \end{aligned} \tag{8.11}$$

First note that $(q/4)! \cdot 4!^{q/4} / \mathcal{R}(a,b,c,d) \leq 2^{O(q)}$ since for any integers i, j, k, l

$$(i+j+k+l)!/(i! \cdot j! \cdot k! \cdot l!) \leq 4^{i+j+k+l}.$$

Next note that $\|M_{0,4}\|_{S_1} = \|M\|_F \leq \|M\|_{S_1} \leq 1$. Combining this with the fact that $\|M_{1,3}\|_{S_1}, \|M\|_{S_1} \leq 1$, we get that $\|M_A\|_{S_1}, \|M_B\|_{S_1}, \|M_C\|_{S_1}, \|M_D\|_{S_1} \leq 1$, since $\|X \otimes Y\|_{S_1} = \|X\|_{S_1} \cdot \|Y\|_{S_1}$ for any (possibly rectangular) matrices X and Y. Further note that Schatten 1-norm is invariant to multiplication by a permutation matrix. Thus the claim follows by applying triangle inequality to the $O_q(q^q)$ terms in Eq. (8.11). ■

8.10.4 Proof of Tetris Theorem

We start with defining a *hypergraphical matrix* which will allow a more intuitive paraphrasing of Theorem 8.10.4. By now, this is an important formalism in the context of SoS, and closely-related objects have been defined in several works, including [DM15], [RRS16], [BHK+16].

Hypergraphical Matrix

Definition 8.10.6. *For symbolic sets* $L = \{\ell_1, \ldots \ell_{q_1}\}, R = \{r_1, \ldots r_{q_2}\}$*, a d-uniform* **template-hypergraph** *represented by* (L, R, E)*, is a d-uniform hypergraph on vertex set* $L \uplus R$ *with E being the set of hyperedges.*

For $I = (i_1, \ldots i_{q_1})[n]^{q_1}, J = (j_1, \ldots j_{q_2}) \in [n]^{q_2}$*, we also define a related object called* **edge-set instantiation** *(and denoted by* $E(I, J)$*) as the set of size-d multisets induced by E on substituting* $\ell_t = i_t$ *and* $r_t = j_t$*.*

Remark. There is a subtle distinction between E and $E(I, J)$ above, in that E is a set of d-sets and $E(I, J)$ is a set of size-d multisets (i.e. $e \in E(I, J)$ can have repeated elements).

Definition 8.10.7. *Given an SoS-symmetric order-d tensor* T *and a d-uniform template-hypergraph* $H = (L, R, E)$ *with* $|L| = q_1, |R| = q_2$*, we define the d-uniform degree-*(q_1, q_2) *hypergraphical matrix* $\mathsf{M}^{\mathsf{T}}_{hyp}(H)$ *as*

$$\mathsf{M}^{\mathsf{T}}_{hyp}(H)[I, J] = \prod_{e \in E(I,J)} \mathsf{T}[e]$$

for any $I \in [n]^{q_1}, J \in [n]^{q_2}$.

In order to represent Theorem 8.10.4 in the language of hypergraphical matrices, we first show how to represent $M^{\otimes q/4}$ and $M_A^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d}$ in this language.

Kronecker Products of Hypergraphical Matrices

We begin with the observation that the kronecker product of hypergraphical matrices yields another hypergraphical matrix (corresponding to the "disjoint-union" of the template-hypergraphs).

Definition 8.10.8. *let* $H = (L, R, E)$, $H' = (L', R', E')$ *be template-hypergraphs with* $|L| = q_1, |R| = q_2, |L'| = q_3, |R'| = q_4$*. Let* $\overline{H} = (\overline{L}, \overline{R}, \overline{E})$ *be a template-hypergraph with* $|\overline{L}| = q_1 + q_3, |\overline{R}| = q_2 + q_4$*, where* $\bar{\ell}_t = \ell_t$ *if* $t \in [q_1]$, $\bar{\ell}_t = \ell'_t$ *if* $t \in [q_1 + 1, q_1 + q_3]$, $\bar{r}_t = r_t$ *if* $t \in [q_2]$, $\bar{r}_t = r'_t$ *if* $t \in [q_2 + 1, q_2 + q_4]$*, and* $\overline{E} = E \uplus E'$*. We call* $\overline{H}$ *the* **disjoint-union** *of* H *and* H'*, which we denote by* $H \uplus H'$*.*

Observation 8.10.9. *Let* T *be an SOS-symmetric order-d tensor and let* $H = (L, R, E)$, $H' = (L', R', E')$ *be template-hypergraphs. Then,*

$$\mathsf{M}^{\mathsf{T}}_{hyp}(H) \otimes \mathsf{M}^{\mathsf{T}}_{hyp}(H') = \mathsf{M}^{\mathsf{T}}_{hyp}(H \uplus H')$$

Remark. Note that the disjoint-union operation on template-hypergraphs does not commute, i.e. $\mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H \uplus H') \neq \mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H' \uplus H)$ (since kronecker-product does not commute).

Now consider a degree-4 SoS-symmetric matrix M (as in the statment of Theorem 8.10.4) and let T be the SoS-symmetric tensor corresponding to M. Then for any $x + y = 4$ we have that $M_{x,y} = \mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H_{x,y})$, where $H_{x,y} = (L, R, E)$ is the template-hypergraph satisfying $L = \{\ell_1, \ldots \ell_x\}, R = \{r_1, \ldots r_y\}$ and $E = \{\{\ell_1, \ldots \ell_x, r_1, \ldots r_y\}\}$. Combining this observation with Observation 8.10.9 yields that $M_A = \mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H_A)$, $M_B = \mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H_B)$, $M_C = \mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H_C)$, $M_D = \mathsf{M}^{\mathsf{T}}_{\text{hyp}}(H_D)$, where $H_A := H_{3,1} \uplus H_{0,4} \uplus H_{3,1}$, $H_B := H_{3,1} \uplus H_{1,3}$, $H_C := H_{2,2}$, and $H_D := H_{4,0} \uplus H_{0,4}$. Lastly, another application of Observation 8.10.9 to the above, yields

Observation 8.10.10. *For a template-hypergraph* H*, let* $H^{\uplus t}$ *denote* $\biguplus_{g \in [t]} H$*. Then,*

(1) $M^{\otimes q/4} = \mathsf{M}^{\mathsf{T}}_{hyp}(H_{2,2}^{\uplus q/4})$.

(2) $M_A^{\otimes a} \otimes M_B^{\otimes b} \otimes M_C^{\otimes c} \otimes M_D^{\otimes d} = \mathsf{M}_{hyp}^{\mathsf{T}}(H(a,b,c,d))$ *where*

$$H(a,b,c,d) \;:=\; H_A^{\uplus a} \uplus H_B^{\uplus b} \uplus H_C^{\uplus c} \uplus H_D^{\uplus d}$$

For technical reasons we also define the following related template-hypergraph:

$$\overline{H}(a,b,c,d) \;:=\; \overline{H}_A^{\uplus a} \uplus H_B^{\uplus b} \uplus H_C^{\uplus c} \uplus H_D^{\uplus d},$$

where $\overline{H}_A$ is the template-hypergraph whose corresponding hypergraphical matrix is M_A^T.

To finish paraphrasing Theorem 8.10.4, we are left with studying the effect of permutations on hypergraphical matrices - which is the content of the following section.

Hypergraphical Matrices under Permutation

Recall that for any matrix $B \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ and $\pi \in \mathbb{S}_q$, B^π is the matrix satisfying $B^\pi[K] := B[\pi(K)]$ where $K \in [n]^{q/2} \times [n]^{q/2}$.

Also recall that for any permutation $\sigma \in \mathbb{S}_{q/2}$, $\overline{\sigma} \in \mathbb{S}_{n^{q/2}}$ denotes the permutation that maps any $i \in [n]^{q/2}$ to $\sigma(i)$ and also that $P_\sigma \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ denotes the $[n]^{q/2} \times [n]^{q/2}$ row-permutation matrix induced by the permutation $\overline{\sigma}$.

We next define a permuted template hypergraph in order to capture how permuting a hypergraphical matrix (in the senses above) can be seen as permutations of the vertex set of the hypergraph.

[Permuted Template-Hypergraph] For any $\pi \in \mathbb{S}_q$ (even q), and a d-uniform template-hypergraph $H = (L, R, E)$ with $|L| = |R| = q/2$, let $H^\pi = (L', R', E)$ denote the template-hypergraph obtained by setting $\ell'_t := k_t$ and $r_t = k_{t+q/2}$ for $t \in [q/2]$, where $K = (k_1, \dots k_q) = \pi(L \oplus R)$.

Similarly for any $\sigma_1, \sigma_2 \in \mathbb{S}_{q/2}$, let $H^{\sigma_1,\sigma_2} = (L', R', E)$ denote the template-hypergraph obtained by setting $\ell'_t := \sigma_1(\ell_t)$ and $r'_t = \sigma_2(r_t)$.

We then straightforwardly obtain

Observation 8.10.11. *For any $\pi \in \mathbb{S}_q$, $\sigma_1, \sigma_2 \in \mathbb{S}_{q/2}$, SoS-symmetric order-$d$ tensor T and any d-uniform template-hypergraph H,*

(1) $\left(\mathsf{M}_{hyp}^{\mathsf{T}}(H)\right)^\pi = \mathsf{M}_{hyp}^{\mathsf{T}}(H^\pi)$

(2) $P_{\sigma_1} \cdot \mathsf{M}_{hyp}^{\mathsf{T}}(H) \cdot P_{\sigma_2}^T = \mathsf{M}_{hyp}^{\mathsf{T}}(H^{\sigma_1,\sigma_2})$

Thus to prove Theorem 8.10.4, it remains to estabish

$$\frac{1}{4!^{q/4} \cdot (q/4)!} \cdot \sum_{\pi \in \mathbb{S}_q} \mathsf{M}_{\text{hyp}}^{\mathsf{T}}((H_{2,2}^{\uplus q/4})^\pi)$$
$$= \sum_{12a+8b+4c+8d=q} \frac{1}{\mathcal{R}(a,b,c,d)} \cdot \sum_{\sigma_1,\sigma_2 \in \mathbb{S}_{q/2}} \mathsf{M}_{\text{hyp}}^{\mathsf{T}}(H(a,b,c,d)^{\sigma_1,\sigma_2}) \;+$$

$$\sum_{12a+8b+4c+8d=q} \frac{1}{\mathcal{R}(a,b,c,d)} \cdot \sum_{\sigma_1,\sigma_2 \in \mathbb{S}_{q/2}} \mathsf{M}^{\mathsf{T}}_{\mathrm{hyp}}(\overline{H}(a,b,c,d)^{\sigma_1,\sigma_2}) \tag{8.12}$$

We will establish this in the next section by comparing the template-hypergraphs generated (with multiplicities) in the LHS with those generated in the RHS.

Proof of Eq. (8.12)

We start with some definitions to track the template-hypergraphs generated in the LHS and RHS of Eq. (8.12). For any $12a + 8b + 4c + 8d = q$, let

$$\mathcal{F}(a,b,c,d) := \left\{ H(a,b,c,d)^{\sigma_1,\sigma_2} \;\middle|\; \sigma_1,\sigma_2 \in \mathbb{S}_{q/2} \right\}$$
$$\overline{\mathcal{F}}(a,b,c,d) := \left\{ \overline{H}(a,b,c,d)^{\sigma_1,\sigma_2} \;\middle|\; \sigma_1,\sigma_2 \in \mathbb{S}_{q/2} \right\}$$
$$\mathcal{F} := \left\{ (H_{2,2}^{\uplus q/4})^{\pi} \;\middle|\; \pi \in \mathbb{S}_q \right\}$$

Firstly, it is easily verified that whenever $(a,b,c,d) \neq (a',b',c',d')$, $\mathcal{F}(a,b,c,d) \cap \mathcal{F}(a',b',c',d') = \varphi$, and that $\mathcal{F}(a,b,c,d) \cap \overline{\mathcal{F}}(a,b,c,d) = \varphi$. It is also easily verified that for any $12a + 8b + 4c + 8d = q$, and any $H \in \mathcal{F}(a,b,c,d)$,

$$\mathcal{R}(a,b,c,d) \;=\; \left| \left\{ (\sigma_1,\sigma_2) \in \mathbb{S}_{q/2}^2 \;\middle|\; H(a,b,c,d)^{\sigma_1,\sigma_2} = H \right\} \right|$$

and for any $H \in \mathcal{F}$,

$$4!^{q/4} \cdot (q/4)! \;=\; \left| \left\{ \pi \in \mathbb{S}_q \;\middle|\; (H_{2,2}^{\uplus q/4})^{\pi} = H \right\} \right| .$$

Thus in order to prove Eq. (8.12), it is sufficient to establish that

$$\mathcal{F} \;=\; \biguplus_{12a+8b+4c+8d=q} (\mathcal{F}(a,b,c,d) \uplus \overline{\mathcal{F}}(a,b,c,d)) \tag{8.13}$$

It is sufficient to establish that

$$\mathcal{F} \;\subseteq\; \biguplus_{12a+8b+4c+8d=q} (\mathcal{F}(a,b,c,d) \uplus \overline{\mathcal{F}}(a,b,c,d)) \tag{8.14}$$

since the other direction is straightforward. To this end, consider any $H = (L, R, E) \in \mathcal{F}$, and for any $x + y = 4$, define

$$s_{x,y} := \left| \left\{ e \in E \;\middle|\; |e \cap L| = x,\; |e \cap R| = y \right\} \right| .$$

Now clearly $H \in \mathcal{F}(a,b,c,d)$ iff

$$s_{0,4} = a + d,\; s_{3,1} = 2a + b,\; s_{1,3} = b,\; s_{2,2} = c,\; s_{4,0} = d. \tag{8.15}$$

and $H \in \overline{\mathcal{F}}(a,b,c,d)$ iff

$$s_{4,0} = a + d,\ s_{1,3} = 2a + b,\ s_{3,1} = b,\ s_{2,2} = c,\ s_{0,4} = d. \tag{8.16}$$

Thus we need only find $12a' + 8b' + 4c' + 8d' = q$, such that Eq. (8.15) or Eq. (8.16) is satisfied.

We will assume w.l.o.g. that $s_{0,4} \geq s_{4,0}$ and show that one can satisfy Eq. (8.15), since if $s_{(0,4)} < s_{(4,0)}$, an identical argument allows one to show that Eq. (8.16) is satisfiable. So let $d' = s_{4,0}$, $c' = s_{2,2}$, $b' = s_{1,3}$ and $a' = (s_{3,1} - s_{1,3})/2$. Since $H \in \mathcal{F}$, it must be true that $4s_{4,0} + 3s_{3,1} + 2s_{2,2} + s_{(1,3)} = q/2$. Thus, $12a' + 8b' + 4c' + 8d' = 8s_{4,0} + 6s_{3,1} + 4s_{2,2} + 2s_{1,3} = q$ as desired. We will next see that (a', b', c', d') and $s_{x,y}$ satisfy Eq. (8.15). We have by construction that $s_{4,0} = d'$, $s_{2,2} = c'$, $s_{1,3} = b'$ and $s_{3,1} = 2a' + b'$. It remains to show that $s_{0,4} = a' + d'$. Now we know that $4s_{4,0} + 3s_{3,1} + 2s_{2,2} + s_{1,3} = q/2$ and $4s_{0,4} + 3s_{1,3} + 2s_{2,2} + s_{3,1} = q/2$. Subtracting the two equations yields $s_{0,4} - s_{4,0} = (s_{3,1} - s_{1,3})/2$. This implies $a' + d' = s_{4,0} + (s_{3,1} - s_{1,3})/2 = s_{0,4}$, and furthermore it implies that a' is non-negative since we assumed $s_{0,4} \geq s_{4,0}$. So Eq. (8.15) is satisfied. Thus we have established Eq. (8.14), which completes the proof of Theorem 8.10.4.

8.11 Open problems

Our work makes progress on polynomial optimization based on new spectral techniques for dealing with higher order matrix representations of polynomials. Several interesting questions in the subject remain open, and below we collect some of the salient ones brought to the fore by our work.

1. What is the largest possible ratio between $\Lambda(f)$ and $\|f\|_2$ for arbitrary homogeneous polynomials of degree d? Recall that we have an upper bound of $O_d(n^{d/2-1})$ and a lower bound of $\Omega_d(n^{d/4-1/2})$, and closing this quadratic gap between these bounds is an interesting challenge. Even a lower bound for $\|\cdot\|_{sp}$ that improves upon the current $\Omega_d(n^{d/4-1/2})$ bound by polynomial factors would be very interesting.

2. A similar goal to pursue would be closing the gap between upper and lower bounds for polynomials with non-negative coefficients.

3. We discussed two relaxations of $\|h\|_2$ — $\Lambda(h)$ which minimizes the maximum eigenvalue $\lambda_{\max}(M_h)$ over matrix representations M_h of h, and $\|h\|_{sp}$ which minimizes the spectral norm $\|M_h\|_2$. How far apart, if at all, can these quantities be for arbitrary polynomials h?

4. We studied three classes of polynomials: arbitrary, those with non-negative coefficients, and sparse. Are there other natural classes of polynomials for which we can give improved SoS-based (or other) approximation algorithms? Can our techniques be used in sub-exponential algorithms for special classes?

5. Despite being such a natural problem for which known algorithms give weak polynomially large approximation factors, the known NP-hardness results for polynomial optimization over the unit sphere only rule out an FPTAS. Can one obtain NP-hardness results for bigger approximation factors?

8.12 Oracle Lower Bound

Khot and Naor [KN08] observed that the problem of maximizing a polynomial over unit sphere can be reduced to computing diameter of centrally symmetric convex body. This observation was also used by So [So11] later. We recall the reduction here: For a convex set K, let K° denote the polar of K, i.e., $K^\circ = \{y : \forall x \in K \langle x, y \rangle \leq 1\}$. For a degree-3 polynomial $P(x, y, z)$ on $3n$ variables, let $\|x\|_P = \|P(x, \cdot, \cdot)\|_{sp}$ where $P(x, \cdot, \cdot)$ is a degree-2 restriction of P with x variables set. Let $\mathbb{B}_P = \{x : \|x\|_P \leq 1\}$. From the definition of polar and $\|\cdot\|_{sp}$, we have:

$$\begin{aligned}\max_{\|x\|_2, \|y\|_2, \|z\|_2 \leq 1} P(x, y, z) &= \max_{x \in \mathbb{B}_2} \|x\|_P \\ &= \max_{x \in \mathbb{B}_P^\circ} \|x\|_2\end{aligned}$$

For general convex bodies, a lower bound for number of queries with "weak separation oracle" for approximating the diameter of the convex body was proved by Brieden et al. [BGK$^+$01] and later improved by Khot and Naor [KN08]. We recall the definition:

Definition 8.12.1. *For a given a convex body P, a* weak separation oracle *A is an algorithm which on input (x, ε) behaves as following:*

- *If $x \in A + \varepsilon\mathbb{B}_2$, A accepts it.*
- *Else A outputs a vector $c \in \mathbb{Q}^n$ with $\|c\|_\infty = 1$ such that for all y such that $y + \varepsilon\mathbb{B}_2 \subset P$ we have $c^T x + \varepsilon \geq c^T y$.*

Let $K_{s,v}$ be the convex set $K_{s,v}^{(n)} = \operatorname{conv}(\mathbb{B}_2^n \cup \{sv, -sv\})$, for unit vector v. Brieden et al. [BGK$^+$01] proved the following theorem:

Theorem 8.12.2. *Let A be a randomized algorithm, for every convex set P, with access to a weak separation oracle for P. Let $\mathcal{K}(n, s) = \{K_{s,u}^{(n)}\}_{u \in \mathbb{S}_2^{n-1}} \cup \{\mathbb{B}_2^n\}$. If for every $K \in \mathcal{K}(n, s)$ and $s = \frac{\sqrt{n}}{\lambda}$, we have:*

$$\mathbf{Pr}\left[A(K) \leq \operatorname{diam}(K) \leq \frac{\sqrt{n}}{\lambda} A(K)\right] \geq \frac{3}{4}$$

where $\operatorname{diam}(K)$ is the diameter of K, then A must use at least $\mathrm{O}(\lambda 2^{\lambda^2/2})$ oracle queries for $\lambda \in [\sqrt{2}, \sqrt{n/2}]$.

Using $\lambda = \log n$, we get that to get $s = \frac{\sqrt{n}}{\log n}$ approximation to diameter, A must use super-polynomial number of queries to the weak separation oracle. We note that this was later improved to give analogous lower bound on the number of queries for an approximation factor $\sqrt{\frac{n}{\log n}}$ by Khot and Naor [KN08].

Below, we show that the family of hard convex bodies considered by Brieden et al. [BGK+01] can be realized as $\{\mathbb{B}_P^\circ\}_{P\in\mathcal{P}}$ by a family of polynomials $\mathcal{P}$ – which, in turn, establishes a lower bound of $\Omega\,\frac{\sqrt{n}}{\log n}$ on the approximation for polynomial optimization, achievable using this approach, for the case of $d = 3$. For an unit vector $u \in \mathbb{S}_2^{n-1}$, let P_u be the polynomial defined as:

$$P_u(x, y, z) = \sum_{i=1}^{n} x_i y_i z_1 + s \cdot (u^T x) y_n z_n.$$

A matrix representation of $P_u(x, \cdot, \cdot)$, with rows indexed by y and columns indexed by z variables is as follows:

$$A_u = \begin{pmatrix} x_1 & 0 & \dots & 0 & 0 \\ x_2 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x_{n-1} & 0 & \dots & 0 & 0 \\ x_n & 0 & \dots & 0 & s\cdot(u^T x) \end{pmatrix} \text{ and so, } A_u^T A_u = \begin{pmatrix} \|x\|_2^2 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & s^2\cdot\left|u^T x\right|^2 \end{pmatrix}.$$

This proves: $\|x\|_{P_u} = \|P_u(x, \cdot, \cdot)\|_{sp} = \|A_u\|_{sp} = \max\{\|x\|_2, s\left|u^T x\right|\}$.

Let $B = \{x : \|x\|_2 \le 1\}$ and $C_u = \{x : s \cdot \left|u^T x\right| \le 1\}$. We note that, $B^\circ = \{y \in \mathbb{R}^n : \|y\|_2 \le 1\}$ and, $C_u^\circ = \{\lambda \cdot u : \lambda \in [-s, s]\} = \operatorname{conv}(()\{-s \cdot u, s \cdot u\})$.

Next, we observe: $\mathbb{B}_{P_u} = B \cap C_u$. It follows from De Morgan's law of polars that: $\mathbb{B}_{P_u}^\circ = (B \cap C_u)^\circ = \operatorname{conv}(()\, B^\circ \cup C_u^\circ) = \operatorname{conv}(()\, \mathbb{B}_2^n \cup \{-s \cdot u, s \cdot u\}) = K_{s,u}^{(n)}$. Finally, we observe that for the polynomial $P_0 = \sum_{i=1}^n x_i y_i z_1$, we have: $\mathbb{B}_{P_0} = \mathbb{B}_2^n$.

Hence for polynomial $Q \in \mathcal{P} = \{P_u\}_{u\in\mathbb{S}_2^{n-1}} \cup \{P_0\}$, no randomized polynomial can approximate $\operatorname{diam}\mathbb{B}_Q$ within factor $\frac{\sqrt{n}}{q}$ without using more than $2^{\Omega(q)}$ number of queries. Since the algorithm of Khot and Naor [KN08] reduces the problem of optimizing polynomial Q to computing $\operatorname{diam}(\mathbb{B}_Q)$, $\mathcal{P}$ shows that their analysis is almost tight.

8.13 Maximizing $|f(x)|$ vs. $f(x)$

Let $f_{\max}$ denote $\sup_{\|x\|=1} f(x)$. Note that for polynomials with odd-degree, we have $\|f\|_2 = f_{\max}$. When the degree is even, a multiplicative approximation for $f_{\max}$ is not possible since $f_{\max}$ may be 0 or even negative. Moreover, even when $f_{\max}$ is positive, any constructive multiplicative approximation of $f_{\max}$ with a factor (say) B, can be turned into a $1+\varepsilon$ approximation by considering $f' = f - C \cdot \|x\|_2^d$, for $C = (1-\varepsilon) \cdot f_{\max}$ (one can use binary search on the values of C and use the solution give by the constructive algorithm to check).

An alternate notion considered in the literature [HLZ10, So11] is that of relative approximation where one bounds the ratio $(\Lambda - f_{\min})/(f_{\max} - f_{\min})$ (known as a relative approximation), where Λ is the estimate by the algorithm, and $f_{\min}$ is defined analogously to $f_{\max}$. While this is notion is stronger than approximating $\|f\|_2$ in some cases, one can use a shift of f as in the example above (by $C \cdot f_{\min}$) to obtain a relative approximation unless $|f_{\max} - f_{\min}| \,/\, |f_{\min}| = n^{-\omega(1)}$.

Chapter 9

Random Polynomial Optimization over the Sphere

It is a well-known fact from random matrix theory that for an $n \times n$ matrix M whose entries are i.i.d Rademacher or standard normal random variables, the maximum value $x^T Mx$ taken by the associated quadratic form on the unit sphere $\|x\|_2 = 1$, is $\Theta(\sqrt{n})$ with high probability. Further, this maximum value can be computed efficiently for any matrix, as it equals the largest eigenvalue of $(M + M^T)/2$, so one can also efficiently certify that the maximum of a random quadratic form is at most $O(\sqrt{n})$.

This chapter is motivated by the analogous question for tensors. Namely, given a random order-d tensor $\mathcal{A}$ who entries are i.i.d random $\pm$ entries, we would like to certify an upper bound on the maximum value $\mathcal{A}_{\max} := \max_{\|x\|=1}\langle \mathcal{A}, x^{\otimes d}\rangle$ taken by the tensor on the unit sphere. This value is at most $O_d(\sqrt{n})$ with high probability [TS14]. This chapter is concerned with both positive and negative results on the efficacy of the SoS hierarchy in approximately certifying the maxima of random tensors. We next state our results formally.

9.1 Our Results

For an order-q tensor $\mathcal{A} \in (\mathbb{R}^n)^{\otimes d}$, the polynomial $\mathcal{A}(x)$ and its maximum on the sphere $\mathcal{A}_{\max}$ are defined as

$$\mathcal{A}(x) := \langle \mathcal{A}, x^{\otimes d}\rangle \qquad \mathcal{A}_{\max} := \sup_{\|x\|=1} \mathcal{A}(x).$$

When the entries of $\mathcal{A}$ are i.i.d Rademacher random variables (or i.i.d. Gaussians), it is known that $\mathcal{A}_{\max} \lesssim \sqrt{n \cdot d \cdot \log d}$ (see [TS14]). We will also use, for a polynomial g, $g_{\max}$ to denote $\sup_{\|x\|=1} g(x)$.

SoS degree = Polynomial Degree.

We study the performance of degree-q SoS on random tensors of order-q. The formal definition and basic properties of SoS relaxations are presented in Chapter 7.

Theorem 9.1.1. *For any even $q \leq n$, let $\mathcal{A} \in (\mathbb{R}^n)^{\otimes q}$ be a q-tensor with independent, Rademacher entries. With high probability, the value B of the degree-q SoS relaxation of $\mathcal{A}_{\max}$ satisfies*

$$2^{-O(q)} \cdot \left(\frac{n}{q}\right)^{q/4-1/2} \leq \frac{B}{\mathcal{A}_{\max}} \leq 2^{O(q)} \cdot \left(\frac{n}{q}\right)^{q/4-1/2}.$$

This improves upon the $O(n^{q/4})$ upper bound by Montanari and Richard [MR14].

SoS Degree $\gg$ Polynomial Degree.

Theorem 9.1.2. *Let $\mathcal{A} \in (\mathbb{R}^n)^{\otimes d}$ be a d-tensor with independent, Rademacher entries. Then for any even q satisfying $d \leq q \leq n$, with high probability, the degree-q SoS certifies an upper bound B on $\mathcal{A}_{\max}$ where w.h.p.,*

$$\frac{B}{\mathcal{A}_{\max}} \leq \left(\frac{\widetilde{O}(n)}{q}\right)^{d/4-1/2}$$

Remark 9.1.3. *Combining our upper bounds with the work of [HSS15] would yield improved tensor-PCA guarantees on higher levels of SoS. Our techniques prove similar results for a more general random model where each coefficient is independently sampled from a centred subgaussian distribution. See the previous version of the paper [BGL16] for details.*

Remark 9.1.4. *Raghavendra, Rao, and Schramm [RRS16] have independently and concurrently obtained similar (but weaker) results to Theorem 9.1.2 for random degree-d polynomials. Specifically, their upper bounds appear to require the assumption that the SoS level q must be less than $n^{1/(3d^2)}$ (our result only assumes $q \leq n$). Further, they certify an upper bound that matches Theorem 9.1.2 only when $q \leq 2^{\sqrt{\log n}}$.*

9.2 Related Work

Upper Bounds. Montanari and Richard [MR14] presented a $n^{O(d)}$-time algorithm that can certify that the optimal value of $\mathcal{A}_{\max}$ for a random d-tensor is at most $O(n^{\frac{\lceil d/2 \rceil}{2}})$ with high probability. Hopkins, Shi, and Steurer [HSS15] improved it to $O(n^{\frac{d}{4}})$ with the same running time. They also asked how many levels of SoS are required to certify a bound of $n^{3/4-\delta}$ for $d = 3$.

Our analysis asymptotically improves the aforementioned bound when q is growing with n, and we prove an essentially matching lower bound (but only for the case $q = d$). Secondly, we consider the case when d is fixed, and give improved results for the performance of degree-q SoS (for large q), thus answering in part, a question posed by Hopkins, Shi and Steurer [HSS15].

Raghavendra, Rao, and Schramm [RRS16] also prove results analogous to Theorem 9.1.2 for the case of *sparse* random polynomials (a model we do not consider in this work, and which appears to pose additional technical difficulties). This implied upper bounds for refuting random instances of constraint satisfaction problems using higher levels of the SoS hierarchy, which were shown to be tight via matching SoS lower bounds in [KMOW17].

Lower Bounds. While we only give lower bounds for the case of $q = d$, subsequent to our work, Hopkins et al. [HKP+17] proved the following theorem, which gives lower bounds for the case of $q \gg d$:

Theorem 9.2.1. *Let f be a degree-d polynomial with i.i.d. gaussian coefficients. If there is some constant $\varepsilon > 0$ such that $q \geq n^{\varepsilon}$, then with high probability over f, the optimum of the level-q SoS relaxation of $f_{\max}$ is at least*

$$f_{\max} \cdot \Omega_d\left((n/q^{O(1)})^{d/4-1/2}\right) .$$

Note that this almost matches our upper bounds from Theorem 9.1.2, modulo the exponent of q. For this same reason, the above result does not completely recover our lower bound in Theorem 9.1.1 for the special case of $q = d$.

Results for worst-case tensors. In Chapter 8 that q-level SoS gives an $(O(n)/q)^{d/2-1}$ approximation to $\|\mathcal{A}\|_2$ in the case of arbitrary d-tensors and an $(O(n)/q)^{d/4-1/2}$ approximation to $\mathcal{A}_{\max}$ in the case of d-tensors with non-negative entries (for technical reasons one can only approximate $\|\mathcal{A}\|_2 = \max\{|\mathcal{A}_{\max}|, |\mathcal{A}_{\min}|\}$ in the former case).

It is interesting to note that the approximation factor in the case of non-negative tensors matches the approximation factor (upto polylogs) we achieve in the random case. Additionally, the gap given by Theorem 9.1.1 for the case of random tensors provides the best degree-q SoS gap for the problem of approximating the 2-norm of arbitrary q-tensors. Hardness results for the arbitrary tensor 2-norm problem is an important pursuit due to its connection to various problems for which subexponential algorithms are of interest.

9.3 Organization

In Section 9.5 we touch upon the main technical ingredients driving our work, and give an overview of the proof of Theorem 9.1.2 and the lower bound in Theorem 9.1.1. We present the proof of Theorem 9.1.2 for the case of even d in Section 9.6, with the more tricky odd d case handled in Section 9.8. The lower bound on the value of SoS-hierarchy claimed in Theorem 9.1.1 is proved in Theorem 9.7.7, and the upper bound in Theorem 9.1.1 also follows based on some techniques in that section.

9.4 Notation and Preliminaries

SoS Relaxations for $\mathcal{A}_{\max}$.
Given an order-q tensor $\mathcal{A}$, our degree-q SoS relaxation for $\mathcal{A}_{\max}$ which we will henceforth denote by $\mathsf{SoS}_q(\mathcal{A}(x))$ is given by,

$$\begin{array}{cc} \text{maximize} & \widetilde{\mathbf{E}}_C[\mathcal{A}(x)] \\ \text{subject to :} & \widetilde{\mathbf{E}}_C \text{ is a degree-}q \end{array}$$

pseudoexpectation

$\widetilde{\mathbf{E}}_C$ respects $C \equiv \left\{\|x\|_2^q = 1\right\}$

Assuming q is divisible by $2d$, we make an observation that is useful in our upper bounds:

$$\mathcal{A}_{\max} \;\leq\; \mathsf{SoS}_q(\mathcal{A}(x)) \;\leq\; \mathsf{SoS}_q\Big(\mathcal{A}(x)^{q/d}\Big)^{d/q} = \;\Lambda\Big(\mathcal{A}(x)^{q/d}\Big)^{d/q} \tag{9.1}$$

where the second inequality follows from Pseudo-Cauchy-Scwarz, and the equality follows from strong duality of the programs given in Section 7.3.

Note. In the rest of this chapter, we will drop the subscript C of the pseudo-expectation operator since throughout we only assume the hypersphere constraint.

9.5 Overview of our Methods

We now give a high level view of the two broad techniques driving this work, followed by a more detailed overview of the proofs.

Higher Order Mass-Shifting. Our approach to upper bounds on a random low degree (say d) polynomial f, is through exhibiting a matrix representation of $f^{q/d}$ that has small operator norm. Such approaches had been used previously for low-degree SoS upper bounds. However when the SoS degree is constant, the set of SoS symmetric positions is also a constant and the usual approach is to shift all the mass towards the diagonal which is of little consequence when the SoS-degree is low. In contrast, when the SoS-degree is large, many non-trivial issues arise when shifting mass across SoS-symmetric positions, as there are many permutations with very large operator norm. In our setting, mass-shifting approaches like symmetrizing and diagonal-shifting fail quite spectacularly to provide good upper bounds. For our upper bounds, we crucially exploit the existence of "good permutations", and moreover that there are $q^q \cdot 2^{-O(q)}$ such good permutations. On averaging the representations corresponding to these good permutations, we obtain a matrix that admits similar spectral preperties to those of a matrix with i.i.d. entries, and with much lower variance (in most of the entries) compared to the naive representations.

Square Moments of Wigner Semicircle Distribution. Often when one is giving SoS lower bounds, one has a linear functional that is not necessarily PSD and a natural approach is to fix it by adding a pseudo-expectation operator with large value on square polynomials (under some normalization). Finding such operators however, is quite a non-trivial task when the SoS-degree is growing. We show that if $x_1, \ldots, x_n$ are independently drawn from the Wigner semicircle distribution, then for any polynomial p of any degree, $\mathbb{E}[p^2]$ is large (with respect to the degree and coefficients of p). Our proof crucially relies on knowledge of the Cholesky decomposition of the moment matrix of the univariate Wigner distribution. This tool was useful to us in giving tight q-tensor lower bounds, and we believe it to be generally useful for high degree SoS lower bounds.

9.5.1 Overview of Upper Bound Proofs

For even d, let $\mathcal{A} \in \mathbb{R}^{[n]^d}$ be a d-tensor with i.i.d. ± 1 entries and let $A \in \mathbb{R}^{[n]^{d/2} \times [n]^{d/2}}$ be the matrix flattening of $\mathcal{A}$, i.e., $A[I,J] = \mathcal{A}[I \oplus J]$ (recall that $\oplus$ denotes tuple concatenation). Also let $f(x) := \mathcal{A}(x) = \langle \mathcal{A}, x^{\otimes d} \rangle$. It is well known that $f_{\max} \leq O(\sqrt{n \cdot d \cdot \log d})$ with high probability [TS14]. For such a polynomial f and any q divisible by d, in order to establish Theorem 9.1.2, by Eq. (9.1) it is sufficient to prove that with high probability,

$$\left(\Lambda\left(f^{q/d}\right)\right)^{d/q} \leq \widetilde{O}\left(\frac{n}{q^{1-2/d}}\right)^{d/4} = \widetilde{O}\left(\frac{n}{q}\right)^{d/4-1/2} \cdot f_{\max}.$$

We give an overview of the proof. Let $d = 4$ for the sake of clarity of exposition. To prove an upper bound on $\Lambda(f^{q/4})$ using degree-q SoS (assume q is a multiple of 4), we define a suitable matrix representation $M := M_{f^{q/4}} \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ of $f^{q/4}$ and bound $\|M\|_2$. Since $\Lambda(f) \leq (\|M\|_2)^{q/4}$ for any representation M, a good upper bound on $\|M\|_2$ certifies that $\Lambda(f)$ is small.

One of the intuitive reasons taking a high power gives a better bound on the spectral norm is that this creates more entries of the matrix that correspond to the same monomial, and distributing the coefficient of this monomial equally among the corresponding entries reduces variance (i.e., $\mathrm{Var}[X]$ is less than $k \cdot \mathrm{Var}[X/k]$ for $k > 1$). In this regard, the most natural representation M of $f^{q/4}$ is the *complete symmetrization*.

$$\begin{aligned}
&M_c[(i_1, \ldots, i_{q/2}), (i_{q/2+1}, \ldots, i_q)] \\
&= \frac{1}{q!} \cdot \sum_{\pi \in \mathbb{S}_q} A^{\otimes q/4}[(i_{\pi(1)}, \ldots, i_{\pi(q/2)}), (i_{\pi(q/2+1)}, \ldots, i_{\pi(q)})] \\
&= \frac{1}{q!} \cdot \sum_{\pi \in \mathbb{S}_q} \prod_{j=1}^{q/4} A[(i_{\pi(2j-1)}, i_{\pi(2j)}), (i_{\pi(q/2+2j-1)}, i_{\pi(q/2+2j)})].
\end{aligned}$$

However, $\|M_c\|_2$ turns out to be much larger than $\Lambda(f)$, even when $q = 8$. One intuitive explanation is that M_c, as a $n^4 \times n^4$ matrix, contains a copy of $\mathsf{Vec}(A)\,\mathsf{Vec}(A)^T$, where $\mathsf{Vec}(A) \in \mathbb{R}^{[n]^4}$ is the vector with $\mathsf{Vec}(A)\,[i_1, i_2, i_3, i_4] = A[(i_1, i_2), (i_3, i_4)]$. Then $\mathsf{Vec}(A)$ is a vector that witnesses $\|M_c\|_2 \geq \Omega(n^2)$, regardless of the randomness of f. Our final representation [1] is the following *row-column independent symmetrization* that simultaneously respects the spectral structure of a random matrix A and reduces the variance. Our M is given by

$$\begin{aligned}
&M[(i_1, \ldots, i_{q/2}), (j_1, \ldots, j_{q/2})] \\
&= \frac{1}{(q/2)!^2} \cdot \sum_{\pi, \sigma \in \mathbb{S}_{q/2}} A^{\otimes q/4}[(i_{\pi(1)}, \ldots, i_{\pi(q/2)}), (j_{\sigma(1)}, \ldots, j_{\sigma(q/2)})] \\
&= \frac{1}{(q/2)!^2} \cdot \sum_{\pi, \sigma \in \mathbb{S}_{q/2}} \prod_{k=1}^{q/4} A[(i_{\pi(2k-1)}, i_{\pi(2k)}), (j_{\sigma(2k-1)}, j_{\sigma(2k)})].
\end{aligned}$$

[1] the independent and concurrent work of [RRS16] uses the same representation

To formally show $\|M\|_2 = \tilde{O}(n/\sqrt{q})^{q/4}$ with high probability, we use the trace method to show

$$\mathbb{E}\left[\mathsf{Tr}(M^p)\right] \le 2^{O(pq\log p)} \frac{n^{pq/4+q/2}}{q^{pq/8}},$$

where $\mathbb{E}[\mathsf{Tr}(M^p)]$ can be written as (let $I^{p+1} := I^1$)

$$\begin{aligned}
&\mathbb{E}\left[\sum_{I^1,\ldots,I^p\in[n]^{q/2}} \prod_{j=1}^{p} M[I^j, I^{j+1}]\right] \\
&= \sum_{I^1,\ldots,I^p} \mathbb{E}\left[\prod_{j=1}^{p}\Big(\sum_{\pi_j,\sigma_j\in \mathsf{S}_{q/2}} \prod_{k=1}^{q/4} A[(I^k_{\pi_j(2k-1)}, I^k_{\pi_j(2k)}), (I^{k+1}_{\sigma_j(2k-1)}, I^{k+1}_{\sigma_j(2k)})]\Big)\right].
\end{aligned}$$

Let $E(I^1,\ldots,I^p)$ be the expectation value for $I^1,\ldots,I^p$ in the right hand side. We study $E(I^1,\ldots,I^p)$ for each $I^1,\ldots,I^p$ by careful counting of the number of permutations on a given sequence with possibly repeated entries. For any $I^1,\ldots,I^p \in [n]^{q/2}$, let $\#\left(I^1,\ldots,I^p\right)$ denote the number of distinct elements of $[n]$ that occur in $I^1,\ldots,I^p$, and for each $s = 1,\ldots,\#\left(I^1,\ldots,I^p\right)$, let $c^s \in (\{0\}\cup[q/2])^p$ denote the number of times that the jth smallest element occurs in $I^1,\ldots,I^p$. When $E(I^1,\ldots,I^p) \neq 0$, it means that for some permutations $\{\pi_j,\sigma_j\}_j$, every term $A[\cdot,\cdot]$ must appear even number of times. This implies that the number of distinct elements in $I^1,\ldots,I^p$ is at most half the maximal possible number $pq/2$. This lemma proves the intuition via graph theoretic arguments.

Lemma 9.5.1. *If $E(I^1,\ldots,I^p)\neq 0$, $\#\left(I^1,\ldots,I^p\right) \le \frac{pq}{4}+\frac{q}{2}$.*

The number of $I^1,\ldots,I^p$ that corresponds to a sequence $c^1,\ldots,c^s$ is at most $\frac{n^s}{s!}\cdot\frac{((q/2)!)^p}{\prod_{\ell\in[p]} c^1_\ell!\cdot c^p_\ell!}$. Furthermore, there are at most $2^{O(pq)}p^{pq/2}$ different choices of $c^1,\ldots,c^s$ that corresponds to some $I^1,\ldots,I^p$. The following technical lemma bounds $E(I^1,\ldots,I^p)$ by careful counting arguments.

Lemma 9.5.2. *For any $I^1,\ldots,I^p$, $E(I^1,\ldots,I^p) \le 2^{O(pq)}\frac{p^{5pq/8}}{q^{3pq/8}}\prod_{\ell\in[p]} c^1_\ell!\ldots c^s_\ell!$.*

Summing over all s and multiplying all possibilities,

$$\begin{aligned}
\mathbb{E}\left[\mathsf{Tr}(M^p)\right] &\le \sum_{s=1}^{pq/4+q/2}\left(2^{O(pq)}p^{pq/2}\right)\cdot\left(\frac{n^s}{s!}\cdot((q/2)!)^p\right)\cdot\left(2^{O(pq)}\frac{p^{5pq/8}}{q^{3pq/8}}\right) \\
&= \max_{1\le s\le pq/4+q/2} 2^{O(pq\log p)}\cdot n^s\cdot\frac{q^{pq/8}}{s!}.
\end{aligned}$$

When $q\le n$, the maximum occurs when $s = pq/4+q/2$, so $\mathbb{E}[\mathsf{Tr}(M^p)] \le 2^{O(pq\log p)}\cdot \frac{n^{pq/4+q/2}}{q^{pq/8}}$ as desired.

9.5.2 Overview of Lower Bound Proofs

Let $\mathcal{A}, A, f$ be as in Section 9.5.1. To prove the lower bound in Theorem 9.1.1, we construct a moment matrix M that is positive semidefinite, SoS-symmetric, $\mathsf{Tr}(\mathsf{M}) = 1$, and $\langle A, \mathsf{M}\rangle \geq 2^{-O(d)} \cdot \frac{n^{d/4}}{d^{d/4}}$. At a high level, our construction is $\mathsf{M} := c_1\mathsf{A} + c_2\mathsf{W}$ for some c_1, c_2, where A contains entries of A only corresponding to the multilinear indices, averaged over all SoS-symmetric positions. This gives a large inner product with A, SoS-symmetry, and nice spectral properties even though it is not positive semidefinite. The most natural way to make it positive semidefinite is adding a copy of the identity matrix, but this will again break the SoS-symmetry.

Our main technical contribution here is the construction of W that acts like a *SoS-symmetrized identity*. It has the minimum eigenvalue at least $\frac{1}{2}$, while the trace being $n^{d/2} \cdot 2^{O(d)}$, so the ratio of the average eigenvalue to the minimum eigenvalue is bounded above by $2^{O(d)}$, which allows us to prove a tight lower bound. To the best of our knowledge, no such bound was known for SoS-symmetric matrices except small values of $d = 3, 4$.

Given $I, J \in [n]^{d/2}$, we let $\mathsf{W}[I, J] := \mathbb{E}[x^{\alpha(I)+\alpha(J)}]$, where $x_1, \ldots, x_n$ are independently sampled from the *Wigner semicircle distribution*, whose probability density function is the semicircle $f(x) = \frac{2}{\pi}\sqrt{1-x^2}$. Since $\mathbb{E}[x_1^{\ell}] = 0$ if ℓ is odd and $\mathbb{E}[x_1^{2\ell}] = \frac{1}{\ell+1}\binom{2\ell}{\ell}$, which is the ℓth Catalan number, each entry of W is bounded by $2^{O(d)}$ and $\mathsf{Tr}(\mathsf{W}) \leq n^{d/2} \cdot 2^{O(d)}$. To prove a lower bound on the minimum eigenvalue, we show that for any degree-ℓ polynomial p with m variables, $\mathbb{E}[p(x_1, \ldots, x_m)^2]$ is large by induction on ℓ and m. We use another property of the Wigner semicircle distribution that if $H \in \mathbb{R}^{(d+1)\times(d+1)}$ is the univariate moment matrix of x_1 defined by $H[i, j] = \mathbb{E}[x_1^{i+j}]$ ($0 \leq i, j \leq d$) and $H = (R^T)R$ is the Cholesky decomposition of H, R is an upper triangular matrix with 1's on the main diagonal. This nice Cholesky decomposition allows us to perform the induction on the number of variables while the guarantee on the minimum eigenvalue is independent of n.

9.6 Upper bounds for even degree tensors

For even d, let $\mathcal{A} \in \mathbb{R}^{[n]^d}$ be a d-tensor with i.i.d. ± 1 entries and let $A \in \mathbb{R}^{[n]^{d/2}\times[n]^{d/2}}$ be the matrix flattening of $\mathcal{A}$, i.e., $A[I, J] = \mathcal{A}[I \oplus J]$ (recall that $\oplus$ denotes tuple concatenation). Also let $f(x) := \mathcal{A}(x) = \langle \mathcal{A}, x^{\otimes d}\rangle$. With high probability $f_{\max} = O(\sqrt{n \cdot d \cdot \log d})$. In this section, we prove that for every q divisible by d, with high probability,

$$\left(\Lambda\left(f^{q/d}\right)\right)^{d/q} \;\leq\; \widetilde{O}\left(\frac{n}{q^{1-2/d}}\right)^{d/4} \;=\; \widetilde{O}\left(\frac{n}{q}\right)^{d/4-1/2} \cdot f_{\max}.$$

To prove it, we use the following matrix representation M of $f^{q/d}$, and show that $\|M\|_2 \leq \tilde{O}_d\left(\left(\frac{n\log^5 n}{q^{1-2/d}}\right)^{q/4}\right)$. Given a tuple $I = (i_1, \ldots, i_q)$, and an integer d that divides q and $1 \leq \ell \leq q/d$, let $I_{\ell;d}$ be the d-tuple $(I_{d(\ell-1)+1}, \ldots, I_{d\ell})$ (i.e., if we divide I into q/d tuples of length d, $I_{\ell;d}$ be the ℓ-th tuple). Furthermore, given a tuple $I = (i_1, \ldots, i_q) \in [n]^q$ and a permutation $\pi \in [n]^q$, let $\pi(I)$ be another q-tuple whose ℓth coordinate is $\pi(i_\ell)$. For $I, J \in [n]^{q/2}$, $M[I, J]$ is formally given by

$$\begin{aligned} M[I,J] &= \frac{1}{q!} \cdot \sum_{\pi,\sigma \in \mathbb{S}_{q/2}} A^{\otimes q/d}[\pi(I), \sigma(J)] \\ &= \frac{1}{q!} \cdot \sum_{\pi,\sigma \in \mathbb{S}_{q/2}} \prod_{\ell=1}^{q/d} A[(\pi(I))_{\ell;d/2}, (\sigma(J))_{\ell;d/2}]. \end{aligned}$$

We perform the trace method to bound $\|M\|_2$. Let p be an even integer, that will be eventually taken as $\Theta(\log n)$. $\mathsf{Tr}(M)$ can be written as (let $I^{p+1} := I^1$)

$$\begin{aligned} &\mathbb{E}\left[\sum_{I^1,\ldots,I^p \in [n]^{q/2}} \prod_{\ell=1}^{p} M[I^\ell, I^{\ell+1}] \right] \\ = &\sum_{I^1,\ldots,I^p} \mathbb{E}\left[\prod_{\ell=1}^{p} (\sum_{\pi_j,\sigma_j \in \mathbb{S}_{q/2}} \prod_{m=1}^{q/d} A[(\pi(I^\ell))_{m;d/2}, (\sigma(I^{\ell+1}))_{m;d/2})]) \right]. \end{aligned}$$

Let $E(I^1, \ldots, I^p) := \mathbb{E}[\prod_{\ell=1}^{p} M[I^\ell, I^{\ell+1}]]$, which is the expected value in the right hand side. To analyze $E(I^1, \ldots, I^p)$, we first introduce notions to classify $I^1, \ldots, I^p$ depending on their intersection patterns. For any $I^1, \ldots, I^p \in [n]^{q/2}$, let e_k denote the k-th smallest element in $\bigcup_{\ell,j} \{i_j^\ell\}$. For any $c^1, \ldots, c^s \in [q/2]^p$, let

$$\begin{aligned} &\mathcal{C}(c^1 \ldots c^s) := \\ &\quad \left\{ (I^1, \ldots, I^p) \;\middle|\; \# \left(I^1, \ldots, I^p \right) = s,\ \forall k \in [s], \ell \in [p],\ e_k \text{ appears } c_\ell^k \text{ times in } I^\ell \right\}. \end{aligned}$$

The following two observations on $c^1, \ldots, c^s$ can be easily proved.

Observation 9.6.1. *If $\mathcal{C}(c^1, \ldots, c^s) \neq \varphi$,*

$$\left| \mathcal{C}(c^1, \ldots, c^s) \right| \leq \frac{n^s}{s!} \times \frac{((q/2)!)^p}{\prod_{\ell \in [p]} c_\ell^1! \ldots c_\ell^s!}.$$

Moreover,

$$\left| \left\{ (c^1, \ldots, c^s) \in ([q/2]^p)^s \;\middle|\; \mathcal{C}(c^1, \ldots, c^s) \neq \varphi \right\} \right| \leq 2^{O(pq)} p^{pq/2}.$$

The following lemma bounds $E(I^1, \ldots, I^p)$ in terms of the corresponding $c_1, \ldots, c_s$.

Lemma 9.6.2. *Consider any $c^1, \ldots, c^s \in [q/2]^p$ and $(I^1, \ldots, I^p) \in \mathcal{C}(c^1, \ldots, c^s)$. We have*

$$E(I^1, \ldots, I^p) \leq 2^{O(pq)} \frac{p^{1/2+1/2d}}{q^{1/2-1/2d}} \prod_{\ell \in [p]} c_\ell^1! \ldots c_\ell^s!$$

Proof. Consider any $c^1, \ldots, c^s \in [q/2]^p$ and $(I^1, \ldots, I^p) \in \mathcal{C}(c^1, \ldots, c^s)$. We have

$$\begin{aligned}
&E(I^1, \ldots, I^p) \\
&= \mathbb{E}\left[\prod_{\ell=1}^{p} M[I^\ell, I^{\ell+1}]\right] \\
&= \sum_{\pi_j, \sigma_j \in \mathbb{S}_{q/2}} \mathbb{E}\left[\prod_{\ell=1}^{p}\prod_{m=1}^{q/d} A[(\pi(I^\ell))_{m;d/2}, (\pi(I^{\ell+1}))_{m;d/2}]\right] \\
&= \left(\frac{\prod_\ell \prod_s (c_\ell^s!)^2}{((q/2)!)^{2p}}\right) \cdot \sum_{(J^\ell, K^\ell \in \mathcal{O}(I^\ell))_{\ell\in[p]}} \mathbb{E}\left[\prod_{\ell=1}^{p}\prod_{m=1}^{q/d} A[J^\ell_{m;d/2}, K^{\ell+1}_{m;d/2}]\right]
\end{aligned} \tag{9.2}$$

Thus, $E(I^1, \ldots, I^p)$ is bounded by the number of choices for $J^1, \ldots, J^p, K^1, \ldots, K^p$ such that $J^\ell, K^\ell \in \mathcal{O}(I^\ell)$ for each $\ell \in [p]$, and $\mathbb{E}[\prod_{\ell=1}^{p}\prod_{m=1}^{q/d} A[J^\ell_{m;d/2}, K^{\ell+1}_{m;d/2}]]$ is nonzero.

Given $J^1, \ldots, J^p$ and $K^1, \ldots, K^p$, consider the (pq/d)-tuple T where each coordinate is indexed by $(\ell, m)_{\ell\in[p], m\in[q/d]}$ and has a d-tuple $T_{\ell,m} := (J^\ell_{m;d/2}) \oplus (K^{\ell+1}_{m;d/2}) \in \mathbb{R}^d$ as a value. Note that $\sum_{\ell,m} \alpha(T_{\ell,m})) = (2o_1, \ldots, 2o_n)$ where o_r is the number of occurences of $r \in [n]$ in $(pq/2)$-tuple $\oplus_{\ell=1}^{p} I^\ell$. The fact that $\mathbb{E}[\prod_{\ell=1}^{p}\prod_{m=1}^{q/d} A[j_{m;d/2}, k_{m;d/2}]] \neq 0$ means that every d-tuple occurs even number of times in T.

We count the number of (pq/d)-tuples $T = (T_{\ell,m})_{\ell\in[p], m\in[q]}$ that $\sum_{\ell,m} \alpha(T_{\ell,m}) = (2o_1, \ldots, 2o_n)$ and every d-tuple occurs an even number of times. Let $Q = (Q_1, \ldots, Q_{pq/2d})$, $R = (R_1, \ldots, R_{pq/2d})$ be two $(pq/2d)$-tuples of d-tuples where for every d-tuple P, the number of occurences of P is the same in Q and R, and $\sum_{\ell=1}^{pq/2d} \alpha(Q_\ell) = \sum_{\ell=1}^{pq/2d} \alpha(R_\ell) = (o_1, \ldots, o_n)$. At most $2^{pq/d}$ tuples T can be made by *interleaving* Q and R — for each (ℓ, m), choose $T_{\ell,m}$ from the first unused d-tuple in either Q or R. Furthermore, every tuple T that meets our condition can be constructed in this way.

Due to the condition $\sum_{\ell=1}^{pq/2d} \alpha(Q_\ell) = (o_1, \ldots, o_n)$, the number of choices for Q is at most the number of different ways to permute $I^1 \oplus \cdots \oplus I^p$, which is at most $(pq/2)!/\prod_{m\in[s]}(\bar{c}^m)!$, where $\bar{c}^m := \sum_{\ell\in[p]} c^m_\ell$ for $m \in [s]$. For a fixed choice of Q, there are at most $(pq/2d)!$ choices of R. Therefore, the number of choices for $(J^\ell, K^\ell \in \mathcal{O}(I^\ell))_{\ell\in[p]}$ with nonzero expected value is at most

$$2^{pq/d} \cdot \frac{(pq/2)!}{\prod_{m\in[s]}(\bar{c}^m)!} \cdot (pq/2d)! = 2^{O(pq)} \cdot \frac{(pq)^{1/2+1/2d}}{\prod_{m\in[s]}(\bar{c}^m)!}.$$

Combining with Eq. (9.2),

$$E(I^1, \ldots, I^p) \leq \left(2^{O(pq)} \frac{(pq)^{1/2+1/2d}}{\prod_{m\in[s]}(\bar{c}^m)!}\right) \cdot \left(\frac{\prod_\ell \prod_s (c_\ell^s!)^2}{((q/2)!)^{2p}}\right) \leq 2^{O(pq)} \cdot \frac{p^{1/2+1/2d}}{q^{1/2-1/2d}} \cdot \prod_\ell \prod_s c^s_\ell!$$

as desired. ∎

Lemma 9.6.3. *For all $I^1, \ldots, I^p \in [n]^{q/2}$, if $E(I^1, \ldots, I^p) \neq 0$, $\#\left(I^1, \ldots, I^p\right) \leq \frac{pq}{4} + \frac{q}{2}$.*

Proof. Note that $E(I^1,\dots,I^p) \neq 0$ implies that there exist $J^1,\dots,J^p,K^1,\dots,K^p$ such that $J^\ell, K^\ell \in \mathcal{O}(I^\ell)$ and every d-tuple occurs exactly even number of times in $((J^\ell_{m;d/2}) \oplus (K^{\ell+1}_{m;d/2}))_{\ell\in[p],m\in[q/d]}$. Consider the graph $G = (V, E)$ defined by

$$V := \bigcup_{\ell\in[p]} \bigcup_{k\in[q/2]} \left\{ I^\ell_k \right\}$$
$$E := \bigcup_{m\in[q/2]} \left\{ \left\{ J^1_m, K^2_m \right\}, \left\{ J^2_m, K^3_m \right\}, \dots, \left\{ J^p_m, K^1_m \right\} \right\}.$$

The even multiplicity condition implies that every element in E has even multiplicity and consequently $|E| \le pq/4$. We next show that E is the union of $q/2$ paths. To this end, we construct $G^1 \in \mathcal{O}(I^1), \dots, G^\ell \in \mathcal{O}(I^\ell)$ as follows:

1. Let $G^2 := K^2$
2. For $3 \le \ell \le p$ do:

   Since $G^\ell \in \mathcal{O}(J^\ell)$, there exists $\pi \in \mathcal{S}_{q/2}$ s.t. $\pi(J^\ell) = G^\ell$.
   Let $G^{\ell+1} := \pi(K^{\ell+1})$.

We observe that by construction,

$$\bigcup_{m\in[q/2]} \left\{ \left\{ J^1_m, G^2_m \right\}, \left\{ G^2_m, G^3_m \right\}, \dots, \left\{ G^p_m, G^1_m \right\} \right\}$$
$$= \bigcup_{m\in[q/2]} \left\{ \left\{ J^1_m, K^2_m \right\}, \left\{ J^2_m, K^3_m \right\}, \dots, \left\{ J^p_m, K^1_m \right\} \right\} = E$$

which establishes that E is a union of $q/2$ paths.

Now since E is the union of $q/2$ paths G has at most $q/2$ connected components, and one needs to add at most $q/2-1$ edges make it connected, we have $|V| \le |E| + (q/2 - 1) + 1 \le pq/4 + q/2$. But $\#\left(I^1,\dots,I^p\right) = |V|$, which completes the proof. ■

Finally, $\mathbb{E}[\mathsf{Tr}(M^p)]$ can be bounded as follows.

$$\begin{aligned}
&\mathbb{E}\left[\mathsf{Tr}(M^p)\right] \\
&= \sum_{I^1,\dots,I^p\in[n]^{q/2}} E(I^1,\dots,I^p) \\
&= \sum_{s\in[pq/4+q/2]} \sum_{\#\left(I^1,\dots,I^p\right)=s} E(I^1,\dots,I^p) && \text{(by Lemma 9.6.3)} \\
&= \sum_{s\in[pq/4+q/2]} \sum_{c^1,\dots,c^s\in[q/2]^p} \sum_{(I^1,\dots,I^p)\in\mathcal{C}(c^1\dots c^s)} E(I^1,\dots,I^p) \\
&= \sum_{s\in[pq/4+q/2]} \sum_{c^1,\dots,c^s\in[q/2]^p} \sum_{(I^1,\dots,I^p)\in\mathcal{C}(c^1\dots c^s)} E(I^1,\dots,I^p)
\end{aligned}$$

$$\le \sum_{s \in [pq/4+q/2]} \sum_{c^1,\dots,c^s \in [q/2]^p}$$

$$\sum_{(I^1,\dots,I^p) \in \mathcal{C}(c^1 \dots c^s)} 2^{O(pq)} \frac{p^{(1/2+1/2d)pq}}{q^{(1/2-1/2d)pq}} \prod_{\ell \in [p]} c^1_\ell! \dots c^s_\ell! \qquad \text{(by Eq. (9.2))}$$

$$\le \sum_{s \in [pq/4+q/2]} 2^{O(pq)} \frac{n^s}{s!} p^{(1+1/2d)pq} q^{pq/2d} \qquad \text{(by Observation 9.6.1)}$$

$$\le \sum_{s \in [pq/4+q/2]} 2^{O(pq)} \frac{n^{pq/4+q/2}}{s!\, q^{pq/4+q/2-s}} p^{(1/2+1/2d)p1} q^{(1/2-1/2d)pq} \qquad \text{(assuming } q \le n\text{)}$$

$$\le \sum_{s \in [pq/4+q/2]} 2^{O(pq)} \frac{n^{pq/4+q/2}\, p^{(1+1/2d)pq}}{q^{(1/4-1/2d)pq}}$$

$$\le 2^{O(pq)} \frac{n^{pq/4+q/2}\, p^{(1+1/2d)pq}}{q^{(1/4-1/2d)pq}}.$$

Choose p to be even and let $p = \Theta(\log n)$. Applying Markov inequality shows that with high probability,

$$\left(\Lambda\left(f^{q/d}\right)\right)^{d/q} \le (\|M\|_2)^{d/q} \le (\mathbb{E}\left[\mathrm{Tr}(M^p)\right])^{d/pq} = O_d\left(\frac{n^{d/4} \cdot (\log n)^{\,d+1/2}}{q^{\,d/4-1/2}}\right).$$

Thus we obtain

Theorem 9.6.4. *For even d, let $\mathcal{A} \in \mathbb{R}^{[n]^d}$ be a d-tensor with i.i.d. ± 1 entries. Then for any even q such that $q \le n$, we have that with probability $1 - n^{\Omega(1)}$,*

$$\frac{SoS_q(\mathcal{A}(x))}{\mathcal{A}_{\max}} \le \left(\frac{\widetilde{O}(n)}{q}\right)^{d/4-1/2}.$$

Remark. For the special case where $q = d$, we prove a stronger upper bound, namely

$$\frac{\mathsf{SoS}_q(\mathcal{A}(x))}{\mathcal{A}_{\max}} \le \left(\frac{O(n)}{q}\right)^{d/4-1/2},$$

the proof of which is implicit in the proofs of Lemma 9.7.5 and Lemma 9.7.6.

9.7 Proof of SoS Lower Bound in Theorem 9.1.1

For even q, let $\mathcal{A} \in \mathbb{R}^{[n]^q}$ be a q-tensor with i.i.d. ± 1 entries and let $A \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ be the matrix flattening of $\mathcal{A}$, i.e., $A[I, J] = \mathcal{A}[I \oplus J]$ (recall that $\oplus$ denotes tuple concatenation). Also let $f(x) := \mathcal{A}(x) = \langle \mathcal{A}, x^{\otimes q} \rangle$. This section proves the lower bound in Theorem 9.1.1, by constructing a moment matrix M that is positive semidefinite, SoS-symmetric, $\mathrm{Tr}(\mathsf{M}) = 1$, and $\langle A, \mathsf{M} \rangle \ge 2^{-O(q)} \cdot \frac{n^{q/4}}{q^{q/4}}$. In Section 9.7.1, we construct the matrix $\widehat{\mathsf{W}}$ that acts as a SoS-symmetrized identity matrix. The moment matrix M is presented in Section 9.7.2.

9.7.1 Wigner Moment Matrix

In this section, we construct an SoS-symmetric and positive semidefinite matrix $\widehat{\mathsf{W}} \in \mathbb{R}^{\mathbb{N}^n_{q/2} \times \mathbb{N}^n_{q/2}}$ such that $\lambda_{\min}(\widehat{\mathsf{W}}) / \operatorname{Tr}\left(\widehat{\mathsf{W}}\right) \geq 1/(2^{q+1} \cdot |\mathbb{N}^n_{q/2}|)$, i.e. the ratio of the minimum eigenvalue to the average eigenvalue is at least $1/2^{q+1}$.

Theorem 9.7.1. *For any positive integer n and any positive even integer q, there exists a matrix $\widehat{\mathsf{W}} \subseteq \mathbb{R}^{\mathbb{N}^n_{q/2} \times \mathbb{N}^n_{q/2}}$ that satisfies the following three properties: (1) $\widehat{\mathsf{W}}$ is degree-q SoS symmetric. (2) The minimum eigenvalue of $\widehat{\mathsf{W}}$ is at least $\frac{1}{2}$. (3) Each entry of $\widehat{\mathsf{W}}$ is in $[0, 2^q]$.*

Theorem 9.7.1 is proved by explicitly constructing independent random variables $x_1, \ldots, x_n$ such that for any n-variate polynomial $p(x_1, \ldots, x_n)$ of degree at most $\frac{q}{2}$, $\mathbb{E}[p^2]$ is bounded away from 0. The proof consists of three parts. The first part shows the existence of a desired distribution for one variable x_i. The second part uses induction to prove that $\mathbb{E}[p^2]$ is bounded away from 0. The third part constructs $\widehat{\mathsf{W}} \subseteq \mathbb{R}^{\mathbb{N}^n_{q/2} \times \mathbb{N}^n_{q/2}}$ from the distribution defined.

Wigner Semicircle Distribution and Hankel Matrix. Let k be a positive integer. In this part, the rows and columns of all $(k+1) \times (k+1)$ matrices are indexed by $\{0, 1, \ldots, k\}$. Let T be a $(k+1) \times (k+1)$ matrix where $T[i,j] = 1$ if $|i - j| = 1$ and $T[i,j] = 0$ otherwise. Let $e_0 \in \mathbb{R}^{k+1}$ be such that $(e_0)_0 = 1$ and $(e_0)_i = 0$ for $1 \leq i \leq k$. Let $R \in \mathbb{R}^{(k+1)\times(k+1)}$ be defined by $R := [e_0, Te_0, T^2e_0, \ldots, T^k e_0]$. Let $R_0, \ldots, R_k$ be the columns or R so that $R_i = T^i e_0$. It turns out that R is closely related to the number of ways to consistently put parentheses. Given a string of parentheses '(' or ')', we call it *consistent* if any prefix has at least as many '(' as ')'. For example, ((())(is consistent, but ())((is not.

Claim 9.7.2. *$R[i,j]$ is the number of ways to place j parentheses '(' or ')' consistently so that there are i more '(' than ')'.*

Proof. We proceed by the induction on j. When $j = 0$, $R[0,0] = 1$ and $R[i,0] = 0$ for all $i \geq 1$. Assume the claim holds up to $j-1$. By the definition $R_j = TR_{j-1}$.

- For $i = 0$, the last parenthesis must be the close parenthesis, so the definition $R[0,j] = R[1, j-1]$ still measures the number of ways to place j parentheses with equal number of '(' and ')'.

- For $i = k$, the last parenthesis must be the open parenthesis, so the definition $R[k,j] = R[k-1, j-1]$ still measures the number of ways to place j parentheses with k more '('.

- For $0 < i < k$, the definition of R gives $R[i,j] = R[i-1, j-1] + R[i+1, j-1]$. Since $R[i-1, j]$ corresponds to placing ')' in the jth position and $R[i+1, j]$ corresponds to placing '(' in the jth position, $R[i,j]$ still measures the desired quantity.

This completes the induction and proves the claim. ∎

Easy consequences of the above claim are (1) $R[i,i] = 1$ for all $0 \le i \le k$, and $R[i,j] = 0$ for $i > j$, and (2) $R[i,j] = 0$ if $i+j$ is odd, and $R[i,j] \ge 1$ if $i \le j$ and $i+j$ is even.

Let $H := (R^T)R$. Since R is upper triangular with 1's on the main diagonal, $H = (R^T)R$ gives the unique Cholesky decomposition, so H is positive definite. It is easy to see that $H[i,j] = \langle R_i, R_j \rangle$ is the total number of ways to place $i+j$ parentheses consistently with the same number of '(' and ')'. Therefore, $H[i,j] = 0$ if $i+j$ is odd, and if $i+j$ is even (let $l := \frac{i+j}{2}$), $H[i,j]$ is the lth Catalan number $C_l := \frac{1}{l+1}\binom{2l}{l}$. In particular, $H[i,j] = H[i',j']$ for all $i+j = i'+j'$. Such H is called a *Hankel matrix*.

Given a sequence of $m_0 = 1, m_1, m_2, \dots$ of real numbers, the *Hamburger moment problem* asks whether there exists a random variable W supported on $\mathbb{R}$ such that $\mathbb{E}[W^i] = m_i$. It is well-known that there exists a unique such W if for all $k \in \mathbb{N}$, the Hankel matrix $H \in \mathbb{R}^{(k+1)\times(k+1)}$ defined by $H[i,j] := \mathbb{E}[W^{i+j}]$ is positive definite [Sim98]. Since our construction of $H \in \mathbb{R}^{(k+1)\times(k+1)}$ ensures its positive definiteness for any $k \in \mathbb{N}$, there exists a unique random variable W such that $\mathbb{E}[W^i] = 0$ if i is odd, $\mathbb{E}[W^i] = C_{\frac{i}{2}}$ if i is even. It is known as the *Wigner semicircle distribution* with radius $R = 2$.

Remark 9.7.3. *Some other distributions (e.g., Gaussian) will give an asymptotically weaker bound. Let G be a standard Gaussian random variable. The quantitative difference comes from the fact that* $\mathbb{E}[W^{2l}] = C_l = \frac{1}{l+1}\binom{2l}{l} \le 2^l$ *while* $\mathbb{E}[G^{2l}] = (2l-1)!! \ge 2^{\Omega(l \log l)}$.

Multivariate Distribution. Fix n and q. Let $k = \frac{q}{2}$. Let $H \in \mathbb{R}^{(k+1)\times(k+1)}$ be the Hankel matrix defined as above, and W be a random variable sampled from the Wigner semicircle distribution. Consider $x_1, \dots, x_n$ where each x_i is an independent copy of $\frac{W}{N}$ for some large number N to be determined later. Our $\widehat{\mathrm{W}}$ is later defined to be $\widehat{\mathrm{W}}[\alpha,\beta] = \mathbb{E}[x^{\alpha+\beta}] \cdot N^q$ so that the effect of the normalization by N is eventually canceled, but large N is needed to prove the induction that involves non-homogeneous polynomials.

We study $\mathbb{E}[p(x)^2]$ for any n-variate (possibly non-homogeneous) polynomial p of degree at most k. For a multivarite polynomial $p = \sum_{\alpha \in \mathbb{N}^n_{\le k}} p_\alpha x^\alpha$, define ℓ_2 norm of p to be $\|p\|_{\ell_2} := \sqrt{\sum_\alpha p_\alpha^2}$. For $0 \le m \le n$ and $0 \le l \le k$, let $\sigma(m,l) := \inf_p \mathbb{E}[p(x)^2]$ where the infimum is taken over polynomials p such that $\|p\|_{\ell_2} = 1$, $\deg(p) \le l$, and p depends only on $x_1, \dots, x_m$.

Lemma 9.7.4. *There exists $N := N(n,k)$ such that $\sigma(m,l) \ge \frac{(1-\frac{m}{2n})}{N^{2l}}$ for all $0 \le m \le n$ and $0 \le l \le k$.*

Proof. We prove the lemma by induction on m and l. When $m = 0$ or $l = 0$, p becomes the constant polynomial 1 or -1, so $\mathbb{E}[p^2] = 1$.

Fix $m, l > 0$ and a polynomial $p = p(x_1, \dots, x_m)$ of degree at most l. Decompose $p = \sum_{i=0}^{l} p_i x_m^i$ where each p_i does not depend on x_m. The degree of p_i is at most $l - i$.

$$\mathbb{E}[p^2] = \mathbb{E}[(\sum_{i=0}^{l} p_i x_m^i)^2] = \sum_{0 \le i,j \le l} \mathbb{E}[p_i p_j]\mathbb{E}[x_m^{i+j}].$$

Let $\Sigma = \text{diag}(1, \frac{1}{N}, \ldots, \frac{1}{N^l}) \in \mathbb{R}^{(l+1)\times(l+1)}$. Let $H_l \in \mathbb{R}^{(l+1)\times(l+1)}$ be the submatrix of H with the first $l+1$ rows and columns. The rows and columns of $(l+1)\times(l+1)$ matrices are still indexed by $\{0,\ldots,l\}$. Define $R_l \in \mathbb{R}^{(l+1)\times(l+1)}$ similarly from R, and r_t $(0 \le t \le l)$ be the tth column of $(R_l)^T$. Note $H_l = (R_l)^T R_l = \sum_{t=0}^{l} r_t r_t^T$. Let $H' = \Sigma H_l \Sigma$ such that $H'[i,j] = \mathbb{E}[x_m^{i+j}]$. Finally, let $P \in \mathbb{R}^{(l+1)\times(l+1)}$ be defined such that $P[i,j] := \mathbb{E}[p_i p_j]$. Then $\mathbb{E}[p^2]$ is equal to

$$\begin{aligned}\mathsf{Tr}\big(PH'\big) &= \mathsf{Tr}(P\Sigma H_l \Sigma) = \mathsf{Tr}\left(P\Sigma(\sum_{t=0}^{l} r_t r_t^T)\Sigma\right)\\ &= \sum_{t=0}^{l} \mathbb{E}[(p_t \frac{1}{N^t} + p_{t+1}\frac{(r_t)_{t+1}}{N^{t+1}} + \cdots + p_l \frac{(r_t)_l}{N^l})^2],\end{aligned}$$

where the last step follows from the fact that $(r_t)_j = 0$ if $j < t$ and $(r_t)_t = 1$. Consider the polynomial

$$q_t := p_t \frac{1}{N^t} + p_{t+1}\frac{(r_t)_{t+1}}{N^{t+1}} + \cdots + p_l \frac{(r_t)_l}{N^l}.$$

Since p_i is of degree at most $l-i$, q_t is of degree at most $l-t$. Also recall that each entry of R is bounded by 2^k. By the triangle inequality,

$$\|q_t\|_{\ell_2} \ge \frac{1}{N^t}\left(\|p_t\|_{\ell_2} - (\|p_{t+1}\|_{\ell_2}\frac{(r_t)_{t+1}}{N} + \cdots + \|p_l\|_{\ell_2}\frac{(r_t)_l}{N^{l-t}})\right) \ge \frac{1}{N^t}\left(\|p_t\|_{\ell_2} - \frac{k2^k}{N}\right),$$

and

$$\|q_t\|_{\ell_2}^2 \ge \frac{1}{N^{2t}}\left(\|p_t\|_{\ell_2}^2 - \frac{2k2^k}{N}\right).$$

Finally,

$$\begin{aligned}\mathbb{E}[p^2] &= \sum_{t=0}^{l} \mathbb{E}[q_t^2]\\ &\ge \sum_{t=0}^{l} \sigma(m-1, l-t)\cdot \|q_t\|_{\ell_2}^2\\ &\ge \sum_{t=0}^{l} \sigma(m-1, l-t)\cdot \frac{1}{N^{2t}}\left(\|p_t\|_{\ell_2}^2 - \frac{2k2^k}{N}\right)\\ &\ge \sum_{t=0}^{l} \frac{(1-\frac{m-1}{2n})}{N^{2l-2t}}\cdot\frac{1}{N^{2t}}\cdot\left(\|p_t\|_{\ell_2}^2 - \frac{2k2^k}{N}\right)\\ &= \frac{(1-\frac{m-1}{2n})}{N^{2l}}\cdot\sum_{t=0}^{l}\left(\|p_t\|_{\ell_2}^2 - \frac{2k2^k}{N}\right)\\ &\ge \frac{(1-\frac{m-1}{2n})}{N^{2l}}\cdot(1-\frac{2K^2 2^k}{N}).\end{aligned}$$

Take $N := 4nK^2 2^k$ so that $\big(1-\frac{m-1}{2n}\big)\cdot\big(1-\frac{2K^2 2^k}{N}\big) \ge 1 - \frac{m-1}{2n} - \frac{2K^2 2^k}{N} = 1 - \frac{m}{2n}$. This completes the induction and proves the lemma. ■

Construction of $\widehat{\mathsf{W}}$. We now prove Theorem 9.7.1. Given n and q, let $k = \frac{q}{2}$, and consider random variables $x_1, \ldots, x_n$ above. Let $\widehat{\mathsf{W}} \in \mathbb{R}^{\mathbb{N}_k^n \times \mathbb{N}_k^n}$ be such that for any $\alpha, \beta \in \mathbb{N}_k^n$, $\widehat{\mathsf{W}}[\alpha, \beta] = \mathbb{E}[x^{\alpha+\beta}] \cdot N^{2k}$. By definition, $\widehat{\mathsf{W}}$ is degree-q SoS symmetric. Since each entry of $\widehat{\mathsf{W}}$ corresponds to a monomial of degree exactly q and each x_i is drawn independently from the Wigner semicircle distribution, each entry of $\widehat{\mathsf{W}}$ is at most the $\frac{q}{2}$th Catalan number $C_{\frac{q}{2}} \leq 2^q$. For any unit vector $p = (p_S)_{S \in \mathbb{N}_k^n} \in \mathbb{R}^{\mathbb{N}_k^n}$, Lemma 9.7.4 shows $p^T \widehat{\mathsf{W}} p = \mathbb{E}[p^2] \cdot N^{2k} \geq \frac{1}{2}$ where p also represents a degree-k homogeneous polynomial $p(x_1, \ldots, x_n) = \sum_{\alpha \in \binom{[n]}{k}} p_\alpha x^\alpha$. Therefore, the minimum eigenvalue of $\widehat{\mathsf{W}}$ is at least $\frac{1}{2}$.

9.7.2 Final Construction

For even d, let $\mathcal{A} \in \mathbb{R}^{[n]^q}$ be a q-tensor with i.i.d. ± 1 entries and let $A \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ be the matrix flattening of $\mathcal{A}$, i.e., $A[I, J] = \mathcal{A}[I \oplus J]$ (recall that $\oplus$ denotes tuple concatenation). Also let $f(x) := \mathcal{A}(x) = \langle \mathcal{A}, x^{\otimes q} \rangle$. Our lower bound on $f_{\max}$ by is proved by constructing a moment matrix $\mathsf{M} \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ that satisfies

- $\mathsf{Tr}(\mathsf{M}) = 1$.
- $\mathsf{M} \succeq 0$.
- M is SoS-symmetric.
- $\langle A, \mathsf{M} \rangle \;\geq\; 2^{-O(q)} \cdot n^{q/4} / q^{q/4}$,

where $A \in \mathbb{R}^{[n]^{q/2} \times [n]^{q/2}}$ is any matrix representation of f (SoS-symmetry of M ensures $\langle A, \mathsf{M} \rangle$ does not depend on the choice of A).

Let A be the SoS-symmetric matrix such that for any $I = (i_1, \ldots, i_{q/2})$ and $J = (j_1, \ldots, j_{q/2})$,

$$\mathsf{A}[I, J] = \begin{cases} \frac{f_{\alpha(I)+\alpha(J)}}{q!}, & \text{if } i_1, \ldots, i_{q/2}, j_1, \ldots, j_{q/2} \text{ are all distinct.} \\ 0 & \text{otherwise.} \end{cases}$$

We bound $\|\mathsf{A}\|_2$ in two steps. Let $\widehat{\mathsf{A}}_Q \in \mathbb{R}^{\mathbb{N}_{q/2}^n \times \mathbb{N}_{q/2}^n}$ be the *quotient matrix* of A defined by

$$\widehat{\mathsf{A}}_Q[\beta, \gamma] := \mathsf{A}[I, J] \cdot \sqrt{|\mathcal{O}(\beta)| \cdot |\mathcal{O}(\gamma)|},$$

where $I, J \in [n]^{q/2}$ are such that $\beta = \alpha(I), \gamma = \alpha(J)$.

Lemma 9.7.5. *With high probability,* $\|\widehat{\mathsf{A}}_Q\|_2 \leq \; 2^{O(q)} \cdot \frac{n^{q/4}}{q^{q/4}}$.

Proof. Consider any $y \in \mathbb{R}^{\mathbb{N}_{q/2}^n}$ s.t. $\|y\| = 1$. Since

$$y^T \cdot \widehat{\mathsf{A}}_Q \cdot y = \sum_{\beta + \gamma \leq \mathbb{1}} \widehat{A}_Q[\beta, \gamma] \cdot y_\beta \cdot y_\gamma$$

$$= \sum_{\beta+\gamma \le \mathbb{1}} y_\beta \cdot y_\gamma \sum_{\substack{\alpha(I)+\alpha(J)\\ =\beta+\gamma}} A[I,J] \cdot \frac{\sqrt{|\mathcal{O}(\beta)||\mathcal{O}(\gamma)|}}{|\mathcal{O}(\beta+\gamma)|}$$

$$= \sum_{I,J\in[n]^{q/2}} A[I,J] \sum_{\substack{\beta+\gamma\le\mathbb{1}\\ \beta+\gamma=\\ \alpha(I)+\alpha(J)}} \frac{\sqrt{|\mathcal{O}(\beta)||\mathcal{O}(\gamma)|}}{|\mathcal{O}(\beta+\gamma)|} \cdot y_\beta \cdot y_\gamma$$

So $y^T \cdot \widehat{\mathsf{A}}_Q \cdot y$ is a sum of independent random variables

$$\sum_{I,J\in[n]^q} A[I,J] \cdot c_{I,J}$$

where each $A[I,J]$ is independently sampled from the Rademacher distribution and

$$c_{I,J} := \sum_{\substack{\beta+\gamma\le\mathbb{1}\\ \beta+\gamma=\\ \alpha(I)+\alpha(J)}} \frac{\sqrt{|\mathcal{O}(\beta)||\mathcal{O}(\gamma)|}}{|\mathcal{O}(\beta+\gamma)|} \cdot y_\beta \cdot y_\gamma \,.$$

Fix any $I,J \in [n]^{q/2}$ and let $\alpha := \alpha(I)+\alpha(J)$. By Cauchy-Schwarz,

$$c_{I,J}^2 \;\le\; \left(\sum_{\beta+\gamma=\alpha} \frac{|\mathcal{O}(\beta)||\mathcal{O}(\gamma)|}{|\mathcal{O}(\alpha)|^2}\right) \cdot \Big(\sum_{\beta+\gamma=\alpha} y_\beta^2 \cdot y_\gamma^2\Big) \;\le\; \frac{2^{O(q)}}{|\mathcal{O}(\alpha)|} \cdot \sum_{\beta+\gamma=\alpha} y_\beta^2 \cdot y_\gamma^2 \;=:\; c_\alpha^2\,, \tag{9.3}$$

since there are at most $2^{O(q)}$ choices of β and γ with $\beta+\gamma=\alpha$, and $|\mathcal{O}(\beta)|\cdot|\mathcal{O}(\gamma)| \le |\mathcal{O}(\alpha)|$. Therefore, $y^T \cdot \widehat{\mathsf{A}}_Q \cdot y$ is the sum of independent random variables that are centred and always lie in the interval $[-1,+1]$. Furthermore, by Eq. (9.3), the total variance is

$$\sum_{I,J\in[n]^{q/2}} c_{I,J}^2 \;\le\; \sum_{\alpha\in\mathbb{N}_q^n} c_\alpha^2 \cdot |\mathcal{O}(\alpha)| \;\le\; 2^{O(q)} \cdot \sum_{\beta,\gamma\in\mathbb{N}_{q/2}^n} y_\beta^2 \cdot y_\gamma^2 \;=\; 2^{O(q)} \cdot \Big(\sum_{\beta\in\mathbb{N}_{q/2}^n} y_\beta^2\Big)^2 \;=\; 2^{O(q)}$$

The claim then follows from combining standard concentration bounds with a union bound over a sufficiently fine net of the unit sphere in $|\mathbb{N}_{q/2}^n| \le 2^{O(q)} \cdot \frac{n^{q/2}}{q^{q/2}}$ dimensions. ■

Lemma 9.7.6. *For any SoS-symmetric* $\mathsf{A} \in \mathbb{R}^{[n]^{q/2}\times[n]^{q/2}}$, $\|\mathsf{A}\|_2 \le \|\widehat{\mathsf{A}}_Q\|_2$.

Proof. For any $u,v \in \mathbb{R}^{[n]^{q/2}}$ s.t. $\|u\| = \|v\| = 1$, we have

$$\begin{aligned} & u^T \mathsf{A} v \\ &= \sum_{I,J\in[n]^{q/2}} \mathsf{A}[I,J] u_I v_J \\ &= \sum_{I,J\in[n]^{q/2}} \frac{\widehat{\mathsf{A}}_Q[\alpha(I),\alpha(J)]}{\sqrt{|\mathcal{O}(I)|\,|\mathcal{O}(J)|}} \cdot u_I v_J \end{aligned}$$

$$
\begin{aligned}
&= \sum_{\alpha,\beta\in\mathbb{N}^n_{q/2}} \frac{\mathsf{A}[\alpha,\beta]}{\sqrt{|\mathcal{O}(\alpha)|\,|\mathcal{O}(\beta)|}} \langle u|_{\mathcal{O}(\alpha)}, \mathbf{1}\rangle \langle v|_{\mathcal{O}(\beta)}, \mathbf{1}\rangle \\
&= a^T \widehat{\mathsf{A}}_Q\, b \qquad \text{where } a_\alpha := \frac{\langle u|_{\mathcal{O}(\alpha)}, \mathbf{1}\rangle}{\sqrt{|\mathcal{O}(\alpha)|}},\ b_\alpha := \frac{\langle v|_{\mathcal{O}(\alpha)}, \mathbf{1}\rangle}{\sqrt{|\mathcal{O}(\alpha)|}} \\
&\le \|\widehat{\mathsf{A}}_Q\|_2 \|a\| \cdot \|b\| \\
&= \|\widehat{\mathsf{A}}_Q\|_2 \sqrt{\sum_{\alpha\in\mathbb{N}^n_{q/2}} \frac{\langle u|_{\mathcal{O}(\alpha)}, \mathbf{1}\rangle^2}{|\mathcal{O}(\alpha)|}} \sqrt{\sum_{\alpha\in\mathbb{N}^n_{q/2}} \frac{\langle v|_{\mathcal{O}(\alpha)}, \mathbf{1}\rangle^2}{|\mathcal{O}(\alpha)|}} \\
&\le \|\widehat{\mathsf{A}}_Q\|_2 \sqrt{\sum_{\alpha\in\mathbb{N}^n_{q/2}} \|u|_{\mathcal{O}(\alpha)}\|^2} \sqrt{\sum_{\alpha\in\mathbb{N}^n_{q/2}} \|u|_{\mathcal{O}(\alpha)}\|^2} \qquad \text{(by Cauchy-Schwarz)} \\
&\le \|\widehat{\mathsf{A}}_Q\|_2 \|u\| \cdot \|v\| = \|\widehat{\mathsf{A}}_Q\|_2.
\end{aligned}
$$

■

The above two lemmas imply that $\|\mathsf{A}\|_2 \le \|\widehat{\mathsf{A}}_Q\|_2 \le 2^{O(q)} \cdot \frac{n^{q/4}}{q^{q/4}}$. Our moment matrix M is defined by

$$
\mathsf{M} := \frac{1}{c_1}\left(\frac{1}{c_2} \cdot \frac{q^{3q/4}}{n^{3q/4}}\mathsf{A} + \frac{\mathsf{W}}{n^{q/2}}\right),
$$

where W is the direct extension of $\widehat{\mathsf{W}}$ constructed in Theorem 9.7.1 — $\mathsf{W}[I,J] := \widehat{\mathsf{W}}[\alpha(I),\alpha(J)]$ for all $I,J \in [n]^{q/2}$, and $c_1, c_2 = 2^{\Theta(q)}$ that will be determined later.

We first consider the trace of M. The trace of A is 0 by design, and the trace of W is $n^{q/2} \cdot 2^{O(q)}$. Therefore, the trace of M can be made 1 by setting c_1 appropriately. Since both A and W are SoS-symmetric, so is M. Since $\mathbb{E}[\mathsf{W}, A] = 0$ and for each $I,J \in [n]^{q/2}$ with $i_1,\ldots,i_{q/2}, j_1,\ldots,j_{q/2}$ all distinct we have $\mathbb{E}[\mathsf{A}[I,J]A[I,J]] = \frac{1}{q!}$, with high probability

$$
\langle A, \mathsf{M}\rangle = \frac{1}{c_1} \cdot \langle A, \left(\frac{1}{c_2} \cdot \frac{q^{3q/4}}{n^{3q/4}}\mathsf{A} + \frac{\mathsf{W}}{n^{q/2}}\right)\rangle \ge 2^{O(-q)} \cdot \frac{q^{3q/4}}{n^{3q/4}} \cdot \frac{n^q}{q^q} = 2^{O(-q)} \cdot \frac{n^{q/4}}{q^{q/4}}.
$$

It finally remains to show that M is positive semidefinite. Take an arbitrary vector $v \in \mathbb{R}^{[n]^{q/2}}$, and let

$$
p = \sum_{\alpha\in\mathbb{N}^n_{q/2}} x^\alpha p_\alpha = \sum_{\alpha\in\mathbb{N}^n_{q/2}} x^\alpha \cdot \left(\sum_{I\in[n]^{q/2}:\alpha(I)=\alpha} v_I\right)
$$

be the associated polynomial. If $p = 0$, SoS-symmetry of M ensures $v\mathsf{M}v^T = 0$. Normalize v so that $\|p\|_{\ell_2} = 1$. First, consider another vector $v_m \in [n]^{q/2}$ such that

$$
(v_m)_I = \begin{cases} \frac{p^{\alpha(I)}}{(q/2)!}, & \text{if } i_1,\ldots,i_{q/2} \text{ are all distinct.} \\ 0 & \text{otherwise.} \end{cases}
$$

Then

$$
\|v_m\|_2^2 \le \sum_{\alpha\in\mathbb{N}^n_{q/2}} p_\alpha^2/(q/2)! = \frac{1}{(q/2)!},
$$

so $\|v_m\|_2 \leq \frac{2^{O(q)}}{q^{q/4}}$. Since A is SoS-symmetric, has the minimum eigenvalue at least $-2^{O(q)} \cdot \frac{n^{q/4}}{q^{q/4}}$, and has nonzero entries only on the rows and columns $(i_1, \ldots, i_{q/2})$ with all different entries,

$$v^T \mathsf{A} v = (v_m)^T \mathsf{A}(v_m) \geq 2^{-O(q)} \cdot \frac{n^{q/4}}{q^{3q/4}}.$$

We finally compute $v^T \mathsf{W} v$. Let $v_w \in [n]^{q/2}$ be the vector where for each $\alpha \in \mathbb{N}^n_{q/2}$, we choose one $I \in [n]^{q/2}$ arbitrarily and set $(v_w)_I = p_\alpha$ (all other $(v_w)_I$'s are 0). By SoS-symmetry of W,

$$v^T \mathsf{W} v = (v_w)^T \mathsf{W}(v_w) = p^T \widehat{\mathsf{W}} p \geq \frac{1}{2},$$

by Theorem 9.7.1. Therefore,

$$v^T \cdot \mathsf{M} \cdot v \;=\; \frac{1}{c_1} \cdot v^T \cdot \left(\frac{1}{c_2} \cdot \frac{q^{3q/4}}{n^{3q/4}} \mathsf{A} + \frac{\mathsf{W}}{n^{q/2}} \right) \cdot v \;\geq\; \frac{1}{c_1} \cdot \left(\frac{1}{c_2} \cdot 2^{-O(q)} \cdot \frac{n^{q/4}}{q^{3q/4}} \cdot \frac{q^{3q/4}}{n^{3q/4}} + \frac{1}{2} \cdot \frac{1}{n^{q/2}} \right) \;\geq\; 0,$$

by taking $c_2 = 2^{\Theta(q)}$. So M is positive semidefinite, and this finishes the proof of the lower bound in Theorem 9.1.1

Thus we obtain,

Theorem 9.7.7 (Lower bound in Theorem 9.1.1)**.** *For even $q \leq n$, let $\mathcal{A} \in \mathbb{R}^{[n]^q}$ be a q-tensor with i.i.d. ± 1 entries. Then with probability $1 - n^{\Omega(1)}$,*

$$\frac{SoS_q(\mathcal{A}(x))}{\mathcal{A}_{\max}} \;\geq\; \left(\frac{\Omega(n)}{q} \right)^{q/4 - 1/2}.$$

As a side note, observe that by applying Lemma 9.7.6 and the proof of Lemma 9.7.5 to the SoS-symmetric matrix representation of $f(x) = \mathcal{A}(x)$ (instead of A), we obtain a stronger SoS upper bound (by polylog factors) for the special case of $d = q$:

Theorem 9.7.8 (Upper bound in Theorem 9.1.1)**.** *For even $q \leq n$, let $\mathcal{A} \in \mathbb{R}^{[n]^q}$ be a q-tensor with i.i.d. ± 1 entries. Then with probability $1 - n^{\Omega(1)}$,*

$$\frac{SoS_q(\mathcal{A}(x))}{\mathcal{A}_{\max}} \;\leq\; \left(\frac{O(n)}{q} \right)^{q/4 - 1/2}.$$

9.8 Upper bounds for Odd Degree Tensors

In the interest of clarity, in this section we shall prove Theorem 9.1.2 for the special case of 3-tensors. The proof readily generalizes to the case of all odd degree-d tensors.

9.8.1 Analysis Overview

Let $\mathcal{A} \in \mathbb{R}^{[n]^3}$ be a 3-tensor with i.i.d. uniform ± 1 entries. Assume $q/4$ is a power of 2 as this only changes our claims by constants. For $\ell \in [n]$ let $\overline{T}_\ell$ be an $n \times n$ matrix with i.i.d. uniform ± 1 entries, such that we have

$$f(x) := \langle \mathcal{A}, x^{\otimes 3} \rangle = \sum_{\ell \in [n]} x_\ell \, (x^T \, \overline{T}_\ell \, x) = \sum_{\ell \in [n]} x_\ell \, (x^T \, T_\ell \, x).$$

Let $T_\ell := (\overline{T}_\ell + \overline{T}_\ell^T)/2$. Following Hopkins et. al. [HSS15], let $\mathcal{T} := \sum_{\ell=1}^n T_\ell \otimes T_\ell$. Let $E \in \mathbb{R}^{[n]^2 \times [n]^2}$ be the matrix such that $E[(i,i),(j,j)] = \mathcal{T}[(i,i),(j,j)]$ for any $i,j \in [n]$ and $E[(i,j),(k,l)] = 0$ otherwise. Let $E' \in \mathbb{R}^{[n]^2 \times [n]^2}$ be the matrix such that $E'[(i,j),(i,j)] = E[(i,i),(j,j)] + E[(j,j),(i,i)]$ for any $i,j \in [n]$ and $E'[(i,j),(k,l)] = 0$ otherwise.

Let $T := \mathcal{T} - E \in \mathbb{R}^{[n]^2 \times [n]^2}$ and $\mathrm{A} := T^{\otimes q/4}$. Let $g(x) := (x^{\otimes 2})^T \, \mathcal{T} \, x^{\otimes 2}$ and $h(x) := (x^{\otimes 2})^T \, E \, x^{\otimes 2} = (x^{\otimes 2})^T \, E' \, x^{\otimes 2}$. Let $\widetilde{\mathbf{E}}$ be the pseudo-expectation operator returned by the program above.

We would like to show that there is some matrix representation B of A, such that w.h.p. $\max_{\|y\|=1} y^T \mathrm{B} y$ is small. To this end, consider the following mass shift procedure that we apply to A to get B:

$$\begin{aligned} \forall I, J \in [n]^{q/2}, \quad \mathrm{B}[I,J] &:= \frac{1}{(q/2)!^2} \sum_{\pi,\sigma \in \mathbb{S}_{q/2}} \mathrm{A}[\pi(I), \sigma(J)] \\ &= \frac{1}{|\mathcal{O}(I)| \; |\mathcal{O}(J)|} \sum_{I' \in \mathcal{O}(I), J' \in \mathcal{O}(J)} \mathrm{A}[I', J'] \end{aligned}$$

Below the fold we shall show that $\|\mathrm{B}\|_2^{4/q} = \widetilde{O}(n^{3/2}/\sqrt{q})$ w.h.p. This is sufficient to obtain the desired result since we have

$$\begin{aligned} & \|\mathrm{B}\|_2 \mathbf{I} - B \succeq 0 \\ \Rightarrow\; & \|\mathrm{B}\|_2 \|x\|^q - \langle x^{\otimes q/2}, \mathrm{B}\, x^{\otimes q/2} \rangle \succeq 0 \\ \Rightarrow\; & \|\mathrm{B}\|_2 \|x\|^q - \langle x^{\otimes 2}, T\, x^{\otimes 2} \rangle^{q/4} \succeq 0 \\ \Rightarrow\; & \|\mathrm{B}\|_2 \|x\|^q - (g(x) - h(x))^{q/4} \succeq 0 \\ \Rightarrow\; & \widetilde{\mathbf{E}} \left[(g(x) - h(x))^{q/4} \right] \leq \|\mathrm{B}\|_2 \\ \Rightarrow\; & \widetilde{\mathbf{E}} \left[g(x) - h(x) \right] \leq \|\mathrm{B}\|_2^{4/q} && \text{(Pseudo-Cauchy-Schwarz)} \\ \Rightarrow\; & \widetilde{\mathbf{E}} \left[g(x) \right] \leq \|\mathrm{B}\|_2^{4/q} + \widetilde{\mathbf{E}} \left[h(x) \right] \\ \Rightarrow\; & \widetilde{\mathbf{E}} \left[g(x) \right] \leq \|\mathrm{B}\|_2^{4/q} + 5n && (5n\, \mathbf{I} - E' \succeq 0) \\ \Rightarrow\; & \widetilde{\mathbf{E}} \left[g(x) \right] = \widetilde{O}(n^{3/2}/\sqrt{q}). \end{aligned}$$

Now $$\widetilde{\mathbf{E}} \left[f(x) \right] = \widetilde{\mathbf{E}} \left[\sum_{\ell \in [n]} x_\ell \, (x^T \, T_\ell \, x) \right] \qquad \text{(Following [HSS15])}$$

$$
\begin{aligned}
&\le \widetilde{\mathbf{E}}\left[\|x\|^2\right]^{1/2}\ \widetilde{\mathbf{E}}\left[\sum_{\ell\in[n]}(x^T T_\ell\, x)^2\right]^{1/2} && \text{(Pseudo-Cauchy-Schwarz)}\\
&\le \widetilde{\mathbf{E}}\left[\sum_{\ell\in[n]}(x^T T_\ell\, x)^2\right]^{1/2} && \text{(Pseudo-Cauchy-Schwarz)}\\
&\le \widetilde{\mathbf{E}}\left[\langle x^{\otimes 2}, \mathcal{T}\, x^{\otimes 2}\rangle\right]^{1/2}\\
&= \widetilde{\mathbf{E}}\left[g(x)\right]^{1/2} = \widetilde{O}(n^{3/4}/q^{1/4})
\end{aligned}
$$

9.8.2 Bounding $\|B\|_2^{4/q}$

For any $I^1, \ldots, I^p \in [n]^{q/2}$ let e_k denote the k-th smallest element in $\bigcup_{\ell, j}\{I_j^\ell\}$ and let

$$
\#\left(I^1, \ldots, I^p\right) := \left|\bigcup_{\ell\in[p]}\bigcup_{j\in[n]}\left\{I_j^\ell\right\}\right|.
$$

For any $c^1, \ldots, c^s \in [q/2]^p$, let

$$
\begin{aligned}
&\mathcal{C}(c^1 \ldots c^s) :=\\
&\quad\left\{(I^1, \ldots, I^p)\ \middle|\ \#\left(I^1, \ldots, I^p\right) = s,\ \forall k \in [s], \ell \in [p],\ e_k \text{ appears } c_\ell^k \text{ times in } I^\ell\right\}
\end{aligned}
$$

Observation 9.8.1.

$$
\left|\left\{(c^1, \ldots, c^s) \in ([q/2]^p)^s\ \middle|\ \mathcal{C}(c^1, \ldots, c^s) \neq \varphi\right\}\right| \le 2^{O(pq)} p^{\,pq/2}
$$

$$
\textit{if } \mathcal{C}(c^1, \ldots, c^s) \neq \varphi \textit{ then } \left|\mathcal{C}(c^1, \ldots, c^s)\right| \le \frac{n^s}{s!} \times \frac{((q/2)!)^p}{\prod_{\ell\in[p]} c_\ell^1! \ldots c_\ell^s!}
$$

Lemma 9.8.2. *Consider any $c^1, \ldots, c^s \in [q/2]^p$ and $(I^1, \ldots, I^p) \in \mathcal{C}(c^1, \ldots, c^s)$. We have*

$$
\mathbb{E}\left[\mathrm{B}\left[I^1, I^2\right]\mathrm{B}\left[I^2, I^3\right]\ldots\mathrm{B}\left[I^p, I^1\right]\right] \le 2^{O(pq)} n^{pq/8}\ \frac{p^{\,5pq/8}}{q^{\,3pq/8}}\ \prod_{\ell\in[p]} c_\ell^1! \ldots c_\ell^s!
$$

Proof. Consider any $c^1, \ldots, c^s \in [q/2]^p$ and $(I^1, \ldots, I^p) \in \mathcal{C}(c^1, \ldots, c^s)$. We have

$$
\begin{aligned}
&\mathbb{E}\left[\mathrm{B}\left[I^1, I^2\right]\mathrm{B}\left[I^2, I^3\right]\ldots\mathrm{B}\left[I^p, I^1\right]\right]\\
&= \frac{\prod_\ell c_\ell^1!^2 \ldots c_\ell^s!^2}{((q/2)!)^{2p}} \sum_{J^1, K^1\in\mathcal{O}\left(I^1\right), \ldots, J^p, K^p\in\mathcal{O}(I^p)} \mathbb{E}\left[\mathrm{A}\left[J^1, K^2\right]\mathrm{A}\left[J^2, K^3\right]\ldots\mathrm{A}\left[J^p, K^1\right]\right]
\end{aligned}
$$

$$= \frac{\prod_\ell c_\ell^1!^2 \dots c_\ell^s!^2}{((q/2)!)^{2p}} \cdot$$

$$\sum_{\forall \ell,\, J^\ell, K^\ell \in \mathcal{O}(I^\ell)} \mathbb{E}\left[\prod_{\ell\in[p]} \mathrm{T}\Big[J_1^\ell J_2^\ell \,,\, K_1^{\ell+1} K_2^{\ell+1}\Big] \mathrm{T}\Big[J_3^\ell J_4^\ell \,,\, K_3^{\ell+1} K_4^{\ell+1}\Big] \dots \mathrm{T}\Big[J_{q/2-1}^\ell J_{q/2}^\ell \,,\, K_{q/2-1}^{\ell+1} K_{q/2}^{\ell+1}\Big]\right]$$

$$= \frac{\prod_\ell c_\ell^1!^2 \dots c_\ell^s!^2}{((q/2)!)^{2p}} \sum_{\forall \ell,\, J^\ell, K^\ell \in \mathcal{O}(I^\ell)} \mathbb{E}\left[\prod_{\ell\in[p]} \prod_{g\in[q/4]} \sum_{h\in[n]} T_h\left[J_{2g-1}^\ell, K_{2g-1}^{\ell+1}\right] T_h\left[J_{2g}^\ell, K_{2g}^{\ell+1}\right]\right]$$

$$= \frac{\prod_\ell c_\ell^1!^2 \dots c_\ell^s!^2}{((q/2)!)^{2p}} \sum_{\forall \ell,\, J^\ell, K^\ell \in \mathcal{O}(I^\ell)} \sum_{\forall \ell, g,\, h(\ell,g)\in[n]} \mathbb{E}\left[\prod_{\ell\in[p]} \prod_{g\in[q/4]} T_{h(\ell,g)}\left[J_{2g-1}^\ell, K_{2g-1}^{\ell+1}\right] T_{h(\ell,g)}\left[J_{2g}^\ell, K_{2g}^{\ell+1}\right]\right]$$

$$= \frac{\prod_\ell c_\ell^1!^2 \dots c_\ell^s!^2}{((q/2)!)^{2p}} \sum_{\forall \ell,\, J^\ell, K^\ell \in \mathcal{O}(I^\ell)} \sum_{\biguplus_{u\in[n]} S_u = [p]\times[q/4]} \mathbb{E}\left[\prod_{r\in[n]} \prod_{(\ell,g)\in S_r} T_r\left[J_{2g-1}^\ell, K_{2g-1}^{\ell+1}\right] T_r\left[J_{2g}^\ell, K_{2g}^{\ell+1}\right]\right]$$

$$= \frac{\prod_\ell c_\ell^1!^2 \dots c_\ell^s!^2}{((q/2)!)^{2p}} \left|\mathcal{S}(I^1, \dots, I^p)\right| \qquad \text{where} \tag{9.4}$$

$\mathcal{S}(I^1, \dots, I^p) :=$

$$\Big\{\big(\bigoplus_{\ell\in[p]} (J^\ell, K^\ell)\,,\, (S_1 \dots S_n)\big) \;\Big|\; \bigoplus_{\ell\in[p]} (J^\ell, K^\ell) \in \prod_{\ell\in[p]} \mathcal{O}^2(I^\ell),\ (J_{2g-1}^\ell, K_{2g-1}^{\ell+1}) \neq (J_{2g}^\ell, K_{2g}^{\ell+1}),$$

$$\biguplus_{u\in[n]} S_u = [p]\times[q/4],\ \forall r\in[n],\ \mathcal{I}_{S_r}\Big(\bigoplus\nolimits_\ell (J^\ell, K^\ell)\Big) \text{ has only even multiplicity elements}\Big\},$$

$$\text{and } \mathcal{I}_{S_r}\Big(\bigoplus\nolimits_\ell (J^\ell, K^\ell)\Big) := \bigoplus_{(\ell,g)\in S_r} \Big(\{J_{2g-1}^\ell, K_{2g-1}^{\ell+1}\}\,,\, \{J_{2g}^\ell, K_{2g}^{\ell+1}\}\Big)$$

Thus it remains to estimate the size of $\mathcal{S}(I^1, \dots, I^p)$. We begin with some notation. For a tuple t and a subsequence t_1 of t, let $t \setminus t_1$ denote the subsequence of elements in t that are not in t_1. For a tuple of 2-sets $t = (\{a_1, b_1\}, \dots, \{a_m, b_m\})$, let $\mathbf{atomize}(t)$ denote the tuple $(a_1, b_1, \dots, a_m, b_m)$ (we assume $\forall i,\ a_i < b_i$). Observe that "atomize" is invertible.

For any $(\oplus_\ell (J^\ell, K^\ell), (S_1 \dots S_n)) \in \mathcal{S}(I^1, \dots, I^p)$, observe that $\mathcal{I}_{S_r}(\oplus_\ell (J^\ell, K^\ell))$ (which is of length $2|S_r|$) contains a subsequence I_{S_r} of length $|S_r|$, such that $\mathbf{multiset}(\mathrm{I}_{S_r}) = \mathbf{multiset}(\mathcal{I}_{S_r}(\oplus_\ell (J^\ell, K^\ell)) \setminus \mathrm{I}_{S_r})$. Now we know

$$\mathbf{multiset}\Big(\bigoplus\nolimits_r \mathbf{atomize}\Big(\mathcal{I}_{S_r}\Big(\bigoplus\nolimits_\ell (J^\ell, K^\ell)\Big)\Big)\Big) = \bigsqcup_{\ell\in[p]} \mathbf{multiset}\Big(J^\ell \oplus K^\ell\Big)$$

$$= \bigsqcup_{\ell\in[p]} \mathbf{multiset}\Big(I^\ell \oplus I^\ell\Big)$$

$$\Rightarrow \mathbf{multiset}(\bigoplus\nolimits_r \mathbf{atomize}(\mathrm{I}_{S_r})) = \mathbf{multiset}\Big(\bigoplus\nolimits_r \mathbf{atomize}\Big(\mathcal{I}_{S_r}\Big(\bigoplus\nolimits_\ell (J^\ell, K^\ell)\Big) \setminus \mathrm{I}_{S_r}\Big)\Big)$$

$$= \bigsqcup_{\ell\in[p]} \mathbf{multiset}\Big(I^\ell\Big).$$

$$\text{Thus, } \exists \pi \in \mathbb{S}_{pq/2},\ s.t.\ \bigoplus_r \mathbf{atomize}(\mathrm{I}_{S_r}) = \pi\Big(I^1 \oplus \cdots \oplus I^p\Big) \tag{9.5}$$

For tuples t, t', let $\mathbf{intrlv}(t, t')$ denote the set of all tuples obtained by interleaving the elements in t and t'. By Eq. (9.5), we obtain that for any $(\oplus_\ell(J^\ell, K^\ell), (S_1 \dots S_n)) \in \mathcal{S}(I^1, \dots, I^p)$,

$$\exists \pi \in \mathbb{S}_{pq/2},\ s.t.\ \forall r \in [n],\ \exists \sigma_r \in \mathbb{S}_{|S_r|},\ s.t.\quad \mathcal{I}_{S_r}\Big(\bigoplus_\ell (J^\ell, K^\ell)\Big) \in \mathbf{intrlv}(\mathrm{I}_{S_r}, \sigma_r(\mathrm{I}_{S_r})) \tag{9.6}$$
$$\text{where}\quad \mathrm{I}_{S_r} = \mathbf{atomize}^{-1}\Big(\pi\Big(I^1 \oplus \cdots \oplus I^p\Big)\Big)$$

For any $j \in [s]$, let $\bar{c}^j := \sum_{\ell \in [p]} c^j_\ell$. Now since $|\mathbf{intrlv}(t, t')| \le 2^{|t|+|t'|}$, by Eq. (9.6) we have that for any $\biguplus_{u \in [n]} S_u = [p] \times [q/4]$,

$$\begin{aligned}
\#(S_1, \dots, S_n) &:= \Big|\Big\{\bigoplus_\ell (J^\ell, K^\ell)\ \Big|\ (\bigoplus_\ell (J^\ell, K^\ell), (S_1 \dots S_n)) \in \mathcal{S}\Big(I^1, \dots, I^p\Big)\Big\}\Big| \\
&\le 2^{pq/4}\Big|\mathbb{S}_{|S_1|}\Big| \times \cdots \times \Big|\mathbb{S}_{|S_n|}\Big| \times \Big|\mathcal{O}\Big(I^1 \oplus \cdots \oplus I^p\Big)\Big| \\
&\le 2^{pq/4}\, |S_1|! \dots |S_n|!\, \frac{(pq/2)!}{\bar{c}^1! \dots \bar{c}^s!}
\end{aligned} \tag{9.7}$$

For any $(\oplus_\ell(J^\ell, K^\ell), (S_1 \dots S_n)) \in \mathcal{S}(I^1, \dots, I^p)$, observe that the even multiplicity condition combined with the condition that $(J^\ell_{2g-1}, K^{\ell+1}_{2g-1}) \ne (J^\ell_{2g}, K^{\ell+1}_{2g})$, imply that for each $r \in [n]$, $|S_r| \ne 1$. Thus every non-empty S_r has size at least 2, implying that the number of non-empty sets in $S_1, \dots, S_n$ is at most $pq/8$. Thus we have,

$$\begin{aligned}
&\Big|\mathcal{S}\Big(I^1, \dots, I^p\Big)\Big| \\
&= \sum_{U \subseteq [n], |U| \le pq/8}\ \sum_{\biguplus_{u \in U} S_u = [p] \times [q/4]} \#(S_1, \dots, S_n) && (S_r = \varphi \text{ if } r \notin U) \\
&= \sum_{U \subseteq [n], |U| \le pq/8}\ \sum_{\sum_{u \in U} s_u = pq/4}\ \sum_{|S_r| = s_r,\ \biguplus_{u \in U} S_u = [p] \times [q/4]} \#(S_1, \dots, S_n) && (s_r = 0, S_r = \varphi \text{ if } r \notin U) \\
&\le \sum_{U \subseteq [n], |U| \le pq/8}\ \sum_{\sum_{u \in U} s_u = pq/4} \#(S_1, \dots, S_n) \binom{pq/4}{s_1} \dots \binom{pq/4}{s_n} && (s_r = 0, S_r = \varphi \text{ if } r \notin U) \\
&\le \sum_{U \subseteq [n], |U| \le pq/8}\ \sum_{\sum_{u \in U} s_u = pq/4} 2^{pq/4}\, s_1! \dots s_n!\, \frac{(pq/2)!}{\bar{c}^1! \dots \bar{c}^s!} \binom{pq/4}{s_1} \dots \binom{pq/4}{s_n} && (\text{by Eq. (9.7)}) \\
&\le \sum_{U \subseteq [n], |U| \le pq/8}\ \sum_{\sum_{u \in U} s_u = pq/4} 2^{O(pq)}\, \frac{(pq/2)!}{\bar{c}^1! \dots \bar{c}^s!} (pq)^{s_1 + \dots + s_n} && (s_r = 0, S_r = \varphi \text{ if } r \notin U) \\
&\le \sum_{U \subseteq [n], |U| \le pq/8} 2^{O(pq + |U|)}\, \frac{(pq/2)!}{\bar{c}^1! \dots \bar{c}^s!} (pq)^{pq/4} \\
&\le \sum_{\bar{u} \in [pq/8]}\ \sum_{U \subseteq [n], |U| = \bar{u}} 2^{O(pq + |U|)}\, \frac{(pq/2)!}{\bar{c}^1! \dots \bar{c}^s!} (pq)^{pq/4}
\end{aligned}$$

$$
\begin{aligned}
&\leq \sum_{\overline{u} \in [pq/8]} 2^{O(pq)} \binom{n}{\overline{u}} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!} (pq)^{pq/4} \\
&\leq \sum_{\overline{u} \in [pq/8]} 2^{O(pq+\overline{u})} \frac{n^{\overline{u}}}{\overline{u}^{\overline{u}}} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!} (pq)^{pq/4} \\
&\leq \sum_{\overline{u} \in [pq/8]} 2^{O(pq)} (npq)^{pq/8} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!} \frac{(pq)^{pq/8}}{n^{pq/8-\overline{u}}\, \overline{u}^{\overline{u}}} \\
&\leq \sum_{\overline{u} \in [pq/8]} 2^{O(pq)} (npq)^{pq/8} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!} \frac{(pq)^{\overline{u}}}{\overline{u}^{\overline{u}}} && (\text{ since } pq < n) \\
&\leq \sum_{\overline{u} \in [pq/8]} 2^{O(pq)} (npq)^{pq/8} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!} \\
&\leq 2^{O(pq)} (npq)^{pq/8} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!}
\end{aligned}
$$

$$
\begin{aligned}
\Rightarrow\ &\mathbb{E}\left[\mathrm{B}\left[I^1, I^2\right]\mathrm{B}\left[I^2, I^3\right] \dots \mathrm{B}\left[I^p, I^1\right]\right] \\
&\leq 2^{pq/4} (npq)^{pq/8} \frac{(pq/2)!}{\overline{c}^1! \dots \overline{c}^s!} \frac{\prod_\ell c_\ell^1!^2 \dots c_\ell^s!^2}{((q/2)!)^{2p}} && (\text{by Eq. (9.4)}) \\
&\leq 2^{pq/4} (npq)^{pq/8} (pq/2)! \frac{\prod_\ell c_\ell^1! \dots c_\ell^s!}{((q/2)!)^{2p}} \\
&= 2^{O(pq)} n^{pq/8} \frac{p^{5pq/8}}{q^{3pq/8}} \prod_{\ell \in [p]} c_\ell^1! \dots c_\ell^s!
\end{aligned}
$$

∎

Lemma 9.8.3. *For all $i^1, \dots, i^p \in [n]^{q/2}$, we have*

$$
\begin{aligned}
&(1) \quad \mathbb{E}\left[\mathrm{B}\left[I^1, I^2\right]\mathrm{B}\left[I^2, I^3\right] \dots \mathrm{B}\left[I^p, I^1\right]\right] \geq 0 \\
&(2) \quad \mathbb{E}\left[\mathrm{B}\left[I^1, I^2\right]\mathrm{B}\left[I^2, I^3\right] \dots \mathrm{B}\left[I^p, I^1\right]\right] \neq 0 \quad \Rightarrow \quad \#\left(I^1, \dots, I^p\right) \leq \frac{pq}{4} + \frac{q}{2}
\end{aligned}
$$

Proof. The first claim follows immediately on noting that one is taking expectation of a polynomial of independent centered random variables with all coefficients positive.

For the second claim, note that $\mathbb{E}[\mathrm{B}[I^1, I^2]\mathrm{B}[I^2, I^3] \dots \mathrm{B}[I^p, I^1]] \neq 0$ implies that $\mathcal{S}(I^1, \dots, I^p) \neq \varphi$. Therefore there exists $\oplus_\ell(J^\ell, K^\ell)$ (where $J^\ell, K^\ell \in \mathcal{O}(I^\ell)$) and $\biguplus_{u \in [n]} S_u = [p] \times [q/4]$ such that every element in $\oplus_{r \in [n]} \mathcal{I}_{S_r}(\oplus_\ell(J^\ell, K^\ell))$ has even multiplicity. The rest of the proof follows from the same ideas as in the proof of Lemma 9.6.3. ∎

Lemma 9.8.4.

$$
\|\mathrm{B}\|_2^{4/q} \leq \frac{n^{3/2} \log^5 n}{\sqrt{q}} \quad w.h.p.
$$

Proof. We proceed by trace method. (Note that since T is symmetric, so are A and B).

$$
\begin{aligned}
&\mathbb{E}\left[\mathrm{Tr}(\mathrm{B}^p)\right] \\
&= \sum_{I^1,\dots,I^p \in [n]^{q/2}} \mathbb{E}\left[\mathrm{B}\left[I^1,I^2\right]\mathrm{B}\left[I^2,I^3\right]\dots\mathrm{B}\left[I^p,I^1\right]\right] \\
&= \sum_{s\in[pq/4+q/2]} \sum_{\#\left(i^1,\dots,i^p\right)=s} \mathbb{E}\left[\mathrm{B}\left[I^1,I^2\right]\mathrm{B}\left[I^2,I^3\right]\dots\mathrm{B}\left[I^p,I^1\right]\right] && \text{by Lemma 9.8.3} \\
&= \sum_{s\in[pq/4+q/2]} \sum_{c^1,\dots,c^s\in[q/2]^p} \sum_{(I^1,\dots,I^p)\in\mathcal{C}(c^1\dots c^s)} \mathbb{E}\left[\mathrm{B}\left[I^1,I^2\right]\mathrm{B}\left[I^2,I^3\right]\dots\mathrm{B}\left[I^p,I^1\right]\right] \\
&= \sum_{s\in[pq/4+q/2]} \sum_{c^1,\dots,c^s\in[q/2]^p} \sum_{(I^1,\dots,I^p)\in\mathcal{C}(c^1\dots c^s)} \mathbb{E}\left[\mathrm{B}\left[I^1,I^2\right]\mathrm{B}\left[I^2,I^3\right]\dots\mathrm{B}\left[I^p,I^1\right]\right] \\
&\le \sum_{s\in[pq/4+q/2]} \sum_{c^1,\dots,c^s\in[q/2]^p} \sum_{(I^1,\dots,I^p)\in\mathcal{C}(c^1\dots c^s)} 2^{O(pq)}\, n^{pq/8}\, \frac{p^{5pq/8}}{q^{3pq/8}} \prod_{\ell\in[p]} c^1_\ell!\dots c^s_\ell! && \text{by Eq. (9.4)} \\
&\le \sum_{s\in[pq/4+q/2]} 2^{O(pq)}\, \frac{n^{s+pq/8}}{s!}\, p^{9pq/8} q^{pq/8} && \text{by Observation 9.8.1} \\
&\le \sum_{s\in[pq/4+q/2]} 2^{O(pq)}\, \frac{n^{pq/4+q/2+pq/8}}{s!\, q^{pq/4+q/2-s}}\, p^{9pq/8} q^{pq/8} && (\text{since } q\le n) \\
&\le \sum_{s\in[pq/4+q/2]} 2^{O(pq)}\, \frac{n^{3pq/8+q/2}\, p^{9pq/8}}{q^{pq/8}} \;\le\; 2^{O(pq)}\, \frac{n^{3pq/8+q/2}\, p^{9pq/8}}{q^{pq/8}}.
\end{aligned}
$$

Choose p to be even and let $p = \Theta(\log n)$. Now

$$
\mathbf{Pr}\left[\|\mathrm{B}\|_2^{4/q} \ge n^{3/2}\log^5 n/\sqrt{q}\right] \le \mathbf{Pr}\left[\mathrm{Tr}(\mathrm{B}^p) \ge n^{\Omega(1)}\mathbb{E}\left[\mathrm{Tr}(\mathrm{B}^p)\right]\right].
$$

Applying Markov inequality completes the proof. ∎

Thus we obtain

Theorem 9.8.5. *Let $\mathcal{A} \in \mathbb{R}^{[n]^3}$ be a 3-tensor with i.i.d. ± 1 entries. Then for any even q such that $q \le n$, we have that with probability $1 - n^{\Omega(1)}$,*

$$
\frac{SoS_q(\mathcal{A}(x))}{\mathcal{A}_{\max}} \;\le\; \left(\frac{\widetilde{O}(n)}{q}\right)^{1/4}.
$$

Chapter 10

Future Directions and Perspectives on the Approximability Landscape

10.1 Operator Norms

In this section we discuss operator norms between general normed spaces and strive towards a characterization of their approximability. It will be convenient to have this discussion in the bilinear maximization language i.e., the language of injective tensor norms with two convex sets (Fact 3.4.1 shows how operator norms can be captured by bilinear form maximization). Recall that in Chapter 1 we used $\|\cdot\|_{C_1,C_2}$ to denote injective tensor norm. We will abuse this notation and use $\|\cdot\|_{X,Y}$ to denote $\|\cdot\|_{\mathrm{Ball}(X),\mathrm{Ball}(Y)}$.

We first formalize our goal. Let $(X_n)_{n\in\mathbb{N}}$ and $(Y_m)_{m\in\mathbb{N}}$ be sequences of norms over $\mathbb{R}^n$ and $\mathbb{R}^m$ respectively. For a pair of sequences (X_n, Y_m) we shall say $\|\cdot\|_{X_n,Y_m}$ is computationally approximable if there exists a family of poly-sized circuits approximating $\sup_{\|x\|_{X_n},\|y\|_{Y_m}\leq 1}\langle y\,,Ax\rangle$ within a constant independent of m and n (given an oracle for computing the norms). We are interested in the following questions:

Question. *For which pairs (X_n, Y_m) is the injective tensor norm computationally approximable? Can one connect the approximation factor to the geometry of X_n, Y_m?*

10.1.1 Approximability, Type and Cotype

Towards this characterization, evidence in the literature points to the powerful classification tools of Type and Cotype from functional analysis (see Section 3.5 for a quick introduction to these notions). Below we list various pieces of evidence that inform our conjectures for the approximability landscape:

- In the $\|\cdot\|_{\ell_p^n,\ell_q^m}$ case there are constant factor approximation algorithms whenever $p,q\geq 2$ (and whenever $p=1$ or $q=1$) and there is complexity theoretic evidence that no constant factor approximation algorithm exists in the other regimes. Note that when $p,q\geq 2$, $C_2(\ell_{p^*}^n), C_2(\ell_{q^*}^m)$ are bounded by absolute constants.

- Nesterov [Nes98] and Naor and Schechtman[1] independently observed that a natural computable convex programming relaxation has an integrality gap bounded by K_G when X, Y are 1-unconditional norms with 2-convexity constant $M^{(2)}(\cdot) = 1$ (a rich class of norms). For a 1-unconditional norm X, the Cotype-2 constant $C_2(X^*)$ is within a universal constant of $M^{(2)}(X)$ (see [LT13]). Thus if one extended the above result to the case when $M^{(2)}(X_n), M^{(2)}(Y_m)$ are constants independent of n, m, then one would automatically obtain such a result for the case when Cotype-2 constants of X_n^*, Y_m^* are independent of n, m.

- Naor, Regev and Vidick [NRV13] algorithmicized Haagerup's [Haa85] proof of the non-commutative Grothendieck inequality to give a computable convex programming relaxation with constant integrality gap for the case of X and Y being the Schatten-∞ norm. This is an important new example as the Schatten-∞ norm is not 1-unconditional (more generally it isn't a Banach lattice) but is also the dual of the Schatten-1 norm that has Cotype-2 bounded by an absolute constant independent of the dimension.

- In Section 3.8, Section 5.3 and Section 5.8 we see that factorization through Hilbert space is closely connected to approximation algorithms based on convex programming. Thus it is interesting to note that Pisier's remarkably general factorization theorem (see Section 5.8) applies whenever X^*, Y^* are Cotype-2.

Given the above, some natural questions are:

Question 1. *If $C_2(X_n^*), C_2(Y_m^*) = O(1)$ (i.e., independent of m, n), then is $\|\cdot\|_{X_n, Y_m}$ computationally approximable?*

Question 2. *Let X_n, Y_m be such that $T_p(X_n), T_p(Y_m) = O(1)$ for some $p > 1$ and either $C_2(X_n^*)$ grows polynomially in n or $C_2(Y_m^*)$ grows polynomially in m. Then can one rule out any constant factor approximation to $\|\cdot\|_{X_n, Y_m}$ (ideally assuming no poly-sized family of circuits captures NP)?*

Towards Question 2. As discussed in Section 2.1, combining quantitative finite dimensional versions of the MP+K theorem (see Theorem 13.12 in [MS09] and [AM85] for examples of such theorems) with hardness results of $p \to q$ operator norms yields quite general hardness results in the spirit of Question 2. Recall from Section 2.1.1, that the only regime of $p \to q$ norms for which we don't yet have satisfactorily strong hardness results is the case of $1 < p \leq 2 \leq q < \infty$ (and moreover it suffices to get hardness results for the $2 \to q$ case where $q > 2$). In all other regimes there is either a polynomial time constant factor approximation algorithm or there is complexity theoretic evidence that constant factor approximation requires time at least $2^{n^{\delta}}$.

Strong inapproximability (SoS gaps or NP-hardness) results for the hypercontractive $2 \to q$ case remain elusive and are closely related to obtaining hardness for polynomial maximization over the sphere as well as the quantum separability problem (relates to long standing open problems in quantum information theory). Consider the class of $2 \to X$

[1] personal communication

operator norms for exactly 2-convex norms X. This class of course contains all hypercontractive $2 \to q$ norms and moreover every operator norm in this class faces the same gadget reduction barrier discussed in Section 2.1.1. In Chapter 6 we show that this barrier can be overcome for certain exactly 2-convex norms (specifically mixed ℓ_p norms, i.e., $\ell_q(\ell_{q'})$ for $q, q' > 2$) and give a reduction from random label cover (for which polynomial level SoS gaps are available).

10.2 Degree ≥ 3

Our understanding of higher degree injective tensor norms is modest even in the special case of optimization over ℓ_p. In fact there are many fundamental problems surrounding optimization over ℓ_2.

10.2.1 Hardness over the Sphere and Dense CSPs

The problem of polynomial optimization has been studied over various compact sets [Las09, DK08], and it is natural to ask how well polynomial time algorithms can *approximate* the optimum value over a given compact set (see [DK08] for a survey). While the maximum of a degree-d polynomial over the simplex admits a PTAS for every fixed d [dKLP06], the problem of optimizing even a degree 3 polynomial over the hypercube does not admit any approximation better than $\exp((\log n)^{1-\varepsilon})$ (for arbitrary $\varepsilon > 0$) assuming NP cannot be solved in time $\exp((\log n)^{O(1)})$ [HV04]. The approximability of polynomial optimization on the sphere is poorly understood in comparison. Vertex based gadget reductions from CSPs face the same sparse optimal solution issue discussed in Section 2.1.1 (see also Chapter 4 for more details). However, considering the family of constraint based gadget reductions (i.e. replacing every arity-k constraint by a degree-k homogeneous gadget polynomial) exposes connections to an independently interesting question. For any degree-k polynomial a necessary condition for having all optimizers (over the sphere) being well-spread is to have $\gg n^{k/2}$ non-zero coefficients. This (and some further inspection of constraint based gadget reductions) prompts the following question:

Question 3. *Can one establish APX-hardness of dense arity-k CSPs ($\gg n^{k/2}$ constraints) over alphabet size q with a sparse predicate ($\ll q^{k/2}$ satisfying assignments)?*

10.2.2 Other Open Problems

There are still many intriguing questions to explore and the landscape is far from being completely understood:

- SoS gaps/APX-Hardness for Best Separable State and APX-Hardness of maximizing polynomials over the sphere.

- Are there subexponential ($2^{n^{\delta}}$) algorithms (via SoS) for operator norm problems (like $2\to 4$ norm, Grothendieck, Max-Cut) like those already known for unique games?

- Close the gap in the approximation achieved by q-levels of SoS for maximization over the sphere. The upper bound from Chapter 8 is $(n/q)^{d/2-1}$ and the best lower bound is roughly $(n/q^{O(1)})^{d/4-1/2}$ due to Hopkins et al. [HKP+17]. This is closely related to the problem of obtaining d-XOR hard instances with maximum gap in the advantage term.

- The PTAS over the simplex due to DeKlerk, Laurent and Parillo [dKLP06] requires d^d/ε^2 levels of SoS. It is interesting to check if this can be improved to $\mathrm{poly}(d)/\varepsilon^2$. Resolving this question negatively would also be interesting as there are reasons to believe it will lead to improved inapproximability results for optimization over the sphere.

Bibliography

[AAM$^+$11] Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein. Inapproximability of densest κ-subgraph from average case hardness. *Manuscript, available at http://www.cs.tau.ac.il/ nogaa/PDFS/dks8.pdf*, 2011. 82

[ABH$^+$05] Sanjeev Arora, Eli Berger, Elad Hazan, Guy Kindler, and Muli Safra. On non-approximability for quadratic programs. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 206–215. IEEE, 2005. 12

[ABI86] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of algorithms*, 7(4):567–583, 1986. 39

[AC$^+$17] Antonio Auffinger, Wei-Kuo Chen, et al. Parisi formula for the ground state energy in the mixed p-spin model. *The Annals of Probability*, 45(6B):4617–4631, 2017. 3

[AK06] Fernando Albiac and Nigel J Kalton. *Topics in Banach space theory, volume 233 of Graduate Texts in Mathematics*. Springer, New York, 2006. 43, 50

[Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 298–307. IEEE, 2003. 82

[AM85] D Amir and VD Milman. A quantitative finite-dimensional krivine theorem. *Israel Journal of Mathematics*, 50(1-2):1–12, 1985. 170

[AMMN06] Noga Alon, Konstantin Makarychev, Yury Makarychev, and Assaf Naor. Quadratic forms on graphs. *Inventiones mathematicae*, 163(3):499–522, 2006. Conference version in STOC '05. 12

[AN04] Noga Alon and Assaf Naor. Approximating the cut-norm via Grothendieck's inequality. In *Proceedings of the 36th annual ACM symposium on Theory of computing*, pages 72–80. ACM, 2004. 49, 54

[BBH$^+$12] Boaz Barak, Fernando GSL Brandao, Aram W Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 307–326. ACM, 2012. 3, 9, 10, 15, 16

[BCG+12] Aditya Bhaskara, Moses Charikar, Venkatesan Guruswami, Aravindan Vijayaraghavan, and Yuan Zhou. Polynomial integrality gaps for strong sdp relaxations of densest k-subgraph. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 388–405. SIAM, 2012. 83

[BdOFV14] Jop Briët, Fernando Mário de Oliveira Filho, and Frank Vallentin. Grothendieck inequalities for semidefinite programs with rank constraint. *Theory Of Computing*, 10(4):77–105, 2014. 64

[BGG+17] Vijay Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani. Weak decoupling, polynomial folds and approximate optimization over the sphere. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 1008–1019. IEEE, 2017. 7

[BGG+18a] V. Bhattiprolu, M. Ghosh, V. Guruswami, E. Lee, and M. Tulsiani. Inapproximability of Matrix $p \to q$ Norms. *ArXiv e-prints*, February 2018. 7

[BGG+18b] Vijay Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani. Approximating operator norms via generalized krivine rounding. *CoRR*, abs/1804.03644, 2018. 7

[BGK+01] Andreas Brieden, Peter Gritzmann, Ravindran Kannan, Victor Klee, László Lovász, and Miklós Simonovits. Deterministic and randomized polynomial-time approximation of radii. *Mathematika*, 48(1-2):63–105, 2001. 94, 142, 143

[BGL16] Vijay Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Certifying random polynomials over the unit sphere via sum of squares hierarchy. *arXiv preprint arXiv:1605.00903*, 2016. 7, 95, 146

[BH13] Fernando GSL Brandao and Aram W Harrow. Quantum de finetti theorems under local measurements with applications. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 861–870. ACM, 2013. 16

[BH15] Fernando GSL Brandao and Aram W Harrow. Estimating operator norms using covering nets. *arXiv preprint arXiv:1509.05065*, 2015. 9, 95

[BHK+16] Boaz Barak, Samuel B Hopkins, Jonathan Kelner, Pravesh K Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. In *Proceedings of the 57th annual IEEE symposium on Foundations of Computer Science*, FOCS '16, pages 428–437, 2016. 94, 105, 137

[Bis11] Punyashloka Biswal. Hypercontractivity and its applications. *arXiv preprint arXiv:1101.2913*, 2011. 9

[BKS12] Boaz Barak, Guy Kindler, and David Steurer. On the optimality of relaxations for average-case and generalized constraint satisfaction problems. *Manuscript, available from http://www.boazbarak.org/research.html*, 2012. 82

[BKS14] Boaz Barak, Jonathan A Kelner, and David Steurer. Rounding sum-of-squares relaxations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 31–40. ACM, 2014. 15, 16, 86, 88, 94, 95

[BKS15] Boaz Barak, Jonathan A Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 143–151. ACM, 2015. 4, 15

[BKS17] Boaz Barak, Pravesh Kothari, and David Steurer. Quantum entanglement, sum of squares, and the log rank conjecture. *arXiv preprint arXiv:1701.06321*, 2017. 16

[BM15] Gérard Biau and David M Mason. High-dimensional p-norms. In *Mathematical Statistics and Limit Theorems*, pages 21–40. Springer, 2015. 37

[BMMN13] Mark Braverman, Konstantin Makarychev, Yury Makarychev, and Assaf Naor. The Grothendieck constant is strictly smaller than Krivine's bound. In *Forum of Mathematics, Pi*, volume 1. Cambridge University Press, 2013. Conference version in FOCS '11. 13

[BRS15] Jop Briët, Oded Regev, and Rishi Saket. Tight hardness of the non-commutative Grothendieck problem. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1108–1122. IEEE, 2015. 12, 15, 28, 33, 34, 35

[BS14] Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. *arXiv preprint arXiv:1404.5236*, 2014. 87

[BT87] Jean Bourgain and Lior Tzafriri. Invertibility of large submatrices with applications to the geometry of banach spaces and harmonic analysis. *Israel journal of mathematics*, 57(2):137–224, 1987. 50

[BV09] S Charles Brubaker and Santosh S Vempala. Random tensors and planted cliques. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 406–419. Springer, 2009. 15

[BV11] Aditya Bhaskara and Aravindan Vijayaraghavan. Approximating matrix p-norms. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 497–511. SIAM, 2011. 11, 27, 29, 40, 41

[Cha12] Sourav Chatterjee. The missing log in large deviations for triangle counts. *Random Structures & Algorithms*, 40(4):437–451, 2012. 104

[DFJP74] Wayne J Davis, Tadeusz Figiel, William B Johnson, and Aleksander Pełczyński. Factoring weakly compact operators. *Journal of Functional Analysis*, 17(3):311–327, 1974. 25

[DK08] Etienne De Klerk. The complexity of optimizing over a simplex, hypercube or sphere: a short survey. *Central European Journal of Operations Research*, 16(2):111–125, 2008. 4, 93, 171

[DK12a] Bobby DeMarco and Jeff Kahn. Upper tails for triangles. *Random Structures & Algorithms*, 40(4):452–459, 2012. 104

[DK12b] Robert DeMarco and Jeff Kahn. Tight upper tail bounds for cliques. *Random Structures & Algorithms*, 41(4):469–487, 2012. 104

[dKL19] Etienne de Klerk and Monique Laurent. Convergence analysis of a lasserre hierarchy of upper bounds for polynomial minimization on the sphere. *arXiv preprint arXiv:1904.08828*, 2019. 15, 16, 94

[dKLP06] Etienne de Klerk, Monique Laurent, and Pablo A Parrilo. A PTAS for the minimization of polynomials of fixed degree over the simplex. *Theoretical Computer Science*, 361(2):210–225, 2006. 99, 171, 172

[dKLS14] Etienne de Klerk, Monique Laurent, and Zhao Sun. Convergence analysis for Lasserre's measure–based hierarchy of upper bounds for polynomial optimization. *arXiv preprint arXiv:1411.6867*, 2014. 15, 92

[dKLS15] Etienne de Klerk, Monique Laurent, and Zhao Sun. An alternative proof of a PTAS for fixed-degree polynomial optimization over the simplex. *Mathematical Programming*, 151(2):433–457, 2015. 99

[DlPG12] Victor De la Pena and Evarist Giné. *Decoupling: from dependence to independence*. Springer Science & Business Media, 2012. 92

[DM15] Yash Deshpande and Andrea Montanari. Improved sum-of-squares lower bounds for hidden clique and hidden submatrix problems. In *COLT*, pages 523–562, 2015. 137

[DW12] Andrew C Doherty and Stephanie Wehner. Convergence of sdp hierarchies for polynomial optimization on the hypersphere. *arXiv preprint arXiv:1210.5048*, 2012. 6, 15, 16, 92, 94

[Fay04] Leonid Faybusovich. Global optimization of homogeneous polynomials on the simplex and on the sphere. In *Frontiers in global optimization*, pages 109–121. Springer, 2004. 15, 94

[Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 534–543. ACM, 2002. 82

[FLM77] Tadeusz Figiel, Joram Lindenstrauss, and Vitali D Milman. The dimension of almost spherical sections of convex bodies. *Acta Mathematica*, 139(1):53–94, 1977. 22

[FV+62] David G Feingold, Richard S Varga, et al. Block diagonally dominant matrices and generalizations of the gerschgorin circle theorem. *Pacific J. Math*, 12(4):1241–1250, 1962. 117

[GLS12] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012. 5

[GM15] Rong Ge and Tengyu Ma. Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, page 829, 2015. 4, 15

[Gro53] Alexander Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. São Paulo* ***8***, 1953. 12, 25, 50

[Gro14] Leonard Gross. Hypercontractivity, logarithmic sobolev inequalities, and applications: a survey of surveys. *Diffusion, quantum theory, and radically elementary mathematics*, pages 45–74, 2014. 3

[GRSW16] Venkatesan Guruswami, Prasad Raghavendra, Rishi Saket, and Yi Wu. Bypassing UGC from some optimal geometric inapproximability results. *ACM Transactions on Algorithms (TALG)*, 12(1):6, 2016. Conference version in SODA '12. 5, 12, 28, 31, 33, 45

[GW95] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995. 3

[Haa81] Uffe Haagerup. The best constants in the khintchine inequality. *Studia Mathematica*, 70(3):231–283, 1981. 39, 49, 64, 68

[Haa85] Uffe Haagerup. The grothendieck inequality for bilinear forms on C*-algebras. *Advances in Mathematics*, 56(2):93–116, 1985. 3, 170

[HKP+17] Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 720–731. IEEE, 2017. 17, 93, 105, 147, 172

[HLZ10] Simai He, Zhening Li, and Shuzhong Zhang. Approximation algorithms for homogeneous polynomial optimization with quadratic constraints. *Mathematical Programming*, 125(2):353–383, 2010. 15, 16, 94, 98, 108, 144

[HM13] Aram W Harrow and Ashley Montanaro. Testing product states, quantum Merlin-Arthur games and tensor optimization. *Journal of the ACM (JACM)*, 60(1):3, 2013. 9, 29, 40

[HNW16] Aram W Harrow, Anand Natarajan, and Xiaodi Wu. Limitations of semidefinite programs for separable states and entangled games. *arXiv preprint arXiv:1612.09306*, 2016. 9, 16

[HSS15] Samuel B Hopkins, Jonathan Shi, and David Steurer. Tensor principal component analysis via sum-of-square proofs. In *Proceedings of The 28th Conference on Learning Theory*, pages 956–1006, 2015. 4, 15, 146, 163

[HV04] Johan Håstad and Srinivasan Venkatesh. On the advantage over a random assignment. *Random Structures & Algorithms*, 25(2):117–149, 2004. 171

[Ind06] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM (JACM)*, 53(3):307–323, 2006. 43

[JOR04] Svante Janson, Krzysztof Oleszkiewicz, and Andrzej Ruciński. Upper tails for subgraph counts in random graphs. *Israel Journal of Mathematics*, 142(1):61–92, 2004. 104

[Kho02] Subhash Khot. Hardness results for coloring 3-colorable 3-uniform hypergraphs. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 23–32. IEEE, 2002. 31

[KMOW17] Pravesh K Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. Sum of squares lower bounds for refuting any csp. *arXiv preprint arXiv:1701.04521*, 2017. 105, 146

[KMTV11] Amit Kumar, Rajsekar Manokaran, Madhur Tulsiani, and Nisheeth K Vishnoi. On lp-based approximability for strict csps. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 1560–1573. SIAM, 2011. 5

[KN08] Subhash Khot and Assaf Naor. Linear equations modulo 2 and the l_1 diameter of convex bodies. *SIAM Journal on Computing*, 38(4):1448–1463, 2008. 16, 94, 98, 142, 143

[KN09] Subhash Khot and Assaf Naor. Approximate kernel clustering. *Mathematika*, 55(1-2):129–165, 2009. 12

[KN11] Subhash Khot and Assaf Naor. Grothendieck-type inequalities in combinatorial optimization. *arXiv preprint arXiv:1108.2464*, 2011. 2

[KN12] Subhash Khot and Assaf Naor. Grothendieck-type inequalities in combinatorial optimization. *Communications on Pure and Applied Mathematics*, 65(7):992–1035, 2012. 8

[KNS10] Guy Kindler, Assaf Naor, and Gideon Schechtman. The UGC hardness threshold of the Lp Grothendieck problem. *Mathematics of Operations Research*, 35(2):267–283, 2010. Conference version in SODA '08. 12, 45, 46

[KO09] Subhash Khot and Ryan O'Donnell. SDP gaps and UGC-hardness for Max-Cut-Gain. *Theory OF Computing*, 5:83–117, 2009. 12

[Kri77] Jean-Louis Krivine. Sur la constante de Grothendieck. *CR Acad. Sci. Paris Ser. AB*, 284(8):A445–A446, 1977. 12, 13, 49

[KV04] Jeong Han Kim and Van H Vu. Divide and conquer martingales and the number of triangles in a random graph. *Random Structures & Algorithms*, 24(2):166–174, 2004. 104

[Kwa72a] Stanislaw Kwapień. Isomorphic characterizations of inner product spaces by orthogonal series with vector valued coefficients. *Stud. Math.*, 44:583–595, 1972. 22, 25, 50

[Kwa72b] Stanislaw Kwapień. On operators factorizable through l_p space. *Mémoires de la Société Mathématique de France*, 31:215–225, 1972. 25

[Las09] Jean Bernard Lasserre. *Moments, positive polynomials and their applications*, volume 1. World Scientific, 2009. 15, 94, 171

[Lau09] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009. 15, 87, 92, 94

[Loz03] Daniel W. Lozier. NIST digital library of mathematical functions. *Annals of Mathematics and Artificial Intelligence*, 38(1):105–119, May 2003. 60, 61, 69

[LP68] Joram Lindenstrauss and Aleksander Pełczyński. Absolutely summing operators in $\mathcal{L}_p$-spaces and their applications. *Studia Mathematica*, 29(3):275–326, 1968. 25

[LS91] László Lovász and Alexander Schrijver. Cones of matrices and set-functions and 0–1 optimization. *SIAM journal on optimization*, 1(2):166–190, 1991. 5

[LS$^+$09] Troy Lee, Adi Shraibman, et al. Lower bounds in communication complexity. *Foundations and Trends® in Theoretical Computer Science*, 3(4):263–399, 2009. 2

[LT13] Joram Lindenstrauss and Lior Tzafriri. *Classical Banach spaces II: function spaces*, volume 97. Springer Science & Business Media, 2013. 170

[LZ16] Eyal Lubetzky and Yufei Zhao. On the variational problem for upper tails in sparse random graphs. *Random Structures & Algorithms*, 2016. 104

[Mau74] Bernard Maurey. Théorèmes de factorisation pour les opérateurs à valeurs dans un espace l_p. *Séminaire Analyse fonctionnelle (dit)*, pages 1–5, 1974. 25, 50

[MR14] Andrea Montanari and Emile Richard. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014. 4, 15, 146

[MS09] Vitali D Milman and Gideon Schechtman. *Asymptotic theory of finite dimensional normed spaces: Isoperimetric inequalities in riemannian manifolds*, volume 1200. Springer, 2009. 170

[Nes98] Yurii Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optimization methods and software*, 9(1-3):141–160, 1998. 12, 13, 24, 76, 170

[NR14] Assaf Naor and Oded Regev. Krivine schemes are optimal. *Proceedings of the American Mathematical Society*, 142(12):4315–4320, 2014. 13, 49, 64

[NRV13] Assaf Naor, Oded Regev, and Thomas Vidick. Efficient rounding for the noncommutative grothendieck inequality. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 71–80. ACM, 2013. 3, 170

[NWY00] Yuri Nesterov, Henry Wolkowicz, and Yinyu Ye. Semidefinite programming relaxations of nonconvex quadratic optimization. In *Handbook of semidefinite programming*, pages 361–419. Springer, 2000. 24

[OZ13] Ryan O'Donnell and Yuan Zhou. Approximability and proof complexity. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1537–1556. Society for Industrial and Applied Mathematics, 2013. 89

[Pie67] Albrecht Pietsch. Absolut p-summierende abbildungen in normierten räumen. *Studia Mathematica*, 28(3):333–353, 1967. 25

[Pis78] Gilles Pisier. Grothendieck's theorem for noncommutative C*-algebras, with an appendix on grothendieck's constants. *Journal of Functional Analysis*, 29(3):397–415, 1978. 3

[Pis80] Gilles Pisier. Un theoreme sur les operateurs lineaires entre espaces de banach qui se factorisent par un espace de hilbert. In *Annales scientifiques de lEcole Normale Superieure*, volume 13, pages 23–43. Elsevier, 1980. 25, 50

[Pis86] Gilles Pisier. *Factorization of linear operators and geometry of Banach spaces*. Number 60. American Mathematical Soc., 1986. 2, 25, 50

[Pis12] Gilles Pisier. Grothendieck's theorem, past and present. *Bulletin of the American Mathematical Society*, 49(2):237–323, 2012. 2, 8

[Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 245–254. ACM, 2008. 5

[Ree91] JA Reeds. A new lower bound on the real Grothendieck constant. Manuscript, 1991. 12

[Ros73] Haskell P Rosenthal. On subspaces of lp. *Annals of Mathematics*, pages 344–373, 1973. 25

[RRS16] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random csps below the spectral threshold. *arXiv preprint arXiv:1605.00058*, 2016. i, 6, 15, 17, 95, 96, 137, 146, 149

[RS09] Prasad Raghavendra and David Steurer. Towards computing the Grothendieck constant. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 525–534. Society for Industrial and Applied Mathematics, 2009. 5, 12, 13

[RS10] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd annual ACM Symposium on Theory of Computing*, pages 755–764. ACM, 2010. 9

[RV15] Oded Regev and Thomas Vidick. Quantum xor games. *ACM Transactions on Computation Theory (ToCT)*, 7(4):15, 2015. 3

[SA90] Hanif D Sherali and Warren P Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990. 5

[Sch87] Gideon Schechtman. More on embedding subspaces of L_p in ℓ_r^n. *Compositio Math*, 61(2):159–169, 1987. 29, 41, 42

[Sch06] Gideon Schechtman. Two observations regarding embedding subsets of euclidean spaces in normed spaces. *Advances in Mathematics*, 200(1):125–135, 2006. 37

[Sim98] Barry Simon. The classical moment problem as a self-adjoint finite difference operator. *Advances in Mathematics*, 137(1):82–203, 1998. 157

[So11] Anthony Man-Cho So. Deterministic approximation algorithms for sphere constrained homogeneous polynomial optimization problems. *Mathematical programming*, 129(2):357–382, 2011. 16, 94, 98, 142, 144

[Ste61] Sergei Borisovich Stechkin. On best lacunary systems of functions. *Izvestiya Rossiiskoi Akademii Nauk. Seriya Matematicheskaya*, 25(3):357–366, 1961. 39

[Ste05] Daureen Steinberg. Computation of matrix norms with applications to robust optimization. *Research thesis, Technion-Israel University of Technology*, 2005. 9, 11, 12, 14, 15

[Tro09] Joel A Tropp. Column subset selection, matrix factorization, and eigenvalue optimization. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 978–986. Society for Industrial and Applied Mathematics, 2009. 50, 75

[TS14] R. Tomioka and T. Suzuki. Spectral norm of random tensors. *ArXiv e-prints*, July 2014. 145, 149

[vB72] Paul van Beek. An application of Fourier methods to the problem of sharpening the Berry-Esseen inequality. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 23(3):187–196, 1972. 46

[Ver17] Roman Vershynin. *High Dimensional Probability*. 2017. 29, 37

[Vu01] Van H Vu. A large deviation result on the number of small subgraphs of a random graph. *Combinatorics, Probability & Computing*, 10(1):79–94, 2001. 104