

ON THE COMPLEXITY OF
UNIQUE GAMES AND GRAPH EXPANSION

DAVID STEURER

A DISSERTATION

PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF
COMPUTER SCIENCE

ADVISOR: SANJEEV ARORA

NOVEMBER 2010

© Copyright by David Steurer, 2010. All rights reserved.

Abstract

Understanding the complexity of approximating basic optimization problems is one of the grand challenges of theoretical computer science. In recent years, a sequence of works established that Khot's Unique Games Conjecture, if true, would settle the approximability of many of these problems, making this conjecture a central open question of the field.

The results of this thesis shed new light on the plausibility of the Unique Games Conjecture, which asserts that a certain optimization problem, called `UNIQUE GAMES`, is hard to approximate in a specific regime.

On the one hand, we give the first confirmation of this assertion for a restricted model of computation that captures the best known approximation algorithms. The results of this thesis also demonstrate an intimate connection between the Unique Games Conjecture and approximability of graph expansion. In particular, we show that the Unique Games Conjecture is true if the expansion of small sets in graphs is hard to approximate in a certain regime. This result gives the first sufficient condition for the truth of the conjecture based on the inapproximability of a natural combinatorial problem.

On the other hand, we develop efficient approximation algorithms for certain classes of `UNIQUE GAMES` instances, demonstrating that several previously proposed variants of the Unique Games Conjecture are false. Finally, we develop a subexponential-time algorithm for `UNIQUE GAMES`, showing that this problem is significantly easier to approximate than NP-hard problems like `MAX 3-SAT`, `MAX 3-LIN`, and `LABEL COVER`, which are unlikely to have subexponential-time algorithm achieving a non-trivial approximation guarantee. This algorithm also shows that the inapproximability results based on the Unique Games Conjecture do not rule out subexponential-time algorithms, opening the possibility for such algorithms for many basic optimization problems like `MAX CUT` and `VERTEX COVER`.

At the heart of our subexponential-time algorithm for `UNIQUE GAMES` lies a novel algorithm for approximating the expansion of graphs across different scales, which might have applications beyond `UNIQUE GAMES`, especially in the design of divide-and-conquer algorithms.

Acknowledgments

I am greatly indebted to my advisor Sanjeev Arora, especially for his constant encouragement, his willingness to spend lots of time to discuss about research with me, and his efforts to improve my writing and presenting. His taste for problems, his commitment to research, and his scientific curiosity and creativity were a great source of inspiration. I couldn't have wished for a better advisor.

I am most thankful also to Boaz Barak, Prasad Raghavendra, and Nisheeth Vishnoi for their guidance, encouragement, and collaboration. They had considerable influence on the content of this thesis. I greatly enjoyed the many discussions with them about new problems and ideas.

Many thanks to Sanjeev Arora, Boaz Barak, Moses Charikar, Bernard Chazelle, and Russell Impagliazzo for serving on my thesis committee.

Over the years, I had the great pleasure to collaborate with many terrific researchers. I owe many thanks to my coauthors Sanjeev Arora, Boaz Barak, Markus Bläser, Constantinos Daskalakis, Benjamin Doerr, Venkatesan Guruswami, Moritz Hardt, Ishay Haviv, Thomas Holenstein, Russell Impagliazzo, Subhash Khot, Alexandra Kolla, Johannes Lengler, Yury Makarychev, William Matthews, Prasad Raghavendra, Anup Rao, Oded Regev, Peter Sanders, Prasad Tetali, Madhur Tulsiani, Nisheeth Vishnoi, Avi Wigderson, and Yuan Zhou. I greatly enjoyed working with them and I have learned a lot from them.

I thank all members of the theory groups in Princeton (at the University and at the Institute for Advanced Study) for creating such a great research environment. Thanks also to the members of the Center for Computational Intractability. The lively center meetings every month and the frequent workshops were great opportunities to learn and discuss about new results.

Many thanks to the Algo group at the Laboratoire de Recherche en Informatique (LRI) for hosting me during the winter of 2008.

Many thanks also to Microsoft Research New England for hosting me as an intern in the summer of 2009. It was there that Prasad Raghavendra and I started working on graph expansion.

Before coming to Princeton, I spent three terrific years at Saarland University and the Max-Planck-Institute Informatik. Many thanks to all theorists there, especially my advisors Peter Sanders (now at Karlsruhe Institute of Technology) and Kurt Mehlhorn.

Special thanks to the competent and helpful administrative staff at the Computer Science Department of Princeton University, especially Mitra Kelly and Melissa Lawson.

Heartfelt thanks to all my friends. They made the time in Princeton very enjoyable.

Finally, I would like to thank my family — my parents Florence and Manfred and my brother Hendrik — for their love and support.

I gratefully acknowledge the National Science Foundation for their financial support.

Contents

1. Introduction	1
1.1. Approximation Complexity	2
1.2. Unique Games Conjecture	4
1.3. Graph Expansion	8
1.4. Small-Set Expansion Hypothesis	10
1.5. Organization of this Thesis	12
2. Background	16
2.1. Optimization and Approximation	16
2.2. Graphs, Expansion, and Eigenvalues	18
2.2.1. Proof of Local Cheeger Bound	21
2.3. Unique Games and Semidefinite Relaxation	23
2.4. Unique Games Conjecture and Reductions	27
I. Algorithms	34
3. Unique Games with Expanding Constraint Graphs	35
3.1. Main Result	35
3.2. Propagation Rounding	39
3.3. Tensoring Trick	40
3.4. Notes	45
4. Parallel Repeated Unique Games	47
4.1. Main Result	47
4.2. Rounding Nonnegative Vectors	49
4.3. From Arbitrary to Nonnegative Vectors	53
4.3.1. Constructing Nearly Feasible Solutions	55
4.3.2. Repairing Nearly Feasible Solutions	57
4.4. Notes	62

Contents

5. Subexponential Approximation Algorithms	64
5.1. Main Results	64
5.1.1. Small Set Expansion	66
5.1.2. Unique Games	67
5.2. Subspace Enumeration	68
5.2.1. Small-Set Expansion	70
5.2.2. Unique Games	72
5.3. Threshold Rank vs. Small-Set Expansion	74
5.4. Low Threshold Rank Decomposition	76
5.5. Putting things together	77
5.6. Notes	78
II. Reductions	81
6. Graph Expansion and the Unique Games Conjecture	82
6.1. Main Results	83
6.1.1. Proof of Theorem 6.2	84
6.2. From SMALL-SET EXPANSION to PARTIAL UNIQUE GAMES	84
6.2.1. Completeness	85
6.2.2. Soundness	87
6.3. From PARTIAL UNIQUE GAMES to UNIQUE GAMES	94
6.3.1. Completeness	96
6.3.2. Soundness	97
6.4. From UNIQUE GAMES to SMALL-SET EXPANSION	99
6.4.1. Completeness	100
6.4.2. Soundness	101
6.5. Notes	105
7. Reductions between Expansion Problems	106
7.1. Main Results	106
7.2. Reduction	108
7.2.1. Completeness	110
7.2.2. Soundness	114
7.3. Notes	117

Contents

III. Lower Bounds	118
8. Limits of Semidefinite Relaxations	119
8.1. Overview	119
8.1.1. Techniques	122
8.1.2. Organization	123
8.2. Proof Sketch	123
8.3. Hierarchies of SDP Relaxations	128
8.3.1. LH_r -Relaxation	128
8.3.2. SA_r -Relaxation	129
8.3.3. Robustness	130
8.4. Weak Gaps for UNIQUE GAMES	130
8.4.1. Properties of Weak Gap Instances	131
8.5. Integrality Gap Instance for UNIQUE GAMES	134
8.6. Integral Vectors	134
8.7. Local Distributions for UNIQUE GAMES	140
8.8. Construction of SDP Solutions for E2LIN(q)	142
8.8.1. Vector Solution	143
8.8.2. Local Distributions	144
8.9. Smoothing	150
8.10. Notes	156
9. Open Problems	158
Bibliography	160

1. Introduction

A central goal of theoretical computer science is to understand the complexity of basic computational tasks, that is, the minimum amount of computational resources (especially, time) needed to carry out such tasks.

We will focus on computational tasks arising from combinatorial optimization problems. Here, the goal is to optimize a given objective function over a finite set of solutions. Problems of this kind are ubiquitous in diverse fields, for example, computer networks, artificial intelligence, operation research, and the natural and social sciences. A simple, illustrating example is **MAX CUT**:

MAX CUT: Given a graph G , find a partition of the vertex set of G into two sets so as to maximize the number of edges of G crossing the partition.

With few exceptions, the complexity of finding optimal solutions to combinatorial optimization problems is understood. Either there is a (low-degree) polynomial-time algorithm that finds an optimal solution for the problem, or there is a (low-degree) polynomial-time reduction from 3-SAT to the problem that shows that finding an optimal solution is NP-hard (as hard as solving 3-SAT). Prominent examples of the first category are **MINIMUM SPANNING TREE** and **MAXIMUM MATCHING**. Problems of the first category typically reduce to some form of convex optimization. Unfortunately, most combinatorial optimization problems fall in the second category, ruling out polynomial-time algorithms — and even subexponential-time algorithms — under standard complexity assumptions. For example, finding an optimal solution for **MAX CUT** is NP-hard [Kar72] (and the reduction from 3-SAT takes linear time). Since 3-SAT is unlikely to have a $\exp(o(n))$ -time algorithm [IPZ01], this reduction also shows that finding an optimal solution for **MAX CUT** is likely to require exponential time — $\exp(\Omega(n))$ — in the worst case.

Since finding optimal solution is intractable for many problems, it is natural to allow for approximation, raising the general question:

What is the complexity of finding approximately optimal solutions?

1. Introduction

Understanding the approximability of optimization problems turns out to be challenging. Even after much progress in the last decades, the approximability of basic problems — in particular, MAX CUT — remains open.

1.1. Approximation Complexity

There are several ways to measure the quality of an approximation. In many cases, the “approximation ratio” is an appropriate measure. (Later, we will come across problems that require less coarse measures of approximation.) We say that an algorithm has *approximation ratio* α if it always computes a solution whose value is within a factor α of the optimal value. (For us, there is no difference between approximation ratios α and $1/\alpha$. Most of the discussion in this chapter will be about maximization problems. In this case, approximation ratios are typically taken to be less than 1.) We allow algorithms to be randomized. In this case, we compare the expected value of the computed solution to the optimal value. (The expectation is only over the randomness of the algorithm and not over a particular input distribution.) All randomized algorithms considered in this thesis can be derandomized without (much) loss in efficiency or approximation.

For MAX CUT, a trivial algorithm achieves approximation ratio $1/2$: A random partition cuts half of the edges in expectations, which is within a factor $1/2$ of the optimal value. For a long time, no better approximation algorithm for MAX CUT was known, raising the question:

Is there an efficient algorithm for MAX CUT with non-trivial approximation ratio?

In a beautiful and influential result, Goemans and Williamson [GW95] answered this question positively, giving an algorithm for MAX CUT with non-trivial approximation ratio $\alpha_{\text{GW}} \approx 0.878$. This work introduced semidefinite programming (SDP) relaxations as a general technique for designing approximation algorithms. Today, this technique is the basis of the best known approximation algorithms for many basic optimization problems. (Often, no other techniques are known to achieve the same approximation as SDP relaxations.)

The next question regarding the approximability of MAX CUT is whether the approximation ratio α_{GW} is best possible or if it can be improved. In fact, until the early 1990s, it was not ruled out that for every constant $\alpha < 1$, MAX CUT has an efficient algorithm with approximation ratio α . (This question is

1. Introduction

not idle because such approximation schemes are known for other NP-hard optimization problems.)

Is it NP-hard to obtain an approximation ratio 0.9999 for MAX CUT?

A breakthrough result, known as PCP Theorem¹ [ALM⁺98, AS98], answered this question positively (possibly for a constant closer to 1 than 0.9999). For a large class of problems, this result showed that the approximation ratios of efficient algorithms are bounded away from 1 (assuming P is different from NP).

These initial hardness of approximation results (for MAX CUT say), based on the PCP Theorem, are not entirely satisfying. First, the hardness results leave open the possibility of efficient algorithms with significantly better approximation ratios than currently known. Second, the initial reductions showing these hardness results were quite inefficient (the blowup of the reduction was a high-degree polynomial, especially because of the use of Raz's Parallel Reduction Theorem [Raz98]).

Numerous works addressed the first issue. In an influential work [Hås01], Håstad showed optimal inapproximability results for several basic optimization problem, including MAX 3-LIN, a generalization of MAX CUT to 3-uniform hypergraph. Like for MAX CUT, the trivial randomized algorithm for MAX 3-LIN achieves approximation ratio $1/2$. Shockingly, this algorithm is optimal for MAX 3-LIN (unlike for MAX CUT). Håstad shows that no efficient algorithm for MAX 3-LIN can achieve a non-trivial approximation ratio (better than $1/2$) unless $P = NP$. For many other problems, Håstad obtains improved, but not necessarily tight hardness results. For example, his results rule out efficient algorithms for MAX CUT with approximation ratio less than $16/17 \approx 0.941$ (using a gadget of Trevisan et al. [TSSW00]). Closing the gap between known approximation algorithms and known hardness results for basic problems like MAX CUT remains one of the outstanding open problems of the field. As discussed in the next section, Khot's Unique Games Conjecture offers a unified way to resolve this question for many problem classes. The hardness results based on the Unique Games Conjecture can be seen as a continuation of Håstad's work.

The second issue (inefficiency of initial hardness reductions) led to the development of efficient PCP constructions, culminating in the work of Moshkovitz and Raz [MR08]. Their result implies that Håstad's reductions (from 3-SAT

¹Here, PCP stands for probabilistically checkable proofs.

1. Introduction

can be carried out in near-linear time, which shows that achieving a non-trivial approximation ratio for MAX 3-LIN has essentially the same complexity as solving 3-SAT (exactly). Concretely, if one assumes the Exponential Time Hypothesis [IPZ01], namely that the decision problem 3-SAT requires time $\exp(\Omega(n))$, then for every constant² $\varepsilon > 0$, achieving an approximation ratio $1/2 + \varepsilon$ for MAX 3-LIN requires time $\exp(n^{1-o(1)})$. One of the main results of this thesis is that reductions based on the Unique Games Conjecture are inherently inefficient, and therefore hardness results based on this conjecture cannot rule out subexponential approximation algorithms.

1.2. Unique Games Conjecture

UNIQUE GAMES³ is the following constraint satisfaction problem, generalizing MAX CUT:

UNIQUE GAMES: Given a variable set V , an alphabet size $k \in \mathbb{N}$, and a list of constraints of the form $x_i - x_j = c \pmod k$ with $i, j \in V$ and $c \in \{1, \dots, k\}$, find an assignment to the variables $(x_i)_{i \in V}$ so as to satisfy as many of the constraints as possible.

Khot [Kho02] conjectured that this problem is hard to approximate, in the sense that for $\varepsilon > 0$, it is NP-hard to distinguish between the case that at least $1 - \varepsilon$ of the constraints can be satisfied and the case that not more than ε of the constraints can be satisfied. (An additional condition of the conjecture is that the alphabet size is allowed to grow only mildly with the input size, say $k = o(\log n)$.)

Unique Games Conjecture: For every constant $\varepsilon > 0$, the following task is NP-hard: Given a UNIQUE GAMES instance with n variables and alphabet size $k = o(\log n)$, distinguish between the cases,

YES: some assignment satisfies at least $1 - \varepsilon$ of the constraints,

NO: no assignment satisfies at least ε of the constraints.

²One can choose ε even slightly subconstant.

³In fact, UNIQUE GAMES typically refers to a slightly more general problem. (See [Section 2.3](#) for the more general definition.) For the purpose of the current discussion, it suffices to consider the special case of UNIQUE GAMES described above.

1. Introduction

Starting with Khot’s work, a sequence of results showed that the truth of this conjecture would imply improved and often optimal hardness of approximation results for many basic problems, including MAX CUT. Perhaps most strikingly, Raghavendra [Rag08] showed that assuming the Unique Games Conjecture a simple SDP relaxation achieves an optimal approximation for every constraint satisfaction problem (this class of problems includes MAX CUT, MAX 3-LIN, and UNIQUE GAMES). We refer to the survey [Kho10] for a more complete account of the known consequences of the Unique Games Conjecture.

Independent of the truth of the Unique Games Conjecture, the above results demonstrate that for many problems, UNIQUE GAMES is a common barrier for improving current algorithms (in the sense that improving the current best algorithm for, say, MAX CUT requires giving an improved algorithm for UNIQUE GAMES). We call approximation problems of this kind UG-hard (as hard as approximating UNIQUE GAMES).

The great number of strong UG-hardness results makes the Unique Games Conjecture a fascinating open question. Towards resolving this conjecture, it makes sense to ask more generally:

What is the complexity of approximating UNIQUE GAMES?

On the one hand, any lower bound on the complexity of UNIQUE GAMES typically translates to corresponding lower bounds for all UG-hard problems. (This translation is essentially without quantitative loss, because UG-hardness reductions tend to be very efficient — the output of the reduction is only a constant factor larger than the input.) On the other hand, any upper bound on the complexity of UNIQUE GAMES opens the possibility of similar upper bounds for other UG-hard problems.

The works in this thesis address several aspects of the above question. First, we consider the following aspect:

What properties make UNIQUE GAMES instances easy?

The goal is to develop efficient algorithms for UNIQUE GAMES that provide good approximation guarantees for instances with certain properties. These properties should not be too restrictive so that more general constraint satisfaction problems (like LABEL COVER) remain hard even restricted to instances with these properties. Results of this kind give concrete explanations why the Unique Games Conjecture seems difficult to prove. Any approach for proving the conjecture has to avoid producing instances with these property. Another

1. Introduction

hope in pursuing this question is that algorithms developed for special classes of instances might be useful for general instances, leading to new upper bounds on the complexity of UNIQUE GAMES and potentially to a refutation of the Unique Games Conjecture. (This hope turns out to be justified; at the end of this section, we discuss a surprising algorithm for UNIQUE GAMES that uses algorithms developed for special classes of instances.)

In this thesis, we identify two properties that make UNIQUE GAMES instances easy. (The properties are not overly restrictive in the sense that the more general LABEL COVER⁴ problem remains hard when restricted to instances with these properties.)

In [Chapter 3](#), we show that instances with expanding constraint graphs are easy. The *constraint graph* of a UNIQUE GAMES instance is a graph on the variable set with an edge between any two variables that appear in the same constraint. Expansion measures how well a graph is connected. (In the next section, we discuss graph expansion in greater detail.) In particular, randomly generated graphs tend to be expanding. Therefore, this result demonstrates that UNIQUE GAMES is easy on certain, commonly studied input distributions.

In [Chapter 4](#), we show that instances of UNIQUE GAMES arising from parallel repetition are easy. Parallel repetition is a commonly used reduction for amplifying approximation hardness (e.g., for LABEL COVER). This result shows that parallel repetition is unlikely to be useful for proving the Unique Games Conjecture. (Previous work [[FKO07](#)] suggested an approach for proving the Unique Games Conjecture based on parallel repetition.)

A major shortcoming of the current knowledge about the Unique Games Conjecture is that only few consequences of a refutation of the conjecture are known. Concretely, the following scenario is not ruled out: there exists an efficient algorithm for UNIQUE GAMES and, at the same time, it is intractable to achieve better approximations for all other UG-hard problems. (In contrast, an efficient algorithm for 3-SAT implies efficient algorithms for all other NP-hard problems.) Hence, the goal is to identify “hard-looking” problems that could be solved efficiently if an efficient algorithm for UNIQUE GAMES existed, raising the question:

What “hard-looking” problems reduce to UNIQUE GAMES?

⁴LABEL COVER is a more general constraint satisfaction problem than UNIQUE GAMES, where the constraints are allowed to be projections instead of permutations. The analog of the Unique Games Conjecture for LABEL COVER is known to be true (e.g., [[MR08](#)]).

1. Introduction

In [Chapter 6](#), we show that a certain graph expansion problem reduces to `UNIQUE GAMES` and is therefore easier to approximate than `UNIQUE GAMES`. (Graph expansion problems are well-studied and have a wide range of applications. See the next section for more discussion.) Our reduction shows that any efficient algorithm for `UNIQUE GAMES` could be turned into an algorithm for this graph expansion problem. This algorithm would have a significantly better approximation guarantee than the current best algorithms for this problem. Before this work, no consequence of a refutation of the Unique Games Conjecture was known for a problem other than `UNIQUE GAMES` itself.

Showing consequences of a refutation of the Unique Games Conjecture gives some evidence for the truth of the conjecture. (Especially, if the consequences are surprising.) Another way to generate evidence for the truth of the conjecture is to prove lower bounds for `UNIQUE GAMES` in restricted models of computation, which leads us to the following question:

What is the complexity of approximating `UNIQUE GAMES` in restricted models of computation?

In the context of approximation problems, it is natural to study models of computations defined by hierarchies of relaxations (typically linear or semidefinite relaxations). Such hierarchies contain relaxations with gradually increasing complexity — from linear complexity to exponential complexity. Relaxations with higher complexity provide better approximations for the optimal value of optimization problems. (The relaxations with highest, i.e., exponential, complexity typically compute the optimal value exactly.)

Starting with the seminal work of Arora, Bollobás, and Lovász [[ABL02](#)], lower bounds on the complexity of many approximation problems were obtained in various hierarchies. We show the first super-polynomial lower bound for `UNIQUE GAMES` in a hierarchy that captures the best known algorithms for all constraint satisfaction problems (see [Chapter 8](#)).

Our lower bounds translate to corresponding lower bounds for most classes of `UG-hard` problems (via the known `UG-hardness` reductions). Even for specific `UG-hard` problem like `MAX CUT` such lower bounds were not known before.

Finally, we address the question about the complexity of approximating `UNIQUE GAMES` directly. In [Chapter 5](#), we give an algorithm for `UNIQUE GAMES` with subexponential running time. The algorithm achieves an approximation guarantee that the Unique Games Conjecture asserts to be `NP-hard` to

1. Introduction

achieve. Concretely, the algorithm distinguishes between the case that at least $1 - \epsilon$ of the constraints are satisfiable and the case that not more than ϵ of the constraints are satisfiable in time $\exp(k n^{\epsilon^{1/3}})$, where n is the number of variables and k is alphabet size. We conclude that UG-hardness results cannot rule out subexponential algorithms (unlike hardness results based on Håstad’s reductions). The algorithm also shows that any reduction from 3-SAT to UNIQUE GAMES proving the Unique Games Conjecture must be inefficient (i.e., blow up the instance size by a high-degree polynomial), assuming 3-SAT does not have subexponential algorithms. In this way, our algorithm rules out certain classes of reductions for proving the Unique Games Conjecture (in particular, reductions of the kind of Håstad’s reductions).

At the heart of our subexponential algorithm for UNIQUE GAMES lies a novel algorithm for approximating the expansion of graphs across different scales. Cheeger’s influential work [Che70], relating eigenvalues and expansion, started a great body of research on approximations for graph expansion. The results of this thesis show an intimate connection between the complexity of UNIQUE GAMES and the complexity of approximating graph expansion. We discuss more of this connection in next sections (especially §1.4).

1.3. Graph Expansion

Graphs are ubiquitous structures in computer science, mathematics, and the natural and social sciences. For example, they are useful for modeling various networks like the internet, genetic networks, and social networks.

A fundamental parameter of a graph is its *expansion*. The expansion⁵ of a vertex set S , denoted $\Phi(S)$, measures how well the vertices in S are connected to the rest of the graph. Formally, $\Phi(S)$ is defined as the number of edges leaving S normalized by the number of the edges incident to the set. (An edge with both endpoints in the set S counts twice in this normalization.) The *expansion* Φ of a graph is then defined as the minimum expansion of a vertex set with volume at most $1/2$. (In the context of expansion, the volume of a set is defined as the fraction of edges incident to the set.)

A natural computational question is to compute the expansion of a graph. We can formulate this question as optimization problem:

SPARSEST CUT: Given a graph, find a vertex set S with volume at most $1/2$ so

⁵In this thesis, expansion always refers to edge expansion. In the context of Markov chains, the quantity we call expansion is often called conductance.

1. Introduction

as to minimize its expansion $\Phi(S)$.

A related problem is BALANCED SEPARATOR, where the goal is to find a set with minimum expansion among all sets with volume between β and $1/2$ for some balance parameter β . We refer to optimization problem of this kind (where we are required to find a set with minimum expansion subject to volume constraints) loosely as *graph expansion problems*.

Unfortunately (but not surprisingly), it is NP-hard to solve graph expansion problems exactly, which motivates the following question:

How well can we approximate graph expansion problems?

In an influential work, Cheeger [Che70] showed a relation, known as Cheeger’s inequality, between the expansion of a graph and the second largest eigenvalue of its adjacency matrix. (Cheeger’s proof is for manifolds instead of graphs. Dodziuk [Dod84] and independently Alon–Milman [AM85] and Alon [Alo86] showed the corresponding inequality for graphs.) Cheeger’s inequality leads to an algorithm for SPARSEST CUT with the following approximation guarantee: Given a graph with expansion ε , the algorithm finds a set with volume at most $1/2$ and expansion at most $2\sqrt{\varepsilon}$. We say that the algorithm achieves an $(\varepsilon, 2\sqrt{\varepsilon})$ -approximation for every $\varepsilon > 0$.

Leighton and Rao [LR99] gave an algorithm for SPARSEST CUT with approximation ratio $O(\log n)$ based on a linear programming relaxation. Arora, Rao, and Vazirani (ARV) [ARV04] improved this approximation ratio to $O(\sqrt{\log n})$ using a semidefinite programming relaxation. Both of these approximations are incomparable to Cheeger’s inequality. For graphs with low expansion $\Phi \ll 1/\log n$, the ARV approximation guarantee is strongest. For expansion $\Phi \gg 1/\log n$, Cheeger’s inequality gives the best known approximation.

A major open question in combinatorial optimization is whether the approximations of ARV and Cheeger’s inequality are optimal in their respective regimes, or if better approximations for SPARSEST CUT are achievable. Even assuming the Unique Games Conjecture, no strong hardness result for SPARSEST CUT is known. (The best known result is that one cannot achieve approximation ratios arbitrarily close to 1 unless 3-SAT has subexponential algorithms [AMS07].) In this thesis, we introduce a hypothesis about the approximability of the expansion of small sets that allows us to show quantitatively tight hardness results for SPARSEST CUT and other expansion problems (in a certain regime).

1. Introduction

In the following discussion, we will focus⁶ on the case that the size of the graph is very large compared to its expansion (for example, $\Phi \gg 1/\log \log n$). In this case, approximation guarantees depending on the instance size are meaningless. In particular, Cheeger's inequality provides the only known non-trivial approximation for SPARSEST CUT in this regime.

1.4. Small-Set Expansion Hypothesis

Cheeger's inequality provides a good approximation for SPARSEST CUT in the sense that it can distinguish between the case that the expansion of a graph is very close to 0, say $\Phi \leq 0.001$, and the case that the expansion is bounded away from 0, say $\Phi \geq 0.1$. For the following natural generalization of SPARSEST CUT, no such approximation is known:

SMALL-SET EXPANSION: Given a graph and a scale δ , find a vertex set S with volume at most δ so as to minimize its expansion $\Phi(S)$.

Concretely, for every constant $\varepsilon > 0$, there exists $\delta > 0$ such that all known polynomial-time algorithms fail to distinguish between the following cases:

Small-set expansion close to 0: there exists a vertex set with volume δ and expansion at most ε ,

Small-set expansion close to 1: every vertex set with volume δ has expansion more than $1 - \varepsilon$.

The lack of efficient algorithms suggests that this problem may be computationally hard. To formalize this possibility, we introduce the following hypothesis:

Small-Set Expansion (SSE) Hypothesis: For every constant $\varepsilon > 0$, there exists a constant $\delta = \delta(\varepsilon) > 0$ such that the following task is NP-hard: Given a graph, distinguish between the cases,

YES: the graph contains a vertex set with volume δ and expansion at most ε ,

⁶ This restriction is of course not without loss of generality. We restrict ourselves to this regime because it turns out to be closely related to the Unique Games Conjecture and because we can prove tight bounds on the approximability of graph expansion problems in this regime. Understanding the approximability of SPARSEST CUT in the regime $\Phi \ll 1/\log n$ remains wide open.

1. Introduction

NO: the graph does not contain a vertex set with volume δ and expansion at most $1 - \varepsilon$.

In [Chapter 6](#), we show that this hypothesis implies the Unique Games Conjecture. (Combined with results of [Chapter 7](#), we can in fact show that this hypothesis is equivalent to a variant of the Unique Games Conjecture, namely [Hypothesis 6.3](#).) This result gives the first (and so far only) sufficient condition for the truth of the Unique Games Conjecture based on the inapproximability of a problem different than UNIQUE GAMES.

Given that the SSE hypothesis is a stronger hardness assumption than the Unique Games Conjecture, we can ask if it has further consequences (beyond the consequences of the Unique Games Conjecture).

What consequences would a confirmation of the SSE hypothesis have?

In [Chapter 7](#), we show that the SSE hypothesis implies quantitatively tight inapproximability results for many graph expansion problems, in particular SPARSEST CUT and BALANCED SEPARATOR. Concretely, our results imply that it is SSE-hard to beat Cheeger’s inequality and achieve an $(\varepsilon, o(\sqrt{\varepsilon}))$ -approximation for SPARSEST CUT. (We say a problem is SSE-hard if an efficient algorithm for the problem implies that the SSE hypothesis is false.)

There is a strong parallel between this result and the known consequences of the Unique Games Conjecture. The Unique Games Conjecture asserts a qualitative⁷ inapproximability for UNIQUE GAMES, which is generalization of MAX CUT. In turn, the conjecture implies tight quantitative inapproximability results for MAX CUT and similar problems. Similarly, the SSE hypothesis asserts a qualitative inapproximability for SMALL-SET EXPANSION, which generalizes SPARSEST CUT. In turn, the hypothesis implies tight quantitative inapproximability results for SPARSEST CUT and other graph expansion problems. The value of results of this kind is that they unify the question of improving known algorithms for a class of problems to a question about the qualitative approximability of a single problem.

The significant consequences of a confirmation of the SSE hypothesis make it an interesting open question to prove or refute the hypothesis. As for UNIQUE GAMES, it makes sense to ask the general question:

What is the complexity of approximating SMALL-SET EXPANSION?

⁷ The Unique Games Conjecture is a qualitative hardness assumption in the sense that it does not specify any precise absolute constants or any concrete functional dependencies.

1. Introduction

On the one hand, it turns out that the lower bounds for `UNIQUE GAMES` in [Chapter 8](#) extend to `SMALL-SET EXPANSION`, confirming the SSE hypothesis in a restricted model of computation (defined by a hierarchy of relaxations that captures the best known approximation algorithms). As for `UG-hard` problems, this lower bound translates to all `SSE-hard` problems. These results give the first quantitatively tight super-polynomial lower bounds for graph expansion problems in a hierarchy that captures the best known algorithms.

The results of [Chapter 6](#) and [Chapter 7](#) also show that the SSE hypothesis is *equivalent* to a variant of the Unique Games Conjecture. This variant asserts that `UNIQUE GAMES` remains hard to approximate even restricted to instances with constraint graphs satisfying a relatively mild, qualitative expansion property. (In particular, the techniques of [Chapter 3](#) — unique games with expanding constraints graphs — fail on instances with this mild expansion property.)

On the other hand, we show that `SMALL-SET EXPANSION` (like `UNIQUE GAMES`) admits approximation algorithms with subexponential running time (see [Chapter 5](#)). This result demonstrates that `SSE-hardness` results do not rule out subexponential algorithms. A concrete possibility is that `SPARSEST CUT` has a subexponential-time algorithm with constant approximation ratio. The subexponential-time algorithm for `SMALL-SET EXPANSION` inspired the subexponential algorithm for `UNIQUE GAMES` discussed previously (in [§1.2](#)).

1.5. Organization of this Thesis

In the following, we outline the structure of this thesis and describe the contributions and results of the individual chapters.

Part 1 — Algorithms

Chapter 3: Unique Games with Expanding Constraint Graphs. We study the approximability of `UNIQUE GAMES` in terms of expansion properties of the underlying constraint graph. If the optimal solution satisfies a $1 - \epsilon$ fraction of the constraints and the constraint graph has spectral gap⁸ λ , the algorithm developed in this chapter finds an assignment satisfying a $1 - O(\epsilon/\lambda)$ fraction of constraints. This result demonstrates that `UNIQUE GAMES` instances

⁸Cheeger’s inequality asserts that spectral gap of a graph is close to its expansion. Concretely, $\lambda/2 \leq \Phi \leq \sqrt{2\lambda}$. Sometimes λ is called *spectral expansion*.

1. Introduction

with spectral gap $\lambda \gg \varepsilon$ cannot be hard in the sense of the Unique Games Conjecture.

Our algorithm is based on a novel analysis of the standard SDP relaxation for UNIQUE GAMES. In contrast to the usual local analysis, our analysis takes into account global properties of the solution to the SDP relaxation.

This chapter is based on joint work with Arora, Khot, Kolla, Tulsiani, and Vishnoi [AKK⁺08].

Chapter 4: Parallel-Repeated Unique Games. We study the approximability of parallel-repeated instances of UNIQUE GAMES. Parallel repetition is commonly used for hardness amplification (e.g., for LABEL COVER, a generalization of UNIQUE GAMES). We give an improved approximation algorithm for parallel-repeated instances of UNIQUE GAMES. The guarantee of our algorithm matches the best known approximation for unrepeated instances [CMM06a]. In this sense, our result demonstrates that parallel repetition fails to amplify hardness (for UNIQUE GAMES).

Our algorithm is based on a novel rounding for the standard SDP relaxation of UNIQUE GAMES. In contrast to previous works, our rounding is able to exploit a tensor-product structure of the solution.

This chapter is based on joint work with Barak, Hardt, Haviv, Rao and Regev [BHH⁺08] and on the work [Ste10b].

Chapter 5: Subexponential Approximation Algorithms. We develop an algorithm for UNIQUE GAMES with subexponential running time. The Unique Games Conjecture asserts that the approximation provided by our algorithm is NP-hard to achieve. A consequence of our algorithm is that known UG-hardness results do not rule subexponential algorithms.

The main ingredient of our algorithm is a general decomposition for graphs. We show that any graph can be efficiently partitioned into induced subgraphs, each with at most n^β eigenvalues above $1 - \eta$. Furthermore, the partition respects all but $O(\eta/\beta^3)^{1/2}$ of the original edges.

This chapter is based on joint work with Arora and Barak [ABS10].

Part 2 — Reductions

Chapter 6: Graph Expansion and the Unique Games Conjecture. We propose a hypothesis (SSE hypothesis) about the approximability of the expansion of small sets in graphs. The reduction from SMALL-SET EXPANSION to UNIQUE

1. Introduction

GAMES developed in this chapter demonstrates that this hypothesis implies the Unique Games Conjecture. This result gives the first sufficient condition for the truth of the Unique Games Conjecture based on a problem different than UNIQUE GAMES.

Furthermore, we show a partial converse of this implication: An refutation of the SSE hypothesis implies that a stronger variant of the Unique Games Conjecture is false. This stronger variant asserts that UNIQUE GAMES remains hard to approximate (in the sense of the Unique Games Conjecture) even restricted to instances with constraint graphs satisfying a relatively mild, qualitative expansion property.

This chapter is based on joint work with Raghavendra [RS10].

Chapter 7: Reductions between Expansion Problems. Based on the reduction from SMALL-SET EXPANSION to UNIQUE GAMES in Chapter 6, we develop reductions from SMALL-SET EXPANSION to other graph expansion problems, in particular BALANCED SEPARATOR. Under the SSE hypothesis, these reductions imply quantitatively tight inapproximability results for these graph expansion problems.

Furthermore, we show that the SSE hypothesis implies a stronger variant of the Unique Games Conjecture, asserting that UNIQUE GAMES restricted to instances with constraint graphs that satisfy a certain expansion property.

Compared to similar reductions from UNIQUE GAMES, a key novelty of these reductions is that, on top of a composition with local gadgets, carefully placed random edges are added to the construction. The random edges are placed such that, on the one hand, the completeness of the reduction is not changed and, on the other hand, the expansion properties in the soundness case are improved.

This chapter is based on joint work with Raghavendra and Tulsiani [RST10b].

Part 3 — Lower Bounds

Chapter 8: Limits of Semidefinite Programming. We show super-polynomial lower bounds for UNIQUE GAMES in certain hierarchies of SDP relaxations (which capture the best known approximation algorithms for constraint satisfaction problems). Previous lower bounds considered only relaxations of fixed polynomial complexity or hierarchies that do not capture the best known approximation algorithms. Via known reductions, our lower

1. Introduction

bounds extend to all known UG-hard problems. (Our construction and analysis also yields super-polynomial lower bounds for SMALL-SET EXPANSION. Via the reductions in [Chapter 7](#), these lower bounds extend to other SSE-hard graph expansion problems, like BALANCED SEPARATOR.)

A simple, but important ingredient of our analysis is a robustness theorem for the hierarchies we consider. This robustness theorem asserts that an “approximate solution” to a relaxation in these hierarchies can be turned to a proper solution without changing the objective value by much.

This chapter is based on joint work with Raghavendra [[RS09a](#)].

2. Background

This chapter serves as a source of references for the following chapters. (However, its content is not assumed in the following.) Interested readers can skim this chapter for the first reading.

2.1. Optimization and Approximation

For simplicity, we restrict the discussion to maximization problems. For minimization problems, we use the same notations, but change signs in the definitions appropriately.

Optimization Problems. A *generic optimization instance* \mathfrak{I} consists of a set Ω of *feasible solutions* and an *objective function* $f: \Omega \rightarrow \mathbb{R}$. (The symbol \mathfrak{I} is a capital “i” in Fraktur font. It is intended to be pronounced as “i”.) If the meaning is clear from the context, we identify the instance \mathfrak{I} with its objective function and write $\mathfrak{I}(x)$ to denote the objective value of x . The (*optimal*) *value* $\text{opt}(\mathfrak{I})$ is defined as maximum objective value of a feasible solution,

$$\text{opt}(\mathfrak{I}) \stackrel{\text{def}}{=} \max_{x \in \Omega} \mathfrak{I}(x).$$

A solution x^* is *optimal* for instance \mathfrak{I} if it achieves the optimal value $\mathfrak{I}(x^*) = \text{opt}(\mathfrak{I})$.

An *optimization problem* Π is formally specified as a set of optimization instances. We say that an optimization problem Π is *combinatorial* if every instance $\mathfrak{I} \in \Pi$ has only a finite number of feasible solutions. For any optimization problem Π , we associate the following computational problem $\text{EXACT-}\Pi$,

EXACT- Π : Given an instance $\mathfrak{I} \in \Pi$, find an optimal solution x^* of \mathfrak{I} .

The computational problem $\text{EXACT-}\Pi$ makes only sense if we also fix an succinct encoding of the instances of Π . For example, the instances of MAX CUT are naturally encoded as graphs.

2. Background

Typically, we use the same notation for the collection of instances and for the computational problem. For example, MAX CUT stands both for the set of MAX CUT instances and for the computational problem of finding the maximum cut in a graph.

Approximation Algorithms. We say that an algorithm \mathcal{A} achieves a (c, s) -*approximation* for an optimization problem Π if given an instance $\mathfrak{I} \in \Pi$ with $\text{opt}(\mathfrak{I}) \geq c$, the algorithm finds a feasible solution $x = \mathcal{A}(\mathfrak{I})$ for the instance \mathfrak{I} with objective value $\mathfrak{I}(x) \geq s$.

To understand the approximation guarantees of an algorithm \mathcal{A} for a problem Π , we want to find for every value $c \in \mathbb{R}$, the largest value s such that the algorithm \mathcal{A} achieves a (c, s) -approximation for Π .

For example, the Goemans–Williamson algorithm achieves a $(1 - \varepsilon, 1 - O(\sqrt{\varepsilon}))$ -approximation for MAX CUT for all $\varepsilon > 0$.

If the objective functions of all instances of a problem Π are nonnegative, we define the *approximation ratio* α of an algorithm \mathcal{A} for Π as the minimum ratio between the value of the solution found by \mathcal{A} and the optimal value,

$$\alpha \stackrel{\text{def}}{=} \min_{\substack{\mathfrak{I} \in \Pi \\ x = \mathcal{A}(\mathfrak{I})}} \frac{\mathfrak{I}(x)}{\text{opt}(\mathfrak{I})}.$$

Gap-Promise Problems. Hardness of approximation results are typically based on gap-preserving reductions between optimization problems. In this context, it is convenient to associate the following kind of promise problems with optimization problems. For an optimization problem Π and constants $c, s \in \mathbb{R}$ with $c \geq s$, define the following promise problem,

(c, s) -GAP- Π : Given an instance $\mathfrak{I} \in \Pi$, distinguish between the cases

Yes $\text{opt}(\mathfrak{I}) \geq c$,

No $\text{opt}(\mathfrak{I}) < s$.

(In some situations, it makes sense to take c and s as functions of instance parameters, e.g., the size of the instance. In this thesis, we focus on the case that c and s are constants, depending only on the problem Π .)

2. Background

2.2. Graphs, Expansion, and Eigenvalues

In this work, we consider undirected, weighted graphs and we allow self-loops. In this case, we can represent a *graph* G with vertex set V as a symmetric distribution over pairs ij with $i, j \in V$. The *edges* of G are the pairs ij in the support of this distribution. (Here, symmetry means that for any two vertices $i, j \in V$, the pair ij has the same probability as the pair ji .) We will assume finite vertex sets.

We write $ij \sim G$ to denote a random edge ij sampled from G . For a vertex $i \in V$, we write $j \sim G(i)$ to denote a random *neighbor* of i in G . (A random neighbor of i is obtained by sampling a random edge of G conditioned on the event that the first endpoint of the edge is i and outputting the second endpoint of that edge.)

Expansion and Expansion Profile. For two vertex sets $S, T \subseteq V$, we define $G(S, T)$ as the fraction of edges going from S to T ,

$$G(S, T) \stackrel{\text{def}}{=} \mathbb{P}_{ij \sim G} \{i \in S, j \in T\}.$$

For a vertex set $S \subseteq V$, we define its (*edge*) *boundary* $\partial_G(S)$ as the fraction of edges leaving S (going from S to the complement of S),

$$\partial_G(S) \stackrel{\text{def}}{=} G(S, V \setminus S).$$

We define its *volume* $\mu_G(S)$ as the fraction of edges going out a vertex in S ,

$$\mu_G(S) \stackrel{\text{def}}{=} G(S, V).$$

The *expansion* $\Phi_G(S)$ is the ratio of these quantities

$$\Phi_G(S) \stackrel{\text{def}}{=} \frac{\partial_G(S)}{\mu_G(S)}.$$

(According to the above expression, the expansion of the empty set is undefined, since both the boundary and the volume is naught. Since we are interested in minimizing the expansion, it is convenient to define the expansion of the empty set to be infinite.)

2. Background

The quantities defined above have the following probabilistic interpretations,

$$\begin{aligned}\partial_G(S) &= \mathbb{P}_{ij \sim G} \{i \in S, j \notin S\}, \\ \mu_G(S) &= \mathbb{P}_{ij \sim G} \{i \in S\}, \\ \Phi_G(S) &= \mathbb{P}_{ij \sim G} \{j \notin S \mid i \in S\}.\end{aligned}$$

Another probabilistic interpretation of the expansion of the set S is the fraction of neighbors outside of S for a typical vertex in S ,

$$\Phi_G(S) = \mathbb{E}_{i \sim G} \left[\mathbb{P}_{j \sim G(i)} \{j \notin S\} \mid i \in S \right].$$

Here, $i \sim G$ denotes a random vertex of G , obtained by sampling a random edge of G and taking its first endpoint. If the graph G is clear from the context, we usually write $i \sim V$ or $i \sim \mu$ to denote a random vertex of G .

The *expansion* Φ_G of the graph G is the minimum expansion of a set with volume at most $1/2$,

$$\Phi_G \stackrel{\text{def}}{=} \min_{\substack{S \subseteq V \\ \mu_G(S) \leq 1/2}} \Phi_G(S).$$

More generally, for $\delta \in [0, 1]$, the *expansion at volume* δ , denoted $\Phi_G(\delta)$, is the minimum expansion of a set with volume at most δ ,

$$\Phi_G(\delta) \stackrel{\text{def}}{=} \min_{\substack{S \subseteq V \\ \mu_G(S) \leq \delta}} \Phi_G(S).$$

The curve $\delta \mapsto \Phi_G(\delta)$ is called the *expansion profile* of G . If the graph G is clear from the context, we typically omit the subscript G for the above notations.

Note that for every $\delta > 0$, the expansion profile satisfies $\Phi(\delta) \cdot \delta' = \Phi(1 - \delta) \cdot (1 - \delta')$ for some $\delta' \leq \delta$. In particular, $\Phi(1 - \delta) \leq \delta / (1 - \delta)$. Therefore, we typically take $\delta \in [0, 1/2]$.

Functions on Graphs, Laplacians, and Cheeger Bounds. For a graph G with vertex set V , we write $L_2(V)$ to denote the space $\{f : V \rightarrow \mathbb{R}\}$ equipped with natural inner product for $f, g \in L_2(V)$,

$$\langle f, g \rangle \stackrel{\text{def}}{=} \mathbb{E}_{i \sim V} f_i g_i.$$

2. Background

(We often write f_i to denote the value $f(i)$ of the function f on vertex i .) This inner product induces the norm $\|f\| := \langle f, f \rangle^{1/2}$. We will also be interested in the norms $\|f\|_\infty = \max_{i \in V} f_i$ and $\|f\|_1 = \mathbb{E}_{i \sim V} |f_i|$.

We identify the graph G with the following linear (Markov) operator on V ,

$$Gf(i) \stackrel{\text{def}}{=} \mathbb{E}_{j \sim G(i)} f(j).$$

The matrix corresponding to this operator is the (weighted) adjacency matrix of G normalized so that every row sums to 1. The operator G is self-adjoint with respect to the inner product on $L_2(V)$ (because $\langle f, Gg \rangle = \mathbb{E}_{ij \sim G} f_i g_j$). Since G is self-adjoint, its eigenvalues are real and its eigenfunctions are an orthogonal basis of $L_2(V)$.

For a vertex set $S \subseteq V$, let $\mathbb{1}_S \in L_2(V)$ be the $\{0, 1\}$ -indicator function of S . (For the all-ones functions $\mathbb{1}_V$, we typically drop the subscript.) Since $G\mathbb{1} = \mathbb{1}$, the all-ones function is an eigenfunction of G with eigenvalue 1. The *Laplacian* L_G of the graph G is the following linear operator on $L_2(V)$,

$$L_G \stackrel{\text{def}}{=} I - G,$$

where I is the identity operator on $L_2(V)$. The Laplacian L_G corresponds to the following quadratic form

$$\langle f, Lf \rangle = \mathbb{E}_{ij \sim G} \frac{1}{2} (f_i - f_j)^2.$$

(We drop the subscript for the Laplacian if the graph is clear from the context.) The following identities hold for all vertex sets $S, T \subseteq V$,

$$\begin{aligned} G(S, T) &= \langle \mathbb{1}_S, G\mathbb{1}_T \rangle, \\ \mu(S) &= \|\mathbb{1}_S\|^2, \quad \partial(S) = \langle \mathbb{1}_S, L\mathbb{1}_S \rangle, \\ \Phi(S) &= \frac{\langle \mathbb{1}_S, L\mathbb{1}_S \rangle}{\|\mathbb{1}_S\|^2}. \end{aligned}$$

Lemma 2.1 (Cheeger's Inequality, [Che70]). *Let $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of G . Suppose the gap between the largest and second largest eigenvalue is $\varepsilon = \lambda_1 - \lambda_2$ (the spectral gap of G). Then,*

$$\varepsilon/2 \leq \Phi_G \leq \sqrt{2\varepsilon}.$$

2. Background

The proof of Cheeger's inequality combined with a truncation argument yields the following lemma. For a self-contained proof, see the next subsection §2.2.1.

Lemma 2.2 (Local Cheeger Bound). *For every function $f \in L_2(V)$, there exists a level set $S \subseteq V$ of the function f^2 with volume $\mu(S) \leq \delta$ and expansion*

$$\Phi(S) \leq \frac{\sqrt{1 - \langle f, Gf \rangle^2 / \|f\|^4}}{1 - \|f\|_1^2 / \delta \|f\|^2} \leq \frac{\sqrt{2\langle f, Lf \rangle / \|f\|^2}}{1 - \|f\|_1^2 / \delta \|f\|^2}.$$

The above Cheeger bound is an important ingredient for the results of Chapter 5 (Subexponential Approximation Algorithms). Local variants of Cheeger's inequality (similar to the above bound) appear in several works, for example, [DI98, GMT06, RST10a].

Spectral Profile. The *spectral profile* of the graph G refers the curve $\delta \mapsto \Lambda_G(\delta)$, where

$$\Lambda_G(\delta) \stackrel{\text{def}}{=} \min_{\substack{f \in L_2(V) \\ \|f\|_1^2 \leq \delta \|f\|^2}} \frac{\langle f, Lf \rangle}{\|f\|^2}.$$

The spectral profile is a lower bound on the expansion profile, $\Lambda_G(\delta) \leq \Phi_G(\delta)$, because $f = \mathbb{1}_S$ satisfies $\|f\|_1^2 = \mu(S)\|f\|^2$ for every vertex set $S \subseteq V$. Lemma 2.2 implies an upper bound of the expansion profile in terms of spectral profile. For all $\delta, \gamma \in [0, 1)$, the following relations hold

$$\Lambda(\delta) \leq \Phi(\delta) \leq \sqrt{2\Lambda((1 - \gamma)\delta)} / \gamma.$$

2.2.1. Proof of Local Cheeger Bound

In this subsection, we prove the following local Cheeger bound.

Lemma (Restatement of Lemma 2.2). *For every function $f \in L_2(V)$, there exists a level set $S \subseteq V$ of the function f^2 with volume $\mu(S) \leq \delta$ and expansion*

$$\Phi(S) \leq \frac{\sqrt{1 - \langle f, Gf \rangle^2 / \|f\|^4}}{1 - \|f\|_1^2 / \delta \|f\|^2} \leq \frac{\sqrt{2\langle f, Lf \rangle / \|f\|^2}}{1 - \|f\|_1^2 / \delta \|f\|^2}.$$

Let $f \in L_2(V)$. Suppose $f^2 \leq 1$. Consider the following distribution over vertex subsets $S \subseteq V$:

2. Background

1. Sample $t \in [0, 1]$ uniformly at random.
2. Output the set $S = \{i \in V \mid f_i^2 > t\}$.

Note that every set S in the support of this distribution is a level set of the function f^2 . In the following lemmas, we establish simple properties of this distribution.

Claim 2.3. *The expected volume of S satisfies $\mathbb{E}_S \mu(S) = \|f\|^2$.*

Proof. We calculate the expected volume as follows

$$\mathbb{E}_S \mu(S) = \mathbb{E}_{i \sim \mu} \mathbb{P}_{t \in [0,1]} \{f_i^2 > t\} = \|f\|^2. \quad \square$$

Claim 2.4. *The second moment of $\mu(S)$ is at most $\mathbb{E}_S \mu(S)^2 \leq \|f\|_1^2$.*

Proof. We bound the expectation of $\mu(S)^2$ as follows

$$\mathbb{E}_S \mu(S)^2 = \mathbb{E}_{i,j \sim \mu} \mathbb{P}\{\min\{f_i^2, f_j^2\} > t\} = \mathbb{E}_{i,j \sim \mu} \min\{f_i^2, f_j^2\} \leq \mathbb{E}_{i,j \sim \mu} f_i f_j = \|f\|_1^2. \quad \square$$

Claim 2.5. *Sets with volume larger than δ contribute to the expected volume at most $\mathbb{E}_S \mu(S) \mathbb{1}_{\mu(S) > \delta} \leq \mathbb{E}_S \mu(S)^2 / \delta$.*

Proof. Immediate because $\mu(S) \mathbb{1}_{\mu(S) > \delta} \leq \mu(S)^2 / \delta$ holds pointwise. \square

Claim 2.6. *The expected boundary of S is bounded by*

$$\mathbb{E}_S G(S, V \setminus S) \leq \|f\|^2 \sqrt{1 - \langle f, Gf \rangle^2 / \|f\|^4}.$$

Proof. We calculate the expected boundary of S and apply the Cauchy–Schwarz inequality,

$$\begin{aligned} \mathbb{E}_S G(S, V \setminus S) &= \mathbb{E}_{ij \sim G} \mathbb{P}\{i \in S \wedge j \notin S\} = \mathbb{E}_{ij \sim G} \mathbb{P}\{f_i^2 > t \geq f_j^2\} \\ &= \mathbb{E}_{ij \sim G} \max\{f_i^2 - f_j^2, 0\} = \frac{1}{2} \mathbb{E}_{ij \sim G} |f_i^2 - f_j^2| = \frac{1}{2} \mathbb{E}_{ij \sim G} |f_i - f_j| \cdot |f_i + f_j| \\ &\leq \left(\mathbb{E}_{ij \sim G} \frac{1}{2} (f_i - f_j)^2 \cdot \mathbb{E}_{ij \sim G} \frac{1}{2} (f_i + f_j)^2 \right)^{1/2} \quad (\text{using Cauchy–Schwarz}) \\ &= \langle f, (I - G)f \rangle^{1/2} \langle f, (I + G)f \rangle^{1/2} = \sqrt{\|f\|^4 - \langle f, Gf \rangle^2}. \quad \square \end{aligned}$$

2. Background

We combine the previous claims to complete the proof of [Lemma 2.2](#). Let S^* be the level set of f^2 with volume at most δ and minimum expansion. Then,

$$\begin{aligned} \Phi(S^*) &\leq \frac{\mathbb{E}_S G(S, V \setminus S) \mathbb{1}_{\mu(S) \leq \delta}}{\mathbb{E}_S \mu(S) \mathbb{1}_{\mu(S) \leq \delta}} \\ &\leq \frac{\mathbb{E}_S G(S, V \setminus S)}{\mathbb{E}_S \mu(S) - \mathbb{E}_S \mu(S)^2 / \delta} \quad (\text{using Claim 2.5}) \\ &\leq \frac{\|f\|^2 \sqrt{1 - \langle f, Gf \rangle^2} / \|f\|^4}{\|f\|^2 - \|f\|_1^2 / \delta} \quad (\text{using Claim 2.3, Claim 2.4, and Claim 2.6}). \end{aligned}$$

Therefore, the set S^* satisfies the conclusion of the local Cheeger bound ([Lemma 2.2](#)).

2.3. Unique Games and Semidefinite Relaxation

Unique Games. A *unique game* \mathfrak{U} with vertex set V and alphabet Σ is a distribution over *constraints* $(u, v, \pi) \in V \times V \times \mathbb{S}_\Sigma$, where \mathbb{S}_Σ denotes the set of permutations of Σ . An *assignment* $x \in \Sigma^V$ *satisfies* a constraint (u, v, π) if $x_v = \pi(x_u)$, that is, the permutation π maps labels for u to labels for v . The *value* $\mathfrak{U}(x)$ of an assignment x for \mathfrak{U} is the fraction of the constraints of \mathfrak{U} satisfied by the assignment x , i.e.,

$$\mathfrak{U}(x) \stackrel{\text{def}}{=} \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U}} \{ \pi(x_u) = x_v \}.$$

(Here, $(u, v, \pi) \sim \mathfrak{U}$ denotes a constraint sampled from \mathfrak{U} .) The *optimal value* $\text{opt}(\mathfrak{U})$ is defined as the maximum value of an assignment, i.e.,

$$\text{opt}(\mathfrak{U}) = \max_{x \in \Sigma^V} \mathfrak{U}(x).$$

We will assume that the distribution over constraints is symmetric, in the sense that a constraint (u, v, π) has the same probability as the constraint (v, u, π^{-1}) .

The value of an assignment x corresponds to the success probability of the following probabilistic verification procedure:

1. Sample a random constraint $(u, v, \pi) \sim \mathfrak{U}$.
2. Verify that $x_v = \pi(x_u)$.

2. Background

Constraint and Label-Extended Graph. We can associate two graphs with a unique game. Let \mathfrak{U} be a unique game with vertex set V and alphabet Σ . The *constraint graph* $G(\mathfrak{U})$ is a graph with vertex set V . Its edge distribution is obtained as follows:

1. Sample a random constraint $(u, v, \pi) \sim \mathfrak{U}$.
2. Output the edge uv .

The *label-extended graph* $\hat{G}(\mathfrak{U})$ is a graph with vertex set $V \times \Sigma$. Its edge distribution is obtained in the following way:

1. Sample a random constraint $(u, v, \pi) \sim \mathfrak{U}$.
2. Sample a random label $i \in \Sigma$.
3. Output an edge between (u, i) and $(v, \pi(i))$.

Sometimes it will be convenient to denote a vertex (u, i) in the label-extended graph by u_i .

An assignment $x \in \Sigma^V$ naturally corresponds to a set $S \subseteq V \times \Sigma$ with cardinality $|S| = |V|$ (and therefore, volume $\mu(S) = 1/|\Sigma|$). The value of the assignment x for the unique game \mathfrak{U} corresponds exactly to the expansion of the corresponding set S in the label-extended graph $\hat{G}(\mathfrak{U})$,

$$\mathfrak{U}(x) = 1 - \Phi(S).$$

This correspondence between expansion and the value of an assignment is the basis of the connection between UNIQUE GAMES and SMALL-SET EXPANSION discussed in this thesis.

Partial Unique Games. It is sometimes convenient to consider partial assignments for unique games. Let \mathfrak{U} be a unique game with vertex set V and alphabet Σ . An assignment $x \in (\Sigma \cup \{\perp\})^V$ is α -*partial* if at least an α fraction of the vertices are labeled (with symbols from Σ), i.e.,

$$\mathbb{P}_{(u,v,\pi) \sim \mathfrak{U}} \{x_u \neq \perp\} \geq \alpha.$$

A partial assignment x satisfies a constraint (u, v, π) if both vertices u and v are labeled and their labels satisfy the constraint $x_v = \pi(x_u)$. For conciseness, we write $x_v = \pi(x_u) \in \Sigma$ to denote the event that the partial assignment x satisfies

2. Background

the constraint (u, v, π) . The *value* $\mathfrak{U}(x)$ of a partial assignment x is the fraction of constraints satisfied by x ,

$$\mathfrak{U}(x) \stackrel{\text{def}}{=} \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U}} \{x_v = \pi(x_u) \in \Sigma\}.$$

The α -*partial value* $\text{opt}_\alpha(\mathfrak{U})$ is the maximum value of a α -partial assignment normalized by the fraction of labeled vertices,

$$\text{opt}_\alpha(\mathfrak{U}) \stackrel{\text{def}}{=} \max \left\{ \frac{1}{\mathbb{P}_{u \sim V} \{x_u \neq \perp\}} \mathfrak{U}(x) \mid x \in (\Sigma \cup \{\perp\})^V, \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U}} \{x_u \neq \perp\} \geq \alpha \right\}. \quad (2.1)$$

Note that $\text{opt}(\mathfrak{U}) = \text{opt}_1(\mathfrak{U})$ and $\text{opt}_\alpha(\mathfrak{U}) \leq \text{opt}_{\alpha'}(\mathfrak{U})$ whenever $\alpha \geq \alpha'$.

SDP Relaxation of Unique Games. With every unique game \mathfrak{U} , we associate a *semidefinite program* $\text{SDP}(\mathfrak{U})$ of polynomial size. (Semidefinite programs are instances of SEMIDEFINITE PROGRAMMING, an efficiently solvable¹ (convex) optimization problem.)

Relaxation 2.7 (Semidefinite Relaxation for UNIQUE GAMES).

Given a unique game \mathfrak{U} with vertex V and alphabet Σ , find a collection of vectors $\{\mathbf{u}_i\}_{u \in V, i \in \Sigma} \subseteq \mathbb{R}^d$ for some $d \in \mathbb{N}$ so as to

$$\text{SDP}(\mathfrak{U}): \quad \text{maximize} \quad \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \langle \mathbf{u}_i, \mathbf{v}_{\pi(i)} \rangle \quad (2.2)$$

$$\text{subject to} \quad \sum_{i \in \Sigma} \|\mathbf{u}_i\|^2 = 1 \quad (u \in V), \quad (2.3)$$

$$\langle \mathbf{u}_i, \mathbf{u}_j \rangle = 0 \quad (u \in V, i \neq j \in \Sigma). \quad (2.4)$$

We let $\text{sdp}(\mathfrak{U})$ denote the optimal value of the relaxation $\text{SDP}(\mathfrak{U})$.

Γ -MAX 2-LIN. An interesting special case of UNIQUE GAMES is the Γ -MAX 2-LIN problem. (Indeed, in many situations it suffices to consider the simpler Γ -MAX 2-LIN problem.) For a finite group Γ , a unique game \mathfrak{U} with alphabet Σ is a Γ -MAX 2-LIN instance, if we can identify the alphabet Σ with the group Γ such

¹ More precisely, for every desired accuracy $\varepsilon > 0$, one can efficiently compute a solution with objective value within ε of the optimal value (even if ε is exponentially small in the instance size.) Since we are interested in approximation algorithms for combinatorial problems, we can ignore this (arbitrarily small) error and assume that we can find an optimal solution efficiently.

2. Background

that for every constraint (u, v, π) in \mathfrak{U} , there exists a group element $c \in \Gamma$ such that $a = \pi(b)$ if and only if $ab^{-1} = c$. For example, in the case that $\Gamma = \mathbb{Z}_k$ is the cyclic group of order k , all constraints of \mathfrak{U} have the form $x_u - x_v = c \pmod k$ for some $c \in \mathbb{Z}_k$.

Parallel Repetition. Let \mathfrak{U} be a unique game with vertex set V and alphabet Σ . For $\ell \in \mathbb{N}$, the ℓ -fold (*parallel*) repetition of \mathfrak{U} , denoted $\mathfrak{U}^{\otimes \ell}$, is a unique game with vertex set V^ℓ and alphabet Σ^ℓ . (We sometimes write \mathfrak{U}^ℓ instead of $\mathfrak{U}^{\otimes \ell}$.) The unique game $\mathfrak{U}^{\otimes \ell}$ corresponds to the following probabilistic verifier for an assignment $X: V^\ell \rightarrow \Sigma^\ell$:

1. Sample ℓ constraints $(u_1, v_1, \pi_1), \dots, (u_\ell, v_\ell, \pi_\ell) \sim \mathfrak{U}$.
2. Verify that $X(v_r) = \pi_r(X(u_r))$ for all $r \in \{1, \dots, \ell\}$.

Note that the constraint distribution of $\mathfrak{U}^{\otimes \ell}$ corresponds to the ℓ -fold product of the constraint distribution of \mathfrak{U} .

It is easy to verify that the optimal value of the repeated unique game $\mathfrak{U}^{\otimes \ell}$ is at least $\text{opt}(\mathfrak{U})^\ell$. This lower bound is not always tight. Indeed, Raz [Raz08] showed that for every small $\varepsilon > 0$, there exists a unique games \mathfrak{U} with $\text{opt}(\mathfrak{U}) \leq 1 - \varepsilon$, but $\text{opt}(\mathfrak{U}^{\otimes \ell}) \geq 1 - O(\sqrt{\ell} \varepsilon)$.

If the constraint graph of the unique game \mathfrak{U} is bipartite, then an approach of Feige and Lovász [FL92] shows that the semidefinite value of the repeated game is just a function of the semidefinite value of the unrepeated game.

Theorem 2.8 ([FL92, MS07, KRT08]). *For every unique game \mathfrak{U} with bipartite constraint graph and every $\ell \in \mathbb{N}$,*

$$\text{sdp}(\mathfrak{U}^{\otimes \ell}) = \text{sdp}(\mathfrak{U})^\ell.$$

We note that in many situations it is possible to reduce general unique games to ones with bipartite constraints.

Squares of Unique Games. In Part II (Reductions), we often consider “squares” of unique games (either implicitly or explicitly). For a unique game \mathfrak{U} with vertex set V and alphabet Σ , its square \mathfrak{U}^2 is a unique game with the same vertex set and alphabet. It corresponds to the following probabilistic verifier for an assignment $x \in \Sigma^V$,

1. Sample a random vertex $u \sim V$. (We sample u according to the marginal distribution of the first vertex in a random constraint of \mathfrak{U} .)

2. Background

2. Sample two random constraints $(u, v, \pi), (u, v', \pi') \sim \mathfrak{U} \mid u$ incident to u . (Here, $\mathfrak{U} \mid u$ denotes the constraint distribution of \mathfrak{U} conditioned on the first vertex being u .)
3. Verify that $\pi^{-1}(x_v) = (\pi')^{-1}(x_{v'})$.

It is easy to check that $\mathfrak{U}(x) \geq 1 - \eta$ implies $\mathfrak{U}^2(x) \geq 1 - 2\eta$. On the other hand, if $\mathfrak{U}^2(x) \geq \zeta$, then one can construct an assignment y such that $\mathfrak{U}(y) \geq \zeta/4$.

Let \mathfrak{U}_\cup be the unique game obtained by sampling with probability $1/2$ a random constraint from \mathfrak{U} and sampling with the remaining probability a trivial constraint (u, u, id) . Note that $\mathfrak{U}_\cup(x) = 1/2(1 + \mathfrak{U}(x))$ for every assignment $x \in (\Sigma \cup \{\perp\})^V$.

The square of the unique game \mathfrak{U}_\cup is very close to the unique game \mathfrak{U} . In particular, $\mathfrak{U}_\cup^2(x) = 1/4 + 1/2\mathfrak{U}(x) + 1/4\mathfrak{U}^2(x)$ for every assignment x . Therefore, if $\mathfrak{U}(x) \geq 1 - \eta$, then also $\mathfrak{U}_\cup^2(x) \geq 1 - \eta$. On the other hand, if $\mathfrak{U}_\cup^2(x) \geq 1 - \eta'$, then $\mathfrak{U}(x) \geq 1 - 2\eta'$. (Similar statements also hold for partial assignments.)

2.4. Unique Games Conjecture and Reductions

Unique Games Conjecture. In [Chapter 1](#), we presented a simplified form of the Unique Games Conjecture. This simplified is formally equivalent to the original formulation of conjecture in [[Kho02](#)]. However, this equivalence is non-trivial (it follows from results of [[KKMO07](#)]).

The following statement of the Unique Games Conjecture is very close to the original formulation in [[Kho02](#)]:

Unique Games Conjecture: For every constant $\varepsilon > 0$, there exists $k = k(\varepsilon) \in \mathbb{N}$ such that given a unique game \mathfrak{U} with alphabet size k , it is NP-hard to distinguish between the cases,

YES: $\text{opt}(\mathfrak{U}) \geq 1 - \varepsilon$,

NO: $\text{opt}(\mathfrak{U}) < \varepsilon$.

In fewer words, for every $\varepsilon > 0$, there exists $k \in \mathbb{N}$ such that the promise problem $(1 - \varepsilon, \varepsilon)$ -GAP-UNIQUE GAMES(k) is NP-hard. (Here, UNIQUE GAMES(k) denotes the optimization problem given by unique games with alphabet size k . See [§2.1](#) for the definition of gap-promise problems.)

To clarify further, we spell out what NP-hardness means in this context: For every constant $\varepsilon > 0$, there exists $k = k(\varepsilon) \in \mathbb{N}$ and a polynomial time reduction from 3-SAT to UNIQUE GAMES such that

2. Background

- every satisfiable 3-SAT instance reduces to a unique game with alphabet size k and optimal value at least $1 - \varepsilon$,
- every unsatisfiable 3-SAT instance reduces to a unique game with alphabet size k and optimal value less than ε .

Consequences of the Unique Games Conjecture. We say a promise problem Π is *UG-hard* if there exists a constant $\varepsilon > 0$ and an efficient reduction from UNIQUE GAMES to Π such that every unique game with value at least $1 - \varepsilon$ reduces to a YES instance of Π and every unique game with value less than ε reduces to a NO instance of Π . The reduction needs to be efficient only for constant alphabet size (in particular, the blow-up of the reduction can be exponential in the alphabet size of the unique game). Assuming the Unique Games Conjecture, every UG-hard problem is also NP-hard.

We briefly discuss the known consequences of the Unique Games Conjecture for constraint satisfaction problems (CSPs). This class of problems contains many basic optimization problems, for example, MAX CUT, UNIQUE GAMES(k) (with fixed alphabet size k), MAX 3-LIN, and MAX 3-SAT. Raghavendra [Rag08] showed that for every CSP Π and every constant c , there exists a constant $s = s_\Pi(c)$ such that for every constant $\varepsilon > 0$, there exists an efficient algorithm that achieves a $(c + \varepsilon, s - \varepsilon)$ -approximation and on the other hand, it is UG-hard to achieve a $(c - \varepsilon, s + \varepsilon)$ -approximation. Hence, Raghavendra’s algorithms achieve essentially optimal approximation guarantees for CSPs assuming the Unique Games Conjecture. In [Ste10a], we show that these approximation guarantees for CSPs can in fact be obtained in quasi-linear time.

Noise Graphs. An important ingredient of UG-hardness reductions are “noise graphs”. Two kinds of noise graphs are relevant for this thesis. The first kind are *Gaussian noise graphs*, denoted U_ρ for a parameter $\rho \in [0, 1]$. The vertex set of these graphs is \mathbb{R}^d for some $d \in \mathbb{N}$. The edge distribution of U_ρ on \mathbb{R}^d is given by two d -dimensional vectors sampled from the Gaussian distribution on \mathbb{R}^{2d} with mean 0 and covariance matrix

$$\begin{pmatrix} I & \rho I \\ \rho I & I \end{pmatrix}.$$

Alternatively, we can describe the edge distribution as follows:

2. Background

1. Sample two standard Gaussian vectors x, y in \mathbb{R}^d independently. (The coordinates of x and y are independent Gaussians with mean 0 and variance 1.)
2. Output an edge between x and $\rho x + \sqrt{1 - \rho^2} y$.

We identify U_ρ with the following linear (Markov) operator on $L_2(\mathbb{R}^d)$ (where \mathbb{R}^d is equipped with the standard Gaussian measure γ),

$$U_\rho f(x) \stackrel{\text{def}}{=} \int_{\mathbb{R}^d} f\left(\rho x + \sqrt{1 - \rho^2} y\right) d\gamma(y).$$

The second kind of noise graph, denote T_ρ for some parameter $\rho \in [0, 1]$, is defined on vertex set Ω^R , where Ω is any (finite) probability space and $R \in \mathbb{N}$. The graph T_ρ on Ω^R has the following edge distribution

1. Sample x_1, \dots, x_R independently from Ω .
2. For every $r \in [R]$, sample y_r as follows: With probability ρ , set $y_r = x_r$ and with the remaining probability $1 - \rho$, sample y_r from Ω .
3. Output the edge xy , where $x = (x_1, \dots, x_R)$ and $y = (y_1, \dots, y_R)$.

The graph T_ρ corresponds to the following linear (Markov) operator on $L_2(\Omega^R)$,

$$T_\rho f(x) \stackrel{\text{def}}{=} \mathbb{E}_{y \sim T_\rho(x)} f(y).$$

Here, $y \sim T_\rho(x)$ denotes that y is a random neighbor of x in the graph T_ρ , i.e., for every coordinate $r \in [R]$, we set $y_r = x_r$ with probability ρ and choose $y_r \sim \Omega$ with probability $1 - \rho$.

We remark that both U_ρ and T_ρ have second largest eigenvalue ρ . Since U_ρ and T_ρ are both product operators (acting independently on coordinates), their spectral decomposition is simple to describe. For $k \in \mathbb{N}$, let Q_k be the linear operator on $L_2(\mathbb{R}^d)$ projecting on the subspace spanned by polynomials of degree k that are orthogonal to all polynomials of degree $k - 1$. Similarly, let P_k be the linear operator on $L_2(\Omega^R)$ projecting on subspace spanned by functions that depend only on k coordinates and are orthogonal to all functions depending on at most $k - 1$ coordinates. The operators U_ρ and T_ρ have

2. Background

the following spectral decompositions

$$U_\rho = \sum_{k \in \mathbb{N}} \rho^k Q_k, \quad (2.5)$$

$$T_\rho = \sum_{k \in [R]} \rho^k P_k. \quad (2.6)$$

Another useful property of the operators U_ρ and T_ρ is the following semigroup structure: For all $\rho, \rho' \in [0, 1]$, it holds that

$$U_\rho U_{\rho'} = U_{\rho\rho'}, \quad (2.7)$$

$$T_\rho T_{\rho'} = T_{\rho\rho'}. \quad (2.8)$$

Influence. Let $f \in L_2(\Omega^R)$ for some finite probability space Ω . The *influence* of coordinate i , denoted $\text{Inf}_i f$, is the squared norm of f projected on the space of functions that depend only on coordinate i and are orthogonal to constant functions. Equivalently, $\text{Inf}_i f$ is the typical variance of $f(x)$ for $x_i \sim \Omega$,

$$\text{Inf}_i f \stackrel{\text{def}}{=} \mathbb{E}_{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_R \sim \Omega} \left(\mathbb{E}_{x_i \sim \Omega} f(x)^2 - \left(\mathbb{E}_{x_i \sim \Omega} f(x) \right)^2 \right).$$

From the above definition, it is straight-forward to check that $\text{Inf}_i f$ is convex in f (using Cauchy–Schwarz).

The *total influence* of a function $\sum_{i \in [R]} \text{Inf}_i f$ has the following expansion in terms of the projections $P_k f$,

$$\sum_{i \in [R]} \text{Inf}_i f = \sum_k k \|P_k f\|^2.$$

A consequence is that “noised” functions can have only few influential coordinates.

Fact 2.9. Let $f \in L_2(\Omega^R)$ for a finite probability space Ω and $R \in \mathbb{N}$. Then, for every $\rho \in [0, 1)$,

$$\sum_{i \in [R]} \text{Inf}_i T_\rho f \leq \frac{1}{\log(1/\rho)} \|f\|^2$$

In particular, if $\|f\|^2 = 1$, the number of coordinates with influence larger than $\tau > 0$ is bounded by $\frac{1}{\tau \log(1/\rho)}$.

Proof. The total influence of $T_\rho f$ is equal to $\sum_{k \in [R]} k \rho^k \|P_k f\|^2 \leq \|f\|^2 \max_k k \rho^k \leq \|f\|^2 / \log(1/\rho)$. \square

2. Background

Gaussian Noise Stability. Let $\Gamma_\rho: [0, 1] \rightarrow [0, 1]$ be the *stability profile* of the Gaussian noise graph U_ρ on \mathbb{R} , that is,

$$\Gamma_\rho(\delta) \stackrel{\text{def}}{=} \max_{S \subseteq \mathbb{R}, \mu(S) \leq \delta} U_\rho(S, S),$$

(Recall that the notation $U_\rho(S, S)$ denotes the probability that a random edge of U_ρ has both endpoints in S .) The stability profile of U_ρ is related to its expansion profile in the following way $\Phi_{U_\rho}(\delta) = 1 - \Gamma_\rho(\delta)/\delta$. A result of C. Borell [Bor85] shows that vertex sets of the form $[t, \infty)$ have maximum stability in U_ρ ,

$$\Gamma_\rho(\delta) = \mathbb{P}_{xy \sim U_\rho} \{x \geq t, y \geq t\} \quad \text{for } t \in \mathbb{R} \text{ such that } \mathbb{P}_{x \sim U_\rho} \{x \geq t\} = \delta.$$

The Gaussian noise stability trivially satisfies $\delta^2 \leq \Gamma_\rho(\delta) \leq \delta$ for all $\delta > 0$. (This relation holds for any regular graph.) Let $\rho = 1 - \varepsilon$. Using basic properties of the Gaussian distribution, one can show the following basic bounds on $\Gamma_\rho(\delta)$. For all $\varepsilon, \delta > 0$, the Gaussian noise stability satisfies $\Gamma_\rho(\delta)/\delta \leq \delta^{\varepsilon/2}$ [KKMO07]. If $\delta \gg 2^{-\varepsilon}$, then $\Gamma_\rho(\delta)/\delta \approx 1 - O(\sqrt{\varepsilon \log(1/\delta)})$ (e.g. [CMM06b]).

Invariance Principle. In this work, the following simple form of the invariance principle suffices: The graph T_ρ acts on bounded functions without influential coordinates in approximately the same way as the graph U_ρ . The formal statement is as follows:

Theorem 2.10 (Invariance Principle, [MOO05]). *For every finite probability space Ω and constants $\rho \in [0, 1)$, $\eta > 0$, there exists constants $\tau, \gamma > 0$ such that for every function $f \in L_2(\Omega^R)$ with $0 \leq f \leq 1$, either*

$$\langle f, T_\rho f \rangle \leq \Gamma_\rho(\mathbb{E} f) + \eta,$$

or $\text{Inf}_i T_{1-\gamma} f > \tau$ for some coordinate $i \in [R]$.

Influence Decoding. To describe this construction, we introduce two notations: (1) For a vertex u of a unique game \mathfrak{U} , we write $(u, v, \pi) \sim U \mid u$ to denote a random constraint incident to u in \mathfrak{U} . (In other words, we condition the constraint distribution on the event that u is the first vertex of the constraint.) (2) For a vector $x \in \Omega^R$ (where Ω can be any set) and a permutation $\pi \in \mathbb{S}_R$ of the coordinates, we write $\pi.x$ to denote the vector in Ω^R obtained by

2. Background

permuting the coordinates of x according to π . Formally, for all coordinates $i \in [R]$, we set $(\pi.x)_{\pi(i)} = x_i$.

The following lemma gives a sufficient condition for when we can decode a good assignment for a unique game based on influential coordinates of functions associated with the vertices of the unique game.

Lemma 2.11 (Influence Decoding). *Let \mathfrak{U} be a unique game with vertex set V and alphabet $[R]$. For some probability space Ω , let $\{f_u\}_{u \in V}$ be a collections of normalized functions in $L_2(\Omega^R)$. Consider functions g_u in $L_2(\Omega^R)$ defined by*

$$g_u(x) = \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}|u} f_v(\pi.x).$$

Then, for all $\gamma, \tau > 0$, there exists $c_{\gamma,\tau} > 0$ (in fact, $c_{\gamma,\tau} = \text{poly}(\gamma, \tau)$) such that

$$\text{opt}(\mathfrak{U}) \geq c_{\gamma,\tau} \cdot \mathbb{P}_{u \sim V} \left\{ \max_{i \in [R]} \text{Inf}_i T_{1-\gamma} g_u > \tau \right\}.$$

Proof. We construct an assignment for the unique game \mathfrak{U} by the following probabilistic procedure:

1. For every vertex $u \in V$, randomly choose one of the following ways to determine the label x_u :

- a) Assign x_u randomly from the set

$$L_f(u) := \{i \in [R] \mid \text{Inf}_i T_{1-\gamma} f_u > \tau\}.$$

(If the set is empty, assign $x_u = \perp$.)

- b) Assign x_u randomly from the set

$$L_g(u) := \{i \in [R] \mid \text{Inf}_i T_{1-\gamma} g_u > \tau\}.$$

(If the set is empty, assign $x_u = \perp$.)

First, we note that the sets $L_f(u)$ and $L_g(u)$ cannot be large. (By [Fact 2.9](#) at most $1/\tau\gamma$ coordinates of say $T_{1-\gamma} f_u$ can have influence larger than τ) To estimate the expected fraction of constraints satisfied by x_u , consider a vertex u such that $\text{Inf}_i T_{1-\gamma} g_u > \tau$ for some label $i \in [R]$. Since $i \in L_g(u)$, we assign $x_u = i$ with probability at least $1/2 \cdot 1/|L_g(u)| \geq \tau\gamma/2$.

2. Background

Next, we estimate the probability that $x_v = \pi(i)$ for a constraint $(u, v, \pi) \sim \mathfrak{U} \mid u$. Using the convexity of Inf_i and the fact that coordinate i for the function $x \mapsto f_u(\pi.x)$ corresponds to coordinate $\pi(i)$ for the function f , we derive that

$$\tau < \text{Inf}_i T_{1-\gamma} g_u \leq \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U} \mid u} \text{Inf}_{\pi(i)} T_{1-\gamma} f_v,$$

which means that $\pi(i) \in L_f(v)$ with probability at least τ (over the choice of $(u, v, \pi) \sim \mathfrak{U} \mid u$). Hence, the assignment x satisfies in expectation at least

$$\begin{aligned} \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U} \mid u} \{x_v = \pi(x_u)\} &\geq \mathbb{P}_{(u,v,\pi) \sim UG \mid u} \{x_u = i \wedge x_v = \pi(i)\} \\ &\geq (\tau\gamma/2) \mathbb{P}_{(u,v,\pi) \sim UG \mid u} \{x_v = \pi(i)\} \\ &\geq (\tau^2\gamma^2/4) \mathbb{P}_{(u,v,\pi) \sim UG \mid u} \{\pi(i) \in L_f(v)\} \geq (\tau^3\gamma^2/4) \end{aligned}$$

of the constraints (u, v, π) incident to u .

The expected value of the assignment x is therefore at least

$$\begin{aligned} \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U}} \{x_v = \pi(x_u)\} &= \mathbb{E}_{u \sim V} \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U} \mid u} \{x_v = \pi(x_u)\} \\ &\geq (\tau^3\gamma^2/4) \mathbb{P}_{u \sim V} \{\exists i. \text{Inf}_i T_{1-\gamma} g_u > \tau\}. \quad \square \end{aligned}$$

Part I.
Algorithms

3. Unique Games with Expanding Constraint Graphs

Many constraint satisfaction problems have natural candidates for hard input distributions. For example, Feige’s [Fei02] Random 3-SAT hypothesis asserts that it is intractable to distinguish near-satisfiable 3-SAT instances and *random* 3-SAT instances of appropriate density (which are far from being satisfiable with high probability). A striking evidence for this hypothesis is that even very strong semidefinite relaxations ($\Omega(n)$ -levels of Lasserre’s hierarchy [Las01]) fail to distinguish between the two cases [Sch08].

In this chapter, we present an efficient algorithm for UNIQUE GAMES that provides a good approximation whenever the underlying constraint graph is mildly expanding. Since random instances tend to be expanding, this algorithm shows that UNIQUE GAMES is easy for many input distributions (for example, Erdős–Renyi graphs with logarithmic average degree or random regular graphs with constant degree).

Given a unique game with optimal value $1 - \varepsilon$, our algorithm finds an assignment of value $1 - O(\varepsilon/\lambda)$, where $\lambda \in [0, 2]$ is the spectral gap of the constraint graph.

In contrast to previous approximation algorithms for UNIQUE GAMES, the approximation guarantee of this algorithm does not degrade with the size of the alphabet or the number of variables of the unique game.

3.1. Main Result

Let \mathfrak{U} be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$. Let $\lambda \in [0, 2]$ be the spectral gap of the constraint graph of \mathfrak{U} . In this section, we present an efficient algorithm that can satisfy a constant fraction of the constraints of \mathfrak{U} if λ is significantly larger than the fraction of constraints violated in an optimal assignment for \mathfrak{U} . The algorithm is based on the standard semidefinite relaxation for UNIQUE GAMES (see [Relaxation 2.7](#)). In the following, we denote this relaxation by $\text{SDP}(\mathfrak{U})$ and its optimal value by

3. Unique Games with Expanding Constraint Graphs

$\text{sdp}(\mathfrak{U})$. (We refer to the optimal value of the relaxation $\text{SDP}(\mathfrak{U})$ sometimes as the *semidefinite value* of \mathfrak{U} .) For the convenience of the reader, we restate the relaxation $\text{SDP}(\mathfrak{U})$.

Given a unique game \mathfrak{U} with vertex V and alphabet Σ , find a collection of vectors $\{\mathbf{u}_i\}_{u \in V, i \in \Sigma} \subseteq \mathbb{R}^d$ for some $d \in \mathbb{N}$ so as to

$$\text{SDP}(\mathfrak{U}): \quad \text{maximize} \quad \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \langle \mathbf{u}_i, \mathbf{v}_{\pi(i)} \rangle \quad (3.1)$$

$$\text{subject to} \quad \sum_{i \in \Sigma} \|\mathbf{u}_i\|^2 = 1 \quad (u \in V), \quad (3.2)$$

$$\langle \mathbf{u}_i, \mathbf{u}_j \rangle = 0 \quad (u \in V, i \neq j \in \Sigma). \quad (3.3)$$

Theorem 3.1. *Given a unique game \mathfrak{U} with spectral gap λ and semidefinite value $\text{sdp}(\mathfrak{U}) \geq 1 - \varepsilon$, we can efficiently compute an assignment x for \mathfrak{U} of value $\mathfrak{U}(x) \geq 1 - O(\varepsilon/\lambda)$.*

The main component of the above theorem is a simple (randomized) rounding algorithm that given a solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$, outputs an assignments x for \mathfrak{U} of value $\mathfrak{U}(x) \geq 1 - O(\varepsilon/\lambda)$.

A solution for $\text{SDP}(\mathfrak{U})$ consists of a vector $\mathbf{u}_i \in \mathbb{R}^d$ for every vertex $u \in V$ and label $i \in \Sigma$. To round this solution, we first pick a random “pivot vertex” $w \in V$ and choose a “seed label” $s \in \Sigma$ for w according to the probability distribution given by the weights $\|\mathbf{w}_1\|^2, \dots, \|\mathbf{w}_k\|^2$. Next, we propagate this initial choice in one step to every other vertex $u \in V$: Roughly speaking, we assign label $i \in \Sigma$ to vertex u if the vector \mathbf{u}_i is “closest” to \mathbf{w}_s .

We remark that unlike previous rounding algorithms for semidefinite relaxations, this rounding is not based on random projections of the solution vectors. (Previous projection-based algorithms in fact cannot achieve the approximation in [Theorem 3.1](#).)

In the following, we give a more precise description of the rounding algorithm used for [Theorem 3.1](#). For any two vertices $u, v \in V$, the algorithm defines a distance $\Delta(u, v)$, which measures the similarity of the two vector collections $\{\mathbf{u}_i\}_{i \in \Sigma}$ and $\{\mathbf{v}_i\}_{i \in \Sigma}$ (similar to the so called *earth mover’s distance*). These distances will play an important role in the analysis of the rounding algorithm.

Algorithm 3.2. (Propagation Rounding)

3. Unique Games with Expanding Constraint Graphs

Input: Vector solution $\{\mathbf{u}_i\}_{u \in V, i \in \Sigma}$ for the relaxation $\text{SDP}(\mathfrak{U})$.

Output: Distribution over $k \cdot n$ assignments for the unique game \mathfrak{U} .

The distribution is specified by the following randomized algorithm:

- For every vertex $u \in V$, let D_u be the distribution over Σ which outputs label $i \in \Sigma$ with probability $\|\mathbf{u}_i\|^2$.
- For any two vertices $u, v \in V$, let

$$\Delta(u, v) \stackrel{\text{def}}{=} \min_{\sigma \in \mathfrak{S}_\Sigma} \sum_{i \in \Sigma} \frac{1}{2} \|\mathbf{u}_i - \mathbf{v}_{\sigma(i)}\|^2$$

and let $\sigma_{v \leftarrow u}$ be a permutation of Σ , for which the minimum $\Delta(u, v)$ is attained.

- *Pivot vertex:* Sample a vertex $w \in V$.
- *Seed label:* Sample a label $s \sim D_w$.
- *Propagation:* For every vertex $u \in V$, set $x_u := \sigma_{u \leftarrow w}(s)$ and output the assignment $x \in \Sigma^V$.

We will show that starting from a solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$, [Algorithm 3.2](#) computes an assignment x that violates at most $O(\varepsilon/\lambda)$ of the constraints of \mathfrak{U} in expectation. (It is easy to derandomize [Algorithm 3.2](#), because the number of choices for the pivot vertex w and the seed label s is at most $k \cdot n$. Thus, we could modify the algorithm such that it outputs a list of $k \cdot n$ assignments $x^{(w,s)}$ for \mathfrak{U} . One of these assignments $x^{(w,s)}$ has value at least $1 - O(\varepsilon/\lambda)$ for \mathfrak{U} .)

Let us consider a constraint (u, v, π) in \mathfrak{U} . The following lemma relates the probability that this constraint is violated to three distance terms. The first term is exactly the (negative) contribution of the constraint (u, v, π) to the value of our solution for $\text{SDP}(\mathfrak{U})$. The remaining terms are averages of the distance $\Delta(u, w)$ and $\Delta(v, w)$ over all choices of vertices $w \in V$. We will prove this lemma in the next section ([§3.2](#)).

Lemma 3.3 (Propagation Rounding). *For every constraint (u, v, π) of \mathfrak{U} ,*

$$\mathbb{P}_x \{x_v \neq \pi(x_u)\} \leq 3 \sum_{i \in \Sigma} \frac{1}{2} \|\mathbf{u}_i - \mathbf{v}_{\pi(i)}\|^2 + 3 \mathbb{E}_{w \in V} \Delta(u, w) + 3 \mathbb{E}_{w \in V} \Delta(v, w).$$

3. Unique Games with Expanding Constraint Graphs

To prove [Theorem 3.1](#) using the lemma above, we need to bound the expected value of $\Delta(u, v)$ for two random vertices $u, v \in V$. It is easy to show that for a random constraint (u, v, π) of \mathfrak{U} , the expectation of $\Delta(u, v)$ is at most ε . We will show that the expansion of the constraint graph implies that the typical value of $\Delta(u, v)$ for two random vertices $u, v \in V$ cannot be much larger than the typical value of $\Delta(u, v)$ for a random constraint (u, v, π) in \mathfrak{U} . Here, the crux is that the distances $\Delta(u, v)$ can be approximated up to constant factors by squared euclidean distances. (We defer the rather technical proof of this fact to [Section 3.3](#). Since the construction that establishes this fact is somewhat mysterious and relies on tensor products of vectors, we refer to it as “tensoring trick”.)

Lemma 3.4 (Tensoring Trick). *There exist unit vectors $X_1, \dots, X_n \in \mathbb{R}^n$ such that for any two vertices $u, v \in V$,*

$$c \cdot \Delta(u, v) \leq \|X_u - X_v\|^2 \leq C \cdot \Delta(u, v).$$

Here, c and C are absolute constants. (We can choose $c = 1/4$ and $C = 20$).

The spectral gap λ of the constraint graph of \mathfrak{U} allows us to relate the typical value of $\|X_u - X_v\|^2$ for two random vertices $u, w \in V$ to the typical value of $\|X_u - X_v\|^2$ for a constraint (u, v, π) in \mathfrak{U} ,

$$\lambda \leq \frac{\mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \|X_u - X_v\|^2}{\mathbb{E}_{u,v \in V} \|X_u - X_v\|^2}.$$

Together with [Lemma 3.4](#), we can show that in expanding constraint graphs, local correlation implies global correlation. Here, *local correlation* refers to the situation that $\Delta(u, v)$ is small for a typical constraint (u, v, π) of \mathfrak{U} . *Global correlation* means that $\Delta(u, v)$ is small for two randomly chosen vertices $u, v \in V$.

Corollary 3.5 (Local Correlation \implies Global Correlation).

$$\mathbb{E}_{u,v \in V} \Delta(u, v) \leq (C/c\lambda) \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \Delta(u, v).$$

By combining [Lemma 3.3](#) and [Corollary 3.5](#), we show in the next lemma that the assignment computed by [Algorithm 3.2](#) satisfies $1 - O(\varepsilon/\lambda)$ of the constraints of \mathfrak{U} in expectation. [Theorem 3.1](#) is implied by this lemma.

Lemma 3.6. *The assignment x computed by [Algorithm 3.2](#) given a solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$ satisfies*

$$\mathbb{E}_x \mathfrak{U}(x) \geq 1 - O(\varepsilon/\lambda).$$

3. Unique Games with Expanding Constraint Graphs

Proof. Using [Lemma 3.3](#) (Propagation Rounding) and the fact that we have a solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$, the expected value of the assignment x is at least

$$\begin{aligned} \mathbb{E}_x \mathfrak{U}(x) &\geq 1 - 3 \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \left[\sum_i \frac{1}{2} \|u_i - v_{\pi(i)}\|^2 - 3 \mathbb{E}_{w \in V} \Delta(u, w) - 3 \mathbb{E}_{w \in V} \Delta(v, w) \right] \\ &= 1 - 3\varepsilon - 6 \mathbb{E}_{u,v \in V} \Delta(u, v). \end{aligned}$$

(Here, we also used that the constraint graph of \mathfrak{U} is regular.) From [Corollary 3.5](#) (Local Correlation \implies Global Correlation), it follows that

$$\mathbb{E}_{u,v \in V} \Delta(u, v) \leq (C/c\lambda) \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \Delta(u, v) \leq (C/c\lambda) \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \sum_i \frac{1}{2} \|u_i - v_{\pi(i)}\|^2 = (C/c\lambda)\varepsilon.$$

We conclude that

$$\mathbb{E}_x \mathfrak{U}(x) \geq 1 - 3\varepsilon - 6(C/c\lambda)\varepsilon = 1 - O(\varepsilon/\lambda). \quad \square$$

3.2. Propagation Rounding

In this section, we will prove [Lemma 3.3](#), which we used to analyze [Algorithm 3.2](#) in the previous section ([§3.1](#)). Let us quickly recall some details of this algorithm. For every vertex pair $u, v \in V$, the algorithm defines a permutation $\sigma_{v \leftarrow u}$ of Σ . The algorithm outputs an assignment x defined as $x_u = \sigma_{u \leftarrow w}(s)$, where w is a random vertex and s is a random label drawn from the probability distribution D_w given by the weights $\|\mathbf{w}_1\|^2, \dots, \|\mathbf{w}_k\|^2$. For $u, v \in V$ and $\sigma \in \mathbb{S}_\Sigma$, we introduce the following notation

$$\Delta(u, v, \sigma) \stackrel{\text{def}}{=} \sum_{i \in \Sigma} \frac{1}{2} \|u_i - v_{\sigma(i)}\|^2.$$

With this notation, [Lemma 3.3](#) asserts the following bound on the probability that x does not satisfy a constraint (u, v, π) of \mathfrak{U} ,

$$\mathbb{P}_x \{x_v \neq \pi(x_u)\} \leq 3\Delta(u, v, \pi) + \mathbb{E}_{w \in V} 3\Delta(w, v, \sigma_{v \leftarrow w}) + \mathbb{E}_{w \in V} 3\Delta(w, u, \sigma_{u \leftarrow w}). \quad (3.4)$$

The following lemma implies [\(3.4\)](#) (and therefore [Lemma 3.3](#)) by averaging over $w \in V$ and setting $\sigma = \sigma_{u \leftarrow w}$ and $\sigma' = \sigma_{v \leftarrow w}$.

3. Unique Games with Expanding Constraint Graphs




Figure 3.1.: Illustration of [Lemma 3.7](#). The permutation π maps labels from vertex u to vertex v . The permutations σ and σ' map labels from vertex w to vertices u and v , respectively.

Lemma 3.7. *For any vertices $u, v, w \in V$ and permutations π, σ, σ' of Σ ,*

$$\mathbb{P}_{s \sim D_w} \left\{ \sigma'(s) \neq \pi \circ \sigma(s) \right\} \leq 3\Delta(u, v, \pi) + 3\Delta(w, u, \sigma') + 3\Delta(w, v, \sigma).$$

Proof. For simplicity, assume that σ and σ' are the identity permutation. (We can achieve this situation by reordering the vector collections $\{\mathbf{u}_i\}_{i \in \Sigma}$ and $\{\mathbf{v}_j\}_{j \in \Sigma}$ and changing π accordingly.) Since the vectors $\mathbf{w}_1, \dots, \mathbf{w}_k$ are pairwise orthogonal, we have $\|\mathbf{w}_s - \mathbf{w}_{s'}\|^2 = \|\mathbf{w}_s\|^2 + \|\mathbf{w}_{s'}\|^2$ for $s \neq s' \in \Sigma$. Therefore,

$$\mathbb{P}_{s \in D_w} \left\{ s \neq \pi(s) \right\} = \sum_{s \in \Sigma} \frac{1}{2} \|\mathbf{w}_s - \mathbf{w}_{\pi(s)}\|^2.$$

By the triangle inequality (and Cauchy–Schwarz),

$$\begin{aligned} \|\mathbf{w}_s - \mathbf{w}_{\pi(s)}\|^2 &\leq \left(\|\mathbf{w}_s - \mathbf{u}_s\| + \|\mathbf{u}_s - \mathbf{v}_{\pi(s)}\| + \|\mathbf{v}_{\pi(s)} - \mathbf{w}_{\pi(s)}\| \right)^2 \\ &\leq 3\|\mathbf{w}_s - \mathbf{u}_s\|^2 + 3\|\mathbf{u}_s - \mathbf{v}_{\pi(s)}\|^2 + 3\|\mathbf{v}_{\pi(s)} - \mathbf{w}_{\pi(s)}\|^2. \end{aligned}$$

Taking these bounds together,

$$\begin{aligned} \mathbb{P}_{s \in D_w} \left\{ s \neq \pi(s) \right\} &\leq 3 \sum_s \frac{1}{2} \|\mathbf{w}_s - \mathbf{u}_s\|^2 + 3 \sum_s \frac{1}{2} \|\mathbf{u}_s - \mathbf{v}_{\pi(s)}\|^2 + 3 \sum_s \frac{1}{2} \|\mathbf{v}_{\pi(s)} - \mathbf{w}_{\pi(s)}\|^2 \\ &= 3\Delta(w, u, \text{id}) + 3\Delta(u, v, \pi) + 3\Delta(w, v, \text{id}). \quad \square \end{aligned}$$

3.3. Tensoring Trick

In this section, we prove [Lemma 3.4](#), which allowed us in [Section 3.1](#) to relate the approximation guarantee of [Algorithm 3.2](#) on a unique game \mathfrak{U} to the spectral gap of the constraint graph of \mathfrak{U} .

3. Unique Games with Expanding Constraint Graphs

Lemma (Restatement of [Lemma 3.4](#)). *There exist unit vectors $X_1, \dots, X_n \in \mathbb{R}^n$ such that for any two vertices $u, v \in V$,*

$$c \cdot \Delta(u, v) \leq \|X_u - X_v\|^2 \leq C \cdot \Delta(u, v).$$

Here, c and C are absolute constants. (We can choose $c = 1/4$ and $C = 20$).

For a vertex $u \in V$, we choose X_u as (weighted) sum of high tensor-powers of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$,

$$X_u := \sum_i \|\mathbf{u}_i\| \bar{\mathbf{u}}_i^{\otimes t}.$$

Here, $t \in \mathbb{N}$ is a large enough odd integer (say $t = 5$), $\bar{\mathbf{u}}_i$ is the unit vector in direction of \mathbf{u}_i , and $\bar{\mathbf{u}}_i^{\otimes t}$ is the t -fold tensor product¹ of $\bar{\mathbf{u}}_i$ with itself. The vector X_u has unit norm, because the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ are pairwise orthogonal and $\sum_i \|\mathbf{u}_i\|^2 = 1$.

We remark that the constructed vectors X_1, \dots, X_n are elements of \mathbb{R}^{n^t} . However, it is straight-forward to find vectors in \mathbb{R}^n with the same pairwise distances and norms (e.g., by passing to the n -dimensional subspace spanned by the vectors X_1, \dots, X_n).

Consider two arbitrary vertices $u, v \in V$. Let $A = (A_{ij})$ be the k -by- k matrix of inner products between the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ and $\mathbf{v}_1, \dots, \mathbf{v}_k$,

$$A_{ij} := \langle \bar{\mathbf{u}}_i, \bar{\mathbf{v}}_j \rangle.$$

Let $p, q \in \mathbb{R}^k$ be the (nonnegative) unit vectors

$$p := (\|\mathbf{u}_1\|, \dots, \|\mathbf{u}_k\|), \quad q := (\|\mathbf{v}_1\|, \dots, \|\mathbf{v}_k\|).$$

With these notations, $\|X_u - X_v\|^2$ and $\Delta(u, v)$ satisfy the following identities,

$$\frac{1}{2} \|X_u - X_v\|^2 = 1 - \sum_{ij} p_i A_{ij}^t q_j, \tag{3.5}$$

$$\Delta(u, v) = 1 - \max_{\sigma \in \mathbb{S}_\Sigma} \sum_i p_i A_{i\sigma(i)} q_{\sigma(i)}. \tag{3.6}$$

To simplify notation, we will prove [Lemma 3.4](#) in this more general setting of matrices A and unit vectors p and q .

¹The *tensor product* of two vectors $x, y \in \mathbb{R}^d$ is a vector $x \otimes y \in \mathbb{R}^{d \times d}$ with $(x \otimes y)_{i,j} = x_i y_j$.

3. Unique Games with Expanding Constraint Graphs

It is enough for the matrix A to satisfy the following property,

$$\forall j. \sum_i A_{ij}^2 \leq 1, \quad \text{and} \quad \forall i. \sum_j A_{ij}^2 \leq 1. \quad (3.7)$$

(In the setting of [Lemma 3.4](#), this property holds, because by orthogonality $\sum_i \langle \bar{\mathbf{u}}_i, \bar{\mathbf{v}}_j \rangle^2 \leq \|\bar{\mathbf{v}}_j\|^2 = 1$ and $\sum_j \langle \bar{\mathbf{u}}_i, \bar{\mathbf{v}}_j \rangle^2 \leq \|\bar{\mathbf{u}}_i\|^2 = 1$.)

We describe how to prove [Lemma 3.4](#) using the lemmas in the rest of the section. The first two lemmas are simple technical observations, which are used in the proofs of the main lemmas, [Lemma 3.10](#) and [Lemma 3.11](#). On the one hand, [Lemma 3.10](#) shows that for any permutation σ ,

$$(1 - \sum_{ij} p_i A_{ij}^t q_j) \leq 4t(1 - \sum_i p_i A_{i\sigma(i)} q_{\sigma(i)}),$$

which implies $\|X_u - X_v\|^2 \leq 4t\Delta(u, v)$. On the other hand, [Lemma 3.11](#) shows that there exists a permutation σ such that

$$(1 - \sum_i p_i A_{i\sigma(i)} q_{\sigma(i)}) \leq 2 \sum_{ij} p_i A_{ij}^t q_j / (1 - 2^{-(t-3)/2}),$$

which means $\Delta(u, v) \leq 2\|X_u - X_v\|^2 / (1 - 2^{-(t-3)/2})$. Together, these two lemmas imply [Lemma 3.4](#).

Lemma 3.8. *Let $p, q \in \mathbb{R}^k$ be two unit vectors and $A \in \mathbb{R}^{k \times k}$ be a matrix with property (3.7). Suppose that for $\gamma \geq 0$ and $s \in \mathbb{N}$ with $s \geq 2$,*

$$\sum_i p_i A_{ii}^s q_i = 1 - \gamma.$$

Then,

$$\sum_{i \neq j} |p_i A_{ij}^s q_j| \leq \gamma.$$

Proof. Since $s \geq 2$ and A satisfies (3.7), every row and column sum of the matrix $B = (|A_{ij}|^s)$ is bounded by 1. It follows that the largest singular value of B is at most 1. Hence, $\sum_{ij} |p_i A_{ij}^s q_j| \leq \|p\| \|q\| = 1$, which implies the lemma. For the convenience of the reader, we give a direct proof that $\sum_{ij} |p_i A_{ij}^s q_j| \leq 1$,

$$\begin{aligned} \sum_{ij} |p_i A_{ij}^s q_j| &\leq \sum_{ij} |A_{ij}^s| \cdot \left(\frac{1}{2} p_i^2 + \frac{1}{2} q_j^2 \right) \\ &= \frac{1}{2} \sum_i p_i^2 \sum_j |A_{ij}^s| + \frac{1}{2} \sum_j q_j^2 \sum_i |A_{ij}^s| \leq \frac{1}{2} \sum_i p_i^2 + \frac{1}{2} \sum_j q_j^2 = 1. \quad \square \end{aligned}$$

3. Unique Games with Expanding Constraint Graphs

Lemma 3.9. Let $p, q \in \mathbb{R}^k$ be two unit vectors and $A \in \mathbb{R}^{k \times k}$ be a matrix with property (3.7). Let I^+ be the set of indices i such that $p_i A_{ii} q_i \geq 0$. Suppose that for $\gamma \geq 0$ and an odd number $s \in \mathbb{N}$,

$$\sum_{i \in I^+} p_i A_{ii}^s q_i = 1 - \gamma. \quad (3.8)$$

Then,

$$\sum_i p_i A_{ii}^s q_i \geq 1 - 2\gamma.$$

Proof. Note that since s is odd, $p_i A_{ii}^s q_i$ has the same sign as $p_i A_{ii} q_i$ for every index i . From (3.8) and the fact that $|A_{ii}| \leq 1$, it follows that $\sum_{i \in I^+} |p_i q_i| \geq 1 - \gamma$. On the other hand, $\sum_i |p_i q_i| \leq \|p\| \cdot \|q\| = 1$ by Cauchy–Schwarz. Hence, $\sum_{i \notin I^+} |p_i q_i| \leq \gamma$. We conclude that

$$\sum_i p_i A_{ii}^s q_i \geq \sum_{i \in I^+} p_i A_{ii}^s q_i - \sum_{i \notin I^+} |p_i q_i| \geq 1 - 2\gamma. \quad \square$$

Lemma 3.10. Let $p, q \in \mathbb{R}^k$ be two unit vectors and $A \in \mathbb{R}^{k \times k}$ be a matrix with property (3.7). Suppose that for $\gamma \geq 0$,

$$\sum_i p_i A_{ii} q_i = 1 - \gamma. \quad (3.9)$$

Then,

$$\sum_{ij} p_i A_{ij}^t q_j \geq 1 - 4t\gamma.$$

Proof. By Lemma 3.8, it is enough to show

$$\sum_i p_i A_{ii}^t q_i \geq 1 - 2t\gamma. \quad (3.10)$$

As in Lemma 3.9, let I^+ be the set of indices i such that $p_i A_{ii} q_i \geq 0$. Note that since t is odd, $p_i A_{ii}^t q_i$ has the same sign as $p_i A_{ii} q_i$. By the convexity of $x \mapsto x^t$ on \mathbb{R}_+ ,

$$\sum_{i \in I^+} p_i A_{ii}^t q_i = \sum_{i \in I^+} |p_i q_i| |A_{ii}|^t \geq \left(\sum_{i \in I^+} |p_i q_i| |A_{ii}| \right)^t \stackrel{(3.9)}{\geq} (1 - \gamma)^t \geq 1 - t\gamma. \quad (3.11)$$

(For the first inequality, we also used that $\sum_i |p_i q_i| \leq \|p\| \cdot \|q\| = 1$, by Cauchy–Schwarz.) By Lemma 3.9, the bound on $\sum_{i \in I^+} p_i A_{ii}^t q_i$ in (3.11) implies (3.10) as desired. \square

3. Unique Games with Expanding Constraint Graphs

Lemma 3.11. *Let $p, q \in \mathbb{R}^k$ be two unit vectors and $A \in \mathbb{R}^{k \times k}$ be a matrix with property (3.7). Suppose that for $\gamma \geq 0$,*

$$\sum_{ij} p_i A_{ij}^t q_j = 1 - \gamma. \quad (3.12)$$

Then, there exists a permutation σ such that

$$\sum_{ij} p_i A_{i\sigma(i)} q_{\sigma(i)} \geq 1 - 2\gamma / (1 - 2^{-(t-3)/2}).$$

Proof. Since A satisfies (3.7), every row and column contains at most one entry larger than $\sqrt{1/2}$ in absolute value. Hence, there exists a permutation σ such that $j = \sigma(i)$ for any two indices i, j with $a_{ij}^2 > 1/2$. By reordering the columns of A , we can assume $\sigma = \text{id}$. Now, A satisfies $a_{ij}^2 \leq 1/2$ for all indices $i \neq j$.

As in Lemma 3.9, let I^+ be the set of indices i with $p_i A_{ii} q_i \geq 0$. Note that $\text{sign}(p_i A_{ii} q_i) = \text{sign}(p_i A_{ii}^3 q_i) = \text{sign}(p_i A_{ii}^t q_i)$ (because t is odd).

Choose $\eta \geq 0$ such that $1 - \eta = \sum_{i \in I^+} p_i A_{ii}^3 q_i$. From Lemma 3.8, it follows that $\sum_{i \neq j} |p_i A_{ij}^3 q_j| \leq \eta$. Therefore, the contribution of the off-diagonal terms to (3.12) is at most

$$\sum_{i \neq j} p_i A_{ij}^t q_j \leq \max_{i \neq j} |A_{ij}|^{t-3} \cdot \sum_{i \neq j} |p_i A_{ij}^3 q_j| \leq \underbrace{2^{-(t-3)/2}}_{\alpha_t} \cdot \eta. \quad (3.13)$$

Combining (3.12) and (3.13) yields

$$\sum_{i \in I^+} p_i A_{ii}^t q_i \geq \sum_{ij} p_i A_{ij}^t q_j - \sum_{i \neq j} p_i A_{ij}^t q_j \geq 1 - \gamma - \alpha_t \eta.$$

On the other hand,

$$\sum_{i \in I^+} p_i A_{ii}^t q_i \leq \sum_{i \in I^+} p_i A_{ii}^3 q_i = 1 - \eta.$$

Together these two bounds on $\sum_{i \in I^+} p_i A_{ii}^t q_i$ imply $\eta \leq \gamma / (1 - \alpha_t)$. We conclude that

$$\sum_{i \in I^+} p_i A_{ii} q_i \geq \sum_{i \in I^+} p_i A_{ii}^t q_i \geq 1 - \gamma / (1 - \alpha_t),$$

which, by Lemma 3.9, implies that $\sum_i p_i A_{ii} q_i \geq 1 - 2\gamma / (1 - \alpha_t)$ as desired. \square

3.4. Notes

The material presented in this chapter is based on the paper “Unique Games on Expanding Constraints Graphs is Easy” [AKK⁺08], joint with Sanjeev Arora, Subhash Khot, Alexandra Kolla, Madhur Tulsiani, and Nisheeth Vishnoi. A preliminary version of the paper appeared at STOC 2008.

Tensoring Trick

The “tensoring trick” presented in §3.3 originates from a construction in Khot and Vishnoi’s influential work [KV05] on integrality gaps in the context of the Unique Games Conjecture. The result presented in Chapter 8 is a continuation of their work. Again the tensoring trick plays an important role in the constructions presented in Chapter 8.

We remark that the conference version of the paper [AKK⁺08] only establishes a weaker version of Lemma 3.4 (which leads to an additional logarithmic factor in the bound of Theorem 3.1). The proof of the tighter bound is only slightly more involved than the proof of the weaker bound presented in [AKK⁺08].

Propagation Rounding

Algorithm 3.2 (Propagation Rounding) is related to one of Trevisan’s algorithms for UNIQUE GAMES [Tre05]. Trevisan’s algorithm fixes a spanning tree of the constraint graph of the unique game and propagates labels along the edges of this tree (as in Algorithm 3.2 the propagation is according to a solution of an SDP relaxation). In contrast to Trevisan’s algorithms, our algorithm propagates labels in one step to every vertex in the graph. (In other words, the propagation is according to a star graph, a spanning tree of the complete graph with radius 1.)

Some ingredients of the analysis of Algorithm 3.2 are inspired by Trevisan’s analysis (especially, parts of the proof of Lemma 3.7).

Subsequent Work

Makarychev and Makarychev [MM09] established the following improvement of Theorem 3.1. Let \mathfrak{U} be a unique game with optimal value $1 - \varepsilon$. Suppose the spectral gap λ of the constraint graph of \mathfrak{U} satisfies $\lambda \gg \varepsilon$. Then, the algorithm in [MM09] efficiently finds an assignment for \mathfrak{U} of value at least

3. Unique Games with Expanding Constraint Graphs

$1 - O(\epsilon/\Phi)$. Since $\Phi \geq \lambda$ (and in fact, Φ can be as large as $\Omega(\sqrt{\lambda})$), their bound is an improvement over the bound in [Theorem 3.1](#). However, it is important to note that their result only applies to instances for which the algorithm in [Theorem 3.1](#) finds an assignment with constant value, say value at least $1/2$. Hence, in the context of the Unique Games Conjecture, the algorithm in [\[MM09\]](#) does not increase the range of “easy instances” for UNIQUE GAMES.

[\[AIMS10\]](#) and independently [\[RS10\]](#) extended the results in the current chapter to “small-set expanders”. The precise notion of small-set expanders in these two works are quantitatively similar, but strictly speaking incomparable. (Surprisingly, the techniques are quite different except for the common root in the results of this chapter.)

Kolla and Tulsiani [\[KT07\]](#) found an alternative proof of (a quantitatively weaker version of) [Theorem 3.1](#). The interesting feature of their proof is that it avoids explicit use of semidefinite programming and instead relies on enumerating certain low-dimensional subspaces (suitably discretized). Kolla [\[Kol10\]](#) noted that this algorithm also works (in slightly super-polynomial time) on interesting instances of UNIQUE GAMES whose constraint graphs do not have large enough spectral gaps, especially the Khot–Vishnoi instances [\[KV05\]](#). (The spectral gap of the constraint graphs of these instances is too small for the bound in [Theorem 3.1](#) to be non-trivial.) The results presented in [Section 5.2](#) (Subspace Enumeration) are inspired by the works [\[KT07\]](#) and [\[Kol10\]](#).

4. Parallel Repeated Unique Games

Proofs of inapproximability results similar to the Unique Games Conjecture are frequently based on hardness-amplifying self-reductions. A simple example is MAX CLIQUE. To approximate the maximum size of a clique in a graph G within a factor $\alpha < 1$, it is enough to approximate the maximum size of a clique in the tensor-product graph $G^{\otimes t}$ within a factor α^t . In this sense, the tensor-product operation amplifies the hardness of MAX CLIQUE.

In the context of UNIQUE GAMES, a natural candidate for a hardness-amplifying self-reduction is *parallel repetition*. (Using parallel repetition, Raz [Raz98] showed that the analog of the Unique Games Conjecture is true for LABEL COVER, a more general problem than UNIQUE GAMES.)

4.1. Main Result

Let \mathfrak{U} be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$. In this section, we lower bound the optimal value $\text{opt}(\mathfrak{U}^\ell)$ of the repeated game in terms of the value of the semidefinite relaxation $\text{SDP}(\mathfrak{U})$. (See Section 2.3 for the definition of parallel repetition.)

Theorem 4.1. *There exists an algorithm that given a parameter $\ell \in \mathbb{N}$ and a unique game \mathfrak{U} with semidefinite value $\text{sdp}(\mathfrak{U}) \geq 1 - \varepsilon$ and alphabet size k , computes in time $\text{poly}(n^\ell)$ an assignment x for the repeated game \mathfrak{U}^ℓ of value*

$$\mathfrak{U}^\ell(x) \geq 1 - O\left(\sqrt{\ell\varepsilon \log k}\right).$$

Before describing the proof, we record two consequences of this theorem. The first consequence is an approximation algorithm for parallel repeated unique games with the following approximation guarantee:

4. Parallel Repeated Unique Games

Corollary 4.2. *Suppose \mathbb{W} is an ℓ -fold repeated unique game with alphabet size K and $\text{opt}(\mathbb{W}) \geq 1 - \eta$. Then, using the algorithm in [Theorem 4.1](#), we can efficiently compute an assignment x for \mathbb{W} of value*

$$\mathbb{W}(x) \geq 1 - O\left(\sqrt{\eta \ell^{-1} \log K}\right). \quad (4.1)$$

For non-repeated games ($\ell = 1$), this algorithm achieves the best known approximation [[CMM06a](#)] (up to constants in the $O(\cdot)$ -notation). For sufficiently many repetitions ($\ell \gg 1$), the algorithm provides strictly better approximations than [[CMM06a](#)]. (Independent of the number of repetitions, the algorithm of [[CMM06a](#)] only guarantees to satisfy $1 - O(\sqrt{\eta \log K})$ of the constraints.)

Let us briefly sketch how [Theorem 4.1](#) implies the above approximation algorithm ([Corollary 4.2](#)). Suppose that $\mathbb{W} = \mathbb{U}^\ell$. Since \mathbb{U} has alphabet size k , the alphabet of \mathbb{W} has size $K = k^\ell$. By a result of Feige and Lovász [[FL92](#)], $\text{sdp}(\mathbb{W}) = \text{sdp}(\mathbb{U})^\ell$. Therefore, (in the relevant range of parameters) $\text{sdp}(\mathbb{U}) \geq 1 - O(\eta/\ell)$. It follows that the assignment obtained via [Theorem 4.1](#) satisfies the desired bound (4.1).

The second consequence of [Theorem 4.1](#) is a tight relation between the amortized value of a unique game and its semidefinite value.

Corollary 4.3.

$$\text{sdp}(\mathbb{U})^{O(\log k)} \leq \sup_{\ell \in \mathbb{N}} \text{opt}(\mathbb{U})^{1/\ell} \leq \text{sdp}(\mathbb{U}).$$

The proof of [Theorem 4.1](#) uses an intermediate relaxation, denoted $\text{SDP}_+(\mathbb{U})$ (defined in [Section 4.2](#)). The optimal value of this relaxation, $\text{sdp}_+(\mathbb{U})$, is sandwiched between $\text{opt}(\mathbb{U})$ and $\text{sdp}(\mathbb{U})$. Furthermore, the relaxation satisfies $\text{sdp}_+(\mathbb{U}^\ell) \geq \text{sdp}_+(\mathbb{U})^\ell$ for every $\ell \in \mathbb{N}$. (Note that both $\text{opt}(\mathbb{U})$ and $\text{sdp}(\mathbb{U})$ have the same property.) In [Section 4.2](#), we will show that the approximation guarantee of the intermediate relaxation $\text{SDP}_+(\mathbb{U})$ is independent of the alphabet size. (In contrast, the approximation guarantee of $\text{SDP}(\mathbb{U})$ degrades with growing alphabet size.)

Theorem 4.4. *Given a solution to $\text{SDP}_+(\mathbb{U})$ of value $1 - \gamma$, we can efficiently compute an assignment x for \mathbb{U} of value*

$$\mathbb{U}(x) \geq 1 - O(\sqrt{\gamma}).$$

4. Parallel Repeated Unique Games

We remark that it is not clear whether optimal solutions for the relaxation $\text{SDP}_+(\mathfrak{U})$ can be computed efficiently (in fact, if the Unique Games Conjecture is true, then this task is NP-hard). To prove [Theorem 4.1](#), we will approximate $\text{SDP}_+(\mathfrak{U})$ using the standard semidefinite relaxation $\text{SDP}(\mathfrak{U})$. (See [Section 4.3](#) for a proof of [Theorem 4.5](#).)

Theorem 4.5. *There exists a polynomial-time algorithm that given a unique game \mathfrak{U} with alphabet size k and semidefinite value $\text{sdp}(\mathfrak{U}) = 1 - \varepsilon$, computes a solution for $\text{SDP}_+(\mathfrak{U})$ of value at least $1 - O(\varepsilon \log k)$.*

Assuming the previous two theorems ([Theorem 4.4](#) and [Theorem 4.5](#)), we can prove [Theorem 4.1](#).

Proof of [Theorem 4.1](#). Since $\text{sdp}(\mathfrak{U}) = 1 - \varepsilon$, [Theorem 4.5](#) allows us to compute a solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - O(\varepsilon \log k)$. By taking tensor products, we obtain a solution for $\text{SDP}_+(\mathfrak{U}^\ell)$ of value $(1 - O(\varepsilon \log k))^\ell \geq 1 - O(\ell \log k)$ (see [Section 4.2](#)). Using [Theorem 4.4](#) (on the unique game \mathfrak{U}^ℓ), we can round this solution for $\text{SDP}_+(\mathfrak{U}^\ell)$ to an assignment x of value $\mathfrak{U}^\ell(x) \geq 1 - O(\ell \log k)^{1/2}$. \square

4.2. Rounding Nonnegative Vectors

In this section, we introduce the intermediate relaxation $\text{SDP}_+(\mathfrak{U})$ of a unique game \mathfrak{U} and prove [Theorem 4.4](#) (restated below).

Relaxation 4.6 (Nonnegative Relaxation for UNIQUE GAMES).

Given a unique game \mathfrak{U} with vertex set V and alphabet Σ , find a collection of nonnegative functions $\{f_{u,i}\}_{u \in V, i \in \Sigma} \subseteq L_2(\Omega)$ for some finite probability space (Ω, μ) so as to

$$\text{SDP}_+(\mathfrak{U}): \quad \text{maximize} \quad \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \langle f_{u,i}, f_{v,\pi(i)} \rangle \quad (4.2)$$

$$\text{subject to} \quad \sum_{i \in \Sigma} \|f_{u,i}\|^2 = 1 \quad (u \in V), \quad (4.3)$$

$$\langle f_{u,i}, f_{u,j} \rangle = 0 \quad (u \in V, i \neq j \in \Sigma). \quad (4.4)$$

An important property of the solutions for [Relaxation 4.6](#) is that for every vertex $u \in V$, the functions $f_{u,1}, \dots, f_{u,k}$ have disjoint support sets. In particular, for every point $\omega \in \Omega$, there exists at most one label such that $f_{u,i}(\omega)$ is positive. This property is the crucial difference from the relaxation $\text{SDP}(\mathfrak{U})$, which allows to prove [Theorem 4.4](#).

4. Parallel Repeated Unique Games

As in the previous section §4.1, let \mathfrak{U} be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$.

Theorem (Restatement of [Theorem 4.4](#)). *Given a solution to $\text{SDP}_+(\mathfrak{U})$ of value $1 - \gamma$, we can efficiently compute an assignment x for \mathfrak{U} of value*

$$\mathfrak{U}(x) \geq 1 - O(\sqrt{\gamma}).$$

Let $\{f_{u,i}\} \subseteq L_2(\Omega)$ be a solution for $\text{SDP}_+(\mathfrak{U})$ of value $\alpha = 1 - \gamma$. Let M be an upper bound on the values of the functions $f_{u,i}^2$. The running time of the algorithm for [Theorem 4.4](#) will depend polynomially on M . It is straightforward to ensure that M is polynomial in n , the number of vertices of \mathfrak{U} .

We remark that for the purpose of proving [Theorem 4.4](#), the reader can replace the function $\varphi(z) = (1 - z)/(1 + z)$ that appears in the next lemmas by the lower bound $\varphi(z) \geq 1 - 2z$. (See [Figure 4.1](#) for a comparison of the two functions.) The more precise bounds show that our rounding for $\text{SDP}_+(\mathfrak{U})$ gives non-trivial guarantees for all $\alpha > 0$ (whereas the bound stated in [Theorem 4.4](#) is non-trivial only for α close enough to 1).

The first step in the proof of [Theorem 4.4](#) is to construct a distribution over partial assignments $A: V \rightarrow \Sigma \cup \{\perp\}$ for \mathfrak{U} . If a vertex is not labeled by a partial assignment A (i.e., A assigns the null symbol \perp to this vertex), then all constraints containing this vertex count as unsatisfied. The following lemma lower bounds the probability that a constraints is satisfied (normalized by the fraction of labeled vertices) in terms of L_1 -distances of the functions $f_{u,i}^2$.

Lemma 4.7. *There is a polynomial-time randomized algorithm that samples a partial assignment $A: V \rightarrow \Sigma \cup \{\perp\}$ such that for every constraint (u, v, π) of \mathfrak{U} ,*

$$\mathbb{P}_A \left\{ A(u) = \pi(A(v)) \mid A(u) \neq \perp \vee A(v) \neq \perp \right\} \geq \varphi \left(\frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1 \right),$$

where $\varphi(z) = (1 - z)/(1 + z)$. Furthermore, $\mathbb{P}_A \{A(u) \neq \perp\} = 1/M$ for $u \in V$.

Proof. Sample $\omega \sim \Omega$ and $t \in [0, M]$. Put $A(u) = i$ if $f_{u,i}^2(\omega) > t$ and $A(u) = \perp$ if $f_{u,j}^2(\omega) \leq t$ for all $j \in \Sigma$. (Note that since $f_{u,1}, \dots, f_{u,k}$ have pairwise disjoint supports, there is at most one label $i \in \Sigma$ with $f_{u,i}^2(\omega) > t$.) The partial assignment

4. Parallel Repeated Unique Games

A satisfies the constraint (u, v, π) of \mathfrak{U} with probability

$$\begin{aligned} \mathbb{P}_A \{A(v) = \pi(A(u)) \in \Sigma\} &= \sum_i \mathbb{P}_A \{A(u) = i \wedge A(v) = \pi(i)\} \\ &= \sum_i \mathbb{E}_{\omega \sim \Omega} \frac{1}{M} \min \{f_{u,i}^2(\omega), f_{v,\pi(i)}^2(\omega)\} \\ &= \frac{1}{M} - \frac{1}{2M} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1. \end{aligned}$$

On the other hand, we can bound the probability that A assigns a label to one of the vertices u and v ,

$$\begin{aligned} \mathbb{P}_A \{A(u) \neq \perp \vee A(v) \neq \perp\} &= \mathbb{E}_{\omega \sim \Omega} \frac{1}{M} \max \left\{ \sum_i f_{u,i}^2(\omega), \sum_i f_{v,\pi(i)}^2(\omega) \right\} \\ &= \frac{1}{M} + \frac{1}{2M} \left\| \sum_i f_{u,i}^2 - \sum_i f_{v,\pi(i)}^2 \right\|_1 \leq \frac{1}{M} + \frac{1}{2M} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1. \end{aligned}$$

Combining these bounds,

$$\mathbb{P}_A \{A(u) = \pi(A(v)) \mid A(u) \neq \perp \vee A(v) \neq \perp\} \geq \frac{1 - \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1}{1 + \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1}. \quad \square$$

The next step in the proof of [Theorem 4.4](#) uses correlated sampling to obtain a distribution over (total) assignments $x \in \Sigma^V$ for \mathfrak{U} .

Lemma 4.8. *We can efficiently sample an assignment $x \in \Sigma$ for \mathfrak{U} such that*

$$\mathbb{E}_x \mathfrak{U}(x) \geq \varphi \left(\mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1 \right),$$

where $\varphi(z) = (1 - z)/(1 + z)$ as before.

Proof. Let $\{A^{(r)}\}_{r \in \mathbb{N}}$ be an infinite sequence of partial assignments independently sampled according to [Lemma 4.7](#). (We will argue at the end of the proof that a small number of samples suffices in expectation.) For every vertex $u \in V$, let $r(u)$ be the smallest index r such that $A^{(r)}(u) \neq \perp$. Assign the label $x_u = A^{(r(u))}(u)$ to vertex u . (Note that $r(u)$ is finite for every vertex with probability 1.) Consider a constraint (u, v, π) of \mathfrak{U} . We claim

$$\begin{aligned} \mathbb{P}_x \{x_v = \pi(x_u)\} &\geq \mathbb{P}_{x,r} \{x_v = \pi(x_u) \wedge r(v) = r(u)\} \\ &= \mathbb{P}_A \{A(v) = \pi(A(u)) \mid A(u) \neq \perp \vee A(v) \neq \perp\} \geq \varphi \left(\frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1 \right). \end{aligned}$$

4. Parallel Repeated Unique Games

For the equality, we use that the partial assignments $\{A^{(r)}\}$ are sampled independently. Therefore, we can condition, say, on the event $\min\{r(u), r(v)\} = 1$ without changing the probability. But this event is the same as the event $A^{(1)}(u) \neq \perp \vee A^{(1)}(v) \neq \perp$. The second inequality follows from [Lemma 4.7](#).

Using the convexity of the function $\varphi(x) = (1-x)/(1+x)$, we can bound the expected value of the assignment x ,

$$\begin{aligned} \mathbb{E}_x \mathfrak{U}(x) &= \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \mathbb{P}\{x_v = \pi(x_u)\} \geq \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \varphi\left(\frac{1}{2} \sum_i \|f_{u,i} - f_{v,\pi(i)}\|_1\right) \\ &\geq \varphi\left(\mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \frac{1}{2} \sum_i \|f_{u,i} - f_{v,\pi(i)}\|_1\right). \end{aligned}$$

To argue about the efficiency of the construction, we can estimate the expectation of $\max_{u \in V} r(u)$. (In this way, we bound the expected number of partial assignments we need to sample.) Since $\mathbb{P}_A\{A(u) \neq \perp\} = 1/M$ for every vertex $u \in V$, the expectation of $\max_{u \in V} r(u)$ is at most $M \cdot n$. \square

In the final step of the proof of [Theorem 4.4](#), we relate L_1 -distances of the functions $f_{u,i}^2$ to the objective value α of our solution for $\text{SDP}_+(\mathfrak{U})$.

Lemma 4.9.

$$\mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1 \leq (1-\alpha)^{1/2} (1+\alpha)^{1/2}.$$

Proof. By Cauchy–Schwarz,

$$\begin{aligned} \mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1 &= \mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \langle |f_{u,i} - f_{v,\pi(i)}|, |f_{u,i} + f_{v,\pi(i)}| \rangle \\ &\leq \left(\mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \|f_{u,i} - f_{v,\pi(i)}\|^2 \right)^{1/2} \left(\mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \|f_{u,i} + f_{v,\pi(i)}\|^2 \right)^{1/2}. \end{aligned}$$

This bound proves the lemma, because $\mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \|f_{u,i} - f_{v,\pi(i)}\|^2 = 1 - \alpha$ and $\mathbb{E}_{(u,v,\pi) \sim G} \frac{1}{2} \sum_i \|f_{u,i} + f_{v,\pi(i)}\|^2 = 1 + \alpha$. \square

[Theorem 4.4](#) follows by combining the previous lemmas ([Lemma 4.8](#) and [Lemma 4.9](#)).

Proof of [Theorem 4.4](#). Using [Lemma 4.8](#), we can sample an assignment x with

$$\mathbb{E}_x \mathfrak{U}(x) \geq \varphi\left(\mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1\right).$$

4. Parallel Repeated Unique Games




Figure 4.1.: Comparison of $1 - 2z$ and $(1 - z)/(1 + z)$

According to [Lemma 4.9](#),

$$\mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}} \frac{1}{2} \sum_i \|f_{u,i}^2 - f_{v,\pi(i)}^2\|_1 \leq (1 - \alpha)^{1/2} (1 + \alpha)^{1/2}$$

Since $\alpha = 1 - \gamma$, we have $(1 - \alpha)^{1/2} (1 + \alpha)^{1/2} \leq \sqrt{2\gamma}$. By monotonicity of $\varphi(z) = (1 - z)/(1 + z)$, we get $\mathbb{E}_x \mathfrak{U}(x) \geq \varphi(\sqrt{2\gamma}) \geq 1 - 2\sqrt{2\gamma}$. Using standard arguments (independently repeated trials), we can efficiently obtain an assignment x with value $\mathfrak{U}(x) \geq 1 - O(\sqrt{\gamma})$ with high probability. \square

4.3. From Arbitrary to Nonnegative Vectors

As in the previous sections, let \mathfrak{U} be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$.

In this section, we show how to obtain approximate solutions for the non-negative relaxation $\text{SDP}_+(\mathfrak{U})$ from an optimal solution to the semidefinite relaxation $\text{SDP}(\mathfrak{U})$.

Theorem (Restatement of [Theorem 4.5](#)). *There exists a polynomial-time algorithm that given a unique game \mathfrak{U} with alphabet size k and semidefinite value $\text{sdp}(\mathfrak{U}) = 1 - \varepsilon$, computes a solution for $\text{SDP}_+(\mathfrak{U})$ of value at least $1 - O(\varepsilon \log k)$.*

Before describing the general construction, we first illustrate the theorem by proving the following special case.

4. Parallel Repeated Unique Games

Proof of Theorem 4.5 for $k = 2$. Let $\{\mathbf{u}_i\}_{u \in V, i \in \{1,2\}} \subseteq \mathbb{R}^d$ be solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$. Consider an orthonormal basis e_1, \dots, e_d of \mathbb{R}^d . Let (Ω, μ) be the uniform distribution over $\{1, \dots, d\}$. For $u \in V$ and $i \in \{1, 2\}$, define a function $f_{u,i} \in L_2(\Omega)$ as follows

$$f_{u,i}(r) = \begin{cases} \sqrt{d} \cdot (\langle e_r, \mathbf{u}_i \rangle - \langle e_r, \mathbf{u}_{3-i} \rangle) & \text{if } \langle e_r, \mathbf{u}_i \rangle > \langle e_r, \mathbf{u}_{3-i} \rangle, \\ 0 & \text{otherwise.} \end{cases}$$

(Note that $3 - i$ is just the label in $\{1, 2\}$ distinct from i .) The construction ensures that $f_{u,1}$ and $f_{u,2}$ have disjoint support and are nonnegative for every vertex $u \in V$. Since \mathbf{u}_1 and \mathbf{u}_2 are orthogonal, the construction also preserves the total mass

$$\|f_{u,1}\|^2 + \|f_{u,2}\|^2 = \sum_{r=1}^d (\langle e_r, \mathbf{u}_1 \rangle - \langle e_r, \mathbf{u}_2 \rangle)^2 = \|\mathbf{u}_1\|^2 + \|\mathbf{u}_2\|^2 - \underbrace{2\langle \mathbf{u}_1, \mathbf{u}_2 \rangle}_0 = 1.$$

Next, we show that the construction approximately preserves distances. Consider a constraint (u, v, π) of \mathfrak{U} . Put $\mathbf{u} = \mathbf{u}_1 - \mathbf{u}_2$ and $\mathbf{v} = \mathbf{v}_{\pi(1)} - \mathbf{v}_{\pi(2)}$. Let I^+ be the set of indices r such that $\text{sign}\langle e_r, \mathbf{u} \rangle = \text{sign}\langle e_r, \mathbf{v} \rangle$, and let I^- be the complement of I^+ . Then,

$$\sum_{i \in \{1,2\}} \|f_{u,i} - f_{v,\pi(i)}\|^2 = \sum_{r \in I^+} \langle e_r, \mathbf{u} - \mathbf{v} \rangle^2 + \sum_{r \in I^-} (\langle e_r, \mathbf{u} \rangle^2 + \langle e_r, \mathbf{v} \rangle^2) \leq \|\mathbf{u} - \mathbf{v}\|^2.$$

(In the last step, we used that $\langle e_r, \mathbf{u} \rangle \langle e_r, \mathbf{v} \rangle < 0$ for $r \in I^-$.) By the triangle inequality, $\|\mathbf{u} - \mathbf{v}\| \leq \sum_{i \in \{1,2\}} \|\mathbf{u}_i - \mathbf{v}_{\pi(i)}\|$. We conclude that $\sum_{i \in \{1,2\}} \|f_{u,i} - f_{v,\pi(i)}\|^2 \leq 2 \sum_{i \in \{1,2\}} \|\mathbf{u}_i - \mathbf{v}_{\pi(i)}\|^2$, which implies that the value of the constructed solution $\{f_{u,i}\}$ for $\text{SDP}_+(\mathfrak{U})$ is at least $1 - 2\varepsilon$. \square

For larger alphabets, the proof of Theorem 4.5 proceeds in two steps. In the first step, we construct a nearly feasible solution for $\text{SDP}_+(\mathfrak{U})$. This step loses the factor $O(\log k)$ that appears in the bound of Theorem 4.5. In the second step, we show how to repair a nearly feasible solution of $\text{SDP}_+(\mathfrak{U})$ while approximately preserving the objective value. The construction for this step is quite delicate, but resembles the simple construction for $k = 2$ to some extent.

4. Parallel Repeated Unique Games

4.3.1. Constructing Nearly Feasible Solutions

We say a collection of nonnegative functions $\{f_{u,i}\}_{u \in V, i \in \Sigma} \subseteq L_2(\Omega)$ is γ -infeasible for $\text{SDP}_+(\mathfrak{U})$ if for every vertex $u \in V$,

$$\sum_i \|f_{u,i}\|^2 = 1, \quad (4.5)$$

$$\sum_{i \neq j} \langle f_{u,i}, f_{u,j} \rangle \leq \gamma. \quad (4.6)$$

In the following, we show that if $\text{sdp}(\mathfrak{U}) = 1 - \varepsilon$, then one can efficiently compute a $1/10$ -infeasible solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - O(\varepsilon \log k)$. (The bound stated in the theorem is more general.)

Theorem 4.10. *If $\text{sdp}(\mathfrak{U}) = 1 - \varepsilon$, then for every $t \geq 1$ we can efficiently compute a $k \cdot 2^{-t}$ -infeasible solution for $\text{SDP}_+(\mathfrak{U})$ of value at least $1 - O(t) \cdot \varepsilon$.*

Let $\{\mathbf{u}_i\}_{u \in V, i \in \Sigma}$ be a solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$. The proof of [Theorem 4.10](#) has two steps. First, we construct nonnegative functions $f_{u,i}$ over the space \mathbb{R}^d (equipped with the d -dimensional Lebesgue measure). These functions will have all the desired properties (nearly feasible and high objective value). Then, we discretize the space and obtain functions defined on a finite probability space, while approximately preserving the desired properties. (The second step is quite tedious and not very interesting. We omit its proof.)

The crux of the proof of [Theorem 4.10](#) is the following mapping from \mathbb{R}^d to nonnegative functions in $L_2(\mathbb{R}^d)$,

$$M_\sigma(x) \stackrel{\text{def}}{=} \|x\| \sqrt{T_{\bar{x}} \phi_\sigma}. \quad (4.7)$$

Here, $\bar{x} = \|x\|^{-1}x$ is the unit vector in direction x , T_h is the translation operator on $L_2(\mathbb{R}^d)$, so that $T_h f(x) = f(x - h)$, and $\phi_\sigma \in L_2(\mathbb{R}^d)$ is the density of the d -dimensional Gaussian distribution $N(0, \sigma^2)^d$ with respect to the Lebesgue measure λ^d on \mathbb{R}^d ,

$$\phi_\sigma(x) \stackrel{\text{def}}{=} \frac{1}{(\sigma\sqrt{2\pi})^d} e^{-\|x\|^2/2\sigma^2}.$$

Since $\int \phi_\sigma d\lambda^d = 1$ (and the measure λ^d is invariant under translation), the mapping M_σ is norm-preserving, so that $\|M_\sigma(x)\| = \|x\|$ for every $x \in \mathbb{R}^d$.

4. Parallel Repeated Unique Games

The mapping M_σ approximately preserves inner products (and thereby distances). The trade-off between the approximation of two scales will be relevant. On the one hand, orthogonal unit vectors are mapped to unit vectors with inner product $2^{-\Omega(\sigma^{-2})}$. On the other hand, the squared distance of vectors is stretched by at most a factor σ^{-2} .

Lemma 4.11. *For any two vectors $x, y \in \mathbb{R}^d$,*

$$\langle M_\sigma(x), M_\sigma(y) \rangle = \|x\| \|y\| e^{-\|\bar{x}-\bar{y}\|^2/8\sigma^2}. \quad (4.8)$$

In particular, $\|M_\sigma(x) - M_\sigma(y)\|^2 \leq \|x - y\|^2/4\sigma^2$ (for $\sigma \leq 1/2$).

Proof. The following identity implies (4.8) (by integration)

$$\sqrt{T_{\bar{x}}\phi_\sigma \cdot T_{\bar{y}}\phi_\sigma} = e^{-\|\bar{x}-\bar{y}\|^2/8\sigma^2} T_{\frac{1}{2}(\bar{x}+\bar{y})}\phi_\sigma.$$

To verify the bound $\|M_\sigma(x) - M_\sigma(y)\|^2 \leq \|x - y\|^2/4\sigma^2$, we first show a corresponding bound for the unit vectors \bar{x} and \bar{y} ,

$$\|M_\sigma(\bar{x}) - M_\sigma(\bar{y})\|^2 = 2 - 2e^{-\|\bar{x}-\bar{y}\|^2/8\sigma^2} \leq \|\bar{x} - \bar{y}\|^2/4\sigma^2.$$

Using that the mapping M_σ is norm-preserving, we can show the desired bound

$$\begin{aligned} \|M_\sigma(x) - M_\sigma(y)\|^2 &= (\|x\| - \|y\|)^2 + \|x\| \|y\| \cdot \|M_\sigma(\bar{x}) - M_\sigma(\bar{y})\|^2 \\ &\leq (\|x\| - \|y\|)^2/4\sigma^2 + \|x\| \|y\| \cdot \|M_\sigma(\bar{x}) - M_\sigma(\bar{y})\|^2/4\sigma^2 = \|x - y\|^2/4\sigma^2. \end{aligned}$$

(The inequality uses the assumption $\sigma \leq 1/2$.) □

To prove [Theorem 4.10](#), we map the vectors \mathbf{u}_i of the solution for $\text{SDP}(\mathfrak{U})$ to nonnegative functions $f_{u,i} := M_\sigma(\mathbf{u}_i)$ (with $\sigma^2 \approx 1/t$). [Lemma 4.11](#) shows that the functions $f_{u,i}$ form a $k \cdot 2^{-t}$ -infeasible solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - O(t) \cdot \varepsilon$.

Proof of [Theorem 4.10](#). Let $\{\mathbf{u}_i\}_{u \in V, i \in \Sigma}$ be a solution for $\text{SDP}(\mathfrak{U})$ of value $1 - \varepsilon$. Put $f_{u,i} = M_\sigma(\mathbf{u}_i)$ (we fix the parameter σ later). First, we verify that $\{f_{u,i}\}_{u \in V, i \in \Sigma}$ is $k \cdot e^{-1/4\sigma^2}$ -infeasible. Since M_σ is norm-preserving, we have

4. Parallel Repeated Unique Games

$\sum_{i \in \Sigma} \|f_{u,i}\|^2 = 1$ for every vertex $u \in V$. On the other hand, since the vectors $\{\mathbf{u}_i\}_{i \in \Sigma}$ are pairwise orthogonal, [Lemma 4.11](#) shows that

$$\sum_{i \neq j} \langle f_{u,i}, f_{u,j} \rangle = e^{-1/4\sigma^2} \left(\left(\sum_i \|\mathbf{u}_i\| \right)^2 - 1 \right) \leq k \cdot e^{-1/4\sigma^2}.$$

(Here, we used Cauchy–Schwarz to bound $\sum_i \|\mathbf{u}_i\|$.) The objective value of the functions $f_{u,i}$ is at least $1 - \varepsilon/4\sigma^2$, because [Lemma 4.11](#) implies that for every constraint (u, v, π) of \mathfrak{U} ,

$$\sum_i \|f_{u,i} - f_{v,\pi(i)}\|^2 \leq \sum_i \|\mathbf{u}_i - \mathbf{v}_{\pi(i)}\|^2 / 4\sigma^2.$$

If we choose $\sigma^2 = 1/(4t \ln 2)$, we obtain a $k \cdot 2^{-t}$ -infeasible solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - O(t)\varepsilon$, as desired. \square

4.3.2. Repairing Nearly Feasible Solutions

In the following, we show how to repair nearly feasible solutions for $\text{SDP}_+(\mathfrak{U})$ while approximately preserving the objective value.

Theorem 4.12. *Given a $1/10$ -infeasible solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - \varepsilon'$, we can efficiently compute a (feasible) solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - O(\varepsilon')$.*

The proof of [Theorem 4.12](#) is quite technical and basically boils down to a careful truncation procedure (which we refer to as “soft truncation”, see [Lemma 4.13](#)). A significantly simpler truncation procedure shows that given a γ -infeasible solution of value $1 - \varepsilon'$, one can efficiently compute a feasible solution of value $1 - O(\varepsilon' + \gamma)$. Due to this additive decrease of the objective value, this simpler truncation is not enough to show [Theorem 4.5](#) ($\text{SDP}(\mathfrak{U})$ vs. $\text{SDP}_+(\mathfrak{U})$). However, it suffices to establish the following slightly weaker result: If $\text{SDP}(\mathfrak{U}) = 1 - \varepsilon$, we can efficiently compute a solution for $\text{SDP}_+(\mathfrak{U})$ of value $1 - O(\varepsilon \log(k/\varepsilon))$.

[Theorem 4.12](#) follows from the following lemma (which asserts the existence of a mapping that we can use to transform a $1/10$ -feasible solution for $\text{SDP}_+(\mathfrak{U})$ to a feasible solution with roughly the same objective value).

Lemma 4.13 (Soft Truncation). *There exists an efficiently computable mapping $Q: L_2(\Omega)^k \rightarrow L_2(\Omega)^k$ with the following properties: Let f_1, \dots, f_k and g_1, \dots, g_k be*

4. Parallel Repeated Unique Games




Figure 4.2.: Illustration of the construction for [Theorem 4.10](#). With each vector u_1, v_1, u_2, v_2 we associate a Gaussian distribution centered at that vector (depicted as disc in the figure). The infeasibility of the solution for the SDP_+ relaxation corresponds to the overlap of discs with the same color.

4. Parallel Repeated Unique Games

nonnegative functions in $L_2(\Omega)$ such that $\sum_i \|f_i\|^2 = \sum_i \|g_i\|^2 = 1$ and $\sum_{i \neq j} \langle f_i, f_j \rangle + \sum_{i \neq j} \langle g_i, g_j \rangle \leq \gamma$. Suppose $(f'_1, \dots, f'_k) = Q(f_1, \dots, f_k)$ and $(g'_1, \dots, g'_k) = Q(g_1, \dots, g_k)$. Then, for every permutation π of $[k]$,

$$\sum_i \|f'_i - g'_{\pi(i)}\|^2 \leq \frac{32}{1-4\gamma} \sum_i \|f_i - g_{\pi(i)}\|^2.$$

Furthermore, $\sum_i \|f'_i\|^2 = \sum_i \|g'_i\|^2 = 1$ and $\text{supp}(f'_i) \cap \text{supp}(f'_j) = \text{supp}(g'_i) \cap \text{supp}(g'_j) = \emptyset$ for all $i \neq j$.

Proof. The mapping Q is the composition of two mappings $Q^{(1)}$ and $Q^{(2)}$. The first mapping $Q^{(1)}$ ensures that the output functions have disjoint support,

$$Q^{(1)}: L_2(\Omega)^k \rightarrow L_2(\Omega)^k, \quad (f_1, \dots, f_k) \mapsto (f'_1, \dots, f'_k), \quad (4.9)$$

$$f'_i(x) = \begin{cases} f_i(x) - \max_{j \neq i} f_j(x) & \text{if } f_i(x) > \max_{j \neq i} f_j(x), \\ 0 & \text{otherwise.} \end{cases} \quad (4.10)$$

The purpose of the second mapping $Q^{(2)}$ is to renormalize the function (so that the squared L_2 norms sum to 1),

$$Q^{(2)}: L_2(\Omega)^k \rightarrow L_2(\Omega)^k, \quad (4.11)$$

$$(f_1, \dots, f_k) \mapsto \left(\frac{1}{\lambda} f_1, \dots, \frac{1}{\lambda} f_k \right), \quad (4.12)$$

$$\text{where } \lambda^2 = \sum_i \|f_i\|^2. \quad (4.13)$$

Let f_1, \dots, f_k and g_1, \dots, g_k be nonnegative functions in $L_2(\Omega)$ satisfying the conditions of the current lemma, i.e.,

$$- \sum_{i \neq j} \langle f_i, f_j \rangle + \sum_{i \neq j} \langle g_i, g_j \rangle \leq \gamma$$

$$- \sum_i \|f_i\|^2 = \sum_i \|g_i\|^2 = 1.$$

Let $(f_1^{(1)}, \dots, f_k^{(1)}) = Q^{(1)}(f_1, \dots, f_k)$ and $(f_1^{(2)}, \dots, f_k^{(2)}) = Q^{(2)}(f_1^{(1)}, \dots, f_k^{(1)})$. Similarly, let $(g_1^{(1)}, \dots, g_k^{(1)}) = Q^{(1)}(g_1, \dots, g_k)$ and $(g_1^{(2)}, \dots, g_k^{(2)}) = Q^{(2)}(g_1^{(1)}, \dots, g_k^{(1)})$. Next, we establish two claims about the effects of the mappings $Q^{(1)}$ and $Q^{(2)}$ ([Claim 4.14](#) and [Claim 4.15](#)). The current lemma ([Lemma 4.13](#)) follows by combining these two claims.

Claim 4.14 (Properties of $Q^{(1)}$). 1. For every permutation π of $[k]$,

$$\sum_i \|f_i^{(1)} - g'_{\pi(i)}\|^2 \leq 8 \sum_i \|f_i - g_{\pi(i)}\|^2.$$

4. Parallel Repeated Unique Games

2. For all $i \neq j$,

$$\text{supp}(f_i^{(1)}) \cap \text{supp}(f_j^{(1)}) = \emptyset.$$

3.

$$1 \geq \sum_i \|f_i^{(1)}\|^2 \geq 1 - 2\gamma.$$

Proof. **Item 2** holds, since by construction $\text{supp}(f_i^{(1)}) = \{x \mid f_i(x) > \max_{j \neq i} f_j(x)\}$. To prove **Item 3**, we observe that $f_i^{(1)}(x)^2 > f_i(x)^2 - 2 \sum_{j \neq i} f_i(x) f_j(x)$ and therefore as desired

$$\sum_i \|f_i^{(1)}\|^2 \geq \sum_i \|f_i\|^2 - 2 \sum_{i \neq j} \langle f_i, f_j \rangle \geq 1 - 2\gamma.$$

To prove **Item 1**, we will show that for every $x \in \Omega$,

$$\sum_i \left(f_i^{(1)}(x) - g_{\pi(i)}^{(1)}(x) \right)^2 \leq 8 \sum_i \left(f_i(x) - g_{\pi(i)}(x) \right)^2. \quad (4.14)$$

Since $Q^{(1)}$ is invariant under permutation of its inputs, we may assume π is the identity permutation. At this point, we can verify (4.14) by an exhaustive case distinction. Fix $x \in \Omega$. Let i_f be the index i that maximizes $f_i(x)$. (We may assume the maximizer is unique.) Let j_f be the index such that $f_{j_f}(x) = \max_{j \neq i_f} f_j(x)$. Similarly, define i_g and j_g such that $g_{i_g}(x) = \max_i g_i(x)$ and $g_{j_g}(x) = \max_{j \neq i_g} g_j(x)$. We may assume that $i_f = 1$ and $j_f = 2$. Furthermore, we may assume $i_g, j_g \in \{1, 2, 3, 4\}$. Notice that the sum on the left-hand side of (4.14) has at most two non-zero terms (corresponding to the indices $i \in \{i_f, i_g\} \subseteq \{1, \dots, 4\}$). Hence, to verify (4.14), it is enough to show

$$\max_{i \in \{1, \dots, 4\}} \left| f_i^{(1)}(x) - g_i^{(1)}(x) \right| \leq 4 \max_{i \in \{1, \dots, 4\}} \left| f_i(x) - g_i(x) \right|. \quad (4.15)$$

Put $\varepsilon = \max_{i \in \{1, \dots, 4\}} |f_i(x) - g_i(x)|$. Let $q_i(a_1, \dots, a_4) := \max\{a_i - \max_{j \neq i} a_j, 0\}$. Note that $f_i^{(1)}(x) = q_i(f_1(x), \dots, f_4(x))$ and $g_i^{(1)}(x) = q_i(g_1(x), \dots, g_4(x))$. The functions q_i are 1-Lipschitz in each of their four inputs. It follows as desired that for every $i \in \{1, \dots, 4\}$,

$$\left| f_i^{(1)}(x) - g_i^{(1)}(x) \right| = \left| q_i(f_1(x), \dots, f_4(x)) - q_i(g_1(x), \dots, g_4(x)) \right| \leq 4\varepsilon. \quad \square$$

Claim 4.15 (Properties of $Q^{(2)}$). 1. For every permutation π of $[k]$,

$$\sum_i \|f_i^{(2)} - g_{\pi(i)}^{(2)}\|^2 \leq \frac{4}{1-4\gamma} \sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2.$$

4. Parallel Repeated Unique Games

2. For all $i \in \Sigma$,

$$\text{supp}(f_i^{(2)}) = \text{supp}(f_i^{(1)}).$$

3.

$$\sum_i \|f_i^{(2)}\|^2 = 1.$$

Proof. Again [Item 2](#) and [Item 3](#) follow immediately by definition of the mapping $Q^{(2)}$. To prove [Item 1](#), let $\lambda_f, \lambda_g > 0$ be the multipliers such that $f_i^{(2)} = f_i^{(1)}/\lambda_f$ and $g_i^{(2)} = g_i^{(1)}/\lambda_g$ for all $i \in [k]$. [Item 1](#) of [Claim 4.14](#) shows that λ_f^2 and λ_g^2 lie in the interval $[1 - 2\gamma, 1]$. We estimate the distances between $f_i^{(2)}$ and $g_{\pi(i)}^{(2)}$ as follows,

$$\begin{aligned} \sum_i \|f_i^{(2)} - g_{\pi(i)}^{(2)}\|^2 &= \sum_i \left\| \frac{1}{\lambda_f} (f_i^{(1)} - g_{\pi(i)}^{(1)}) + \left(\frac{1}{\lambda_f} - \frac{1}{\lambda_g} \right) g_i^{(1)} \right\|^2 \\ &\leq \frac{2}{\lambda_f^2} \sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2 + 2 \left(\frac{1}{\lambda_f} - \frac{1}{\lambda_g} \right)^2 \sum_i \|g_i^{(1)}\|^2 \\ &\quad \text{(using } \|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2\text{)} \\ &\leq \frac{2}{1-2\gamma} \sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2 + 2 \left(\frac{1}{\lambda_f} - \frac{1}{\lambda_g} \right)^2 \quad \left(\text{using } \sum_i \|g_i^{(1)}\|^2 \leq 1 \right). \end{aligned}$$

It remains to upper bound the second term on the right-hand side, $(1/\lambda_f - 1/\lambda_g)^2$. Since the function $x \mapsto 1/x$ is $1/a^2$ -Lipschitz on an interval of the form $[a, \infty)$, we have

$$\begin{aligned} \left| \frac{1}{\lambda_f} - \frac{1}{\lambda_g} \right| &\leq \frac{1}{1-2\gamma} |\lambda_f - \lambda_g| \\ &= \frac{1}{1-2\gamma} \left| \left(\sum_i \|f_i^{(1)}\|^2 \right)^{1/2} - \left(\sum_i \|g_{\pi(i)}^{(1)}\|^2 \right)^{1/2} \right| \\ &\leq \frac{1}{1-2\gamma} \left(\sum_i \left(\|f_i^{(1)}\| - \|g_{\pi(i)}^{(1)}\| \right)^2 \right)^{1/2} \quad \text{(using triangle inequality)} \\ &\leq \frac{1}{1-2\gamma} \left(\sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2 \right)^{1/2} \quad \text{(using triangle inequality)}. \end{aligned}$$

Combining the previous two estimates, we get as desired

$$\sum_i \|f_i^{(2)} - g_{\pi(i)}^{(2)}\|^2 \leq \left(\frac{2}{1-2\gamma} + \frac{2}{(1-2\gamma)^2} \right) \sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2. \quad \square$$

\square

4. Parallel Repeated Unique Games

4.4. Notes

The material presented in this chapter is based on the paper “Rounding Parallel Repetitions of Unique Games” [BHH⁺08], joint with Boaz Barak, Moritz Hardt, Ishay Haviv, Anup Rao, and Oded Regev, and on the paper “Improved Rounding for Parallel-Repeated Unique Games” [Ste10b]. A preliminary version of the first paper appeared at FOCS 2008. The second paper appeared in preliminary form at RANDOM 2010.

Related Works

The techniques in this chapter are natural generalizations of ideas in Raz’s counterexample to strong parallel repetition [Raz08].

Before discussing the connections, we first describe Raz’s result and his construction. He showed that for every $\varepsilon > 0$, there exists a unique game \mathfrak{U} (MAX CUT on an odd cycle of length $1/\varepsilon$) that satisfies $\text{opt}(\mathfrak{U}) \leq 1 - \varepsilon$ and $\text{opt}(\mathfrak{U}^\ell) \geq 1 - O(\sqrt{\ell}\varepsilon)$. To construct the assignment for the repeated game \mathfrak{U}^ℓ , he assigned to every vertex of \mathfrak{U} a distribution over assignments for the unique game \mathfrak{U} . To every vertex of the repeated game \mathfrak{U}^ℓ , he assigned the corresponding product distribution. Finally, he obtained an assignment for \mathfrak{U}^ℓ using Holenstein’s correlated sampling protocol [Hol09].

In retrospect, we can interpret Raz’s construction in terms of techniques presented in this chapter. The distributions he assigns to the vertices of the unique game \mathfrak{U} correspond to a solution to the intermediate relaxation $\text{SDP}_+(\mathfrak{U})$. His construction shows that $\text{sdp}_+(\mathfrak{U}) \geq 1 - O(\varepsilon^2)$. (In fact, his construction roughly corresponds to the proof of the special case of [Theorem 4.5](#) for $k = 2$ that we presented in the beginning of [§4.3](#).) The distributions he considers for the repeated game \mathfrak{U}^ℓ roughly correspond to a (product) solution for the relaxation $\text{SDP}_+(\mathfrak{U}^\ell)$, demonstrating that $\text{sdp}_+(\mathfrak{U}^\ell) \geq 1 - O(\ell\varepsilon^2)$. Finally, he rounds these distributions to an assignment for \mathfrak{U}^ℓ , roughly corresponding to the proof [Theorem 4.4](#) presented in [§4.2](#).

In this light, the main technical contribution of the works [BHH⁺08] and [Ste10b] is the general construction of solutions to the intermediate relaxation $\text{SDP}_+(\mathfrak{U})$ based on a solution to the relaxation $\text{SDP}(\mathfrak{U})$. (Raz’s construction of a solution to $\text{SDP}_+(\mathfrak{U})$ relies on the specific structure of the unique games he considers.)

We note that Kindler et al. [KORW08] extended Raz’s construction to the continuous case, leading to a construction of high-dimensional foams with

4. Parallel Repeated Unique Games

surprisingly small surface area. Alon and Klartag [AK09] unified the results of Raz [Raz08] and Kindler et al. [KORW08] and generalized them further.

Correlated Sampling

In the proof of [Theorem 4.4](#) (see [§4.2](#)), we used a technique called “correlated sampling” to construct an assignment for a unique game from a suitable distribution over partial assignments. This technique was first used by Broder [Bro97] for sketching sets. In the context of rounding relaxations, this technique was introduced by Kleinberg and Tardos [KT02].

Later, Holenstein [Hol09] used this technique in his proof of the Parallel Repetition Theorem. In the context of rounding relaxations of UNIQUE GAMES, this technique was used by [CMM06b] and [KRT08].

5. Subexponential Approximation Algorithms

Impagliazzo, Paturi, and Zane [IPZ01] showed that many NP-hard problems require strongly exponential running time — time $2^{\Omega(n)}$ — assuming 3-SAT requires strongly exponential time. Similarly, efficient PCP construction (the latest one by Moshkovitz and Raz [MR08]) show that, under the same assumption, achieving non-trivial approximation guarantees¹ for problems like MAX 3-LIN and LABEL COVER requires running time $2^{\Omega(n^{1-o(1)})}$.

In this chapter, we show that for UNIQUE GAMES, non-trivial approximation guarantees can be achieved in subexponential time, demonstrating that UNIQUE GAMES is significantly easier to approximate than problems like MAX 3-LIN and LABEL COVER. Concretely, we give an algorithm for UNIQUE GAMES that achieves a $(1 - \varepsilon, \varepsilon)$ -approximation in time exponential in $k \cdot n^{O(\varepsilon^{1/3})}$, where n is the number of vertices and k is the alphabet size.

This algorithm for UNIQUE GAMES is inspired by a (similar, but somewhat simpler) algorithm for SMALL-SET EXPANSION with the following approximation guarantee for every $\beta > 0$: Given a graph G containing a vertex set S with volume $\mu(S) = \delta$ and expansion $\Phi(S) = \varepsilon$, the algorithm finds in time $\exp(n^\beta/\delta)$ a vertex set S' with volume close to δ and expansion $\Phi(S') \leq O(\sqrt{\varepsilon/\beta})$.

5.1. Main Results

In this section, we describe subexponential approximation algorithms for SMALL-SET EXPANSION and UNIQUE GAMES. In the following, we use $\exp(f)$ for a function f to denote a function that is bounded from above by $2^{C \cdot f}$ for some absolute constant $C \in \mathbb{N}$.

Our algorithm for SMALL-SET EXPANSION achieves the following trade-off

¹ Here (as for all of this thesis), we are interested in approximation guarantees independent of the instance size. For MAX 3-LIN, a $(1 - \varepsilon, 1/2 + \varepsilon)$ -approximation for some constant $\varepsilon > 0$ is non-trivial. For LABEL COVER, a $(1, \varepsilon)$ -approximation for some constant ε is non-trivial.

5. Subexponential Approximation Algorithms

between running time and approximation guarantee. (See §5.1.1 for the proof.)

Theorem 5.1 (Subexponential Algorithm for SMALL-SET EXPANSION).

There exists an algorithm that given a graph G on n vertices containing a vertex set with volume at most δ and expansion at most ε , computes in time $\exp(n^\beta/\delta)$ a vertex set S with volume $\mu(S) \leq (1 + \gamma)\delta$ and expansion $\Phi(S) \leq O(\sqrt{\varepsilon/\beta\gamma^9})$. Here, β is a parameter of the algorithm that can be chosen as small as $\log \log n / \log n$.

We remark that for $\beta, \gamma \geq \Omega(1)$ (say $\beta = \gamma = 0.001$), the approximation guarantee of our algorithm for SMALL-SET EXPANSION matches up to constant factors the approximation guarantee of Cheeger's bound for SPARSEST CUT.

Our algorithm for UNIQUE GAMES achieves the following trade-off between running time and approximation guarantee.

Theorem 5.2 (Subexponential Algorithm for UNIQUE GAMES).

There exists an algorithm that given a unique game \mathfrak{U} with n vertices, alphabet size k , and optimal value $\text{opt}(\mathfrak{U}) \geq 1 - \varepsilon$, computes in time $\exp(kn^\beta)$ an assignment x of value $\mathfrak{U}(x) \geq 1 - O(\sqrt{\varepsilon/\beta^3})$. Here, β is a parameter of the algorithm that can be chosen as small as $\log \log n / \log n$.

An immediate consequence of the above theorem is that we can compute in time exponential in $n^{\tilde{O}(\varepsilon^{1/3})}$ an assignment that satisfies a constant fraction (say, $1/2$) of the constraints of \mathfrak{U} (where \mathfrak{U} is as in the statement of the theorem). See §5.1.2 for a proof of this special case of Theorem 5.2. (Section 5.5 contains the proof of the general case of Theorem 5.2.)

Both algorithms rely on the notion of threshold rank of a matrix (and the related notion of soft threshold rank). For a symmetric matrix A and $\tau \geq 0$, we define the *threshold rank* of A at τ , denoted $\text{rank}_\tau(A)$, as the number of eigenvalues of A larger than τ in absolute value. The threshold rank at 0 coincides with the usual notion of rank. In general, the threshold rank at τ satisfies

$$\text{rank}_\tau(A) = \min_{\|B\| \leq \tau} \text{rank}(A - B),$$

where B is a symmetric matrix with the same dimension as A and $\|B\|$ denotes the spectral norm of B (the largest eigenvalue in absolute value).

The algorithm for SMALL-SET EXPANSION is significantly simpler than the algorithm for UNIQUE GAMES, but it contains the main ideas.

5. Subexponential Approximation Algorithms

5.1.1. Small Set Expansion

Let G be a graph with vertex set $V = [n]$. For simplicity, assume that G is regular. We identify G with its stochastic adjacency matrix. For $\eta \in [0, 1)$, let U_η denote the subspace of $L_2(V)$ spanned by the eigenfunctions of G with eigenvalue larger than $1 - \eta$. Note that $\dim U_\eta = \text{rank}_{1-\eta/2}(\frac{1}{2}I + \frac{1}{2}G) \leq \text{rank}_{1-\eta}(G)$.

The next lemma formalizes the following idea: Suppose G contains a set with low expansion. Then, given the projection of the indicator function of this set onto U_η (for appropriate η), we can reconstruct a set with roughly the same volume and expansion. (For the reconstruction, we just try all level sets of the projection.) The lemma also asserts that the reconstruction is robust, i.e., it succeeds even given a function in U_η close to the true projection.

Lemma 5.3. *Let $\varepsilon, \delta, \eta > 0$ such that ε/η is sufficiently small. Suppose the graph G contains a vertex set with volume at most δ and expansion at most ε . Then, every $\varepsilon/100$ -net of the subspace U_η contains a function with a level set $S \subseteq V$ satisfying $\mu(S) \leq 1.1\delta$ and $\Phi(S) \leq O(\varepsilon)$.*

(A subset $N \subseteq U$ is an ε -net of (the unit ball of) a subspace U if for every vector $x \in U$ with norm $\|x\| \leq 1$, the net contains a vector $y \in N$ with $\|x - y\| \leq \varepsilon$.)

An immediate consequence of the previous lemma is an approximation algorithm for SMALL-SET EXPANSION with running time exponential in $\dim U_\eta$. (Here, we use that for a subspace of dimension d , one can construct ε -nets of size $(1/\varepsilon)^{O(d)}$.)

The next lemma shows how to find small sets with low expansion when the dimension of U_η is large. In this case, there exists a vertex and number of steps t such that one can find a set with small volume and low expansion among the level sets of the probability distribution given by the t -step lazy random walk started at this vertex.

Lemma 5.4. *Suppose $\dim U_\eta > n^\beta/\delta$ for n^β sufficiently large. Then, there exist a vertex $i \in V$ and a number $t \in \{1, \dots, \eta^{-1} \log n\}$ such that a level set S of $(\frac{1}{2}I + \frac{1}{2}G)^t \mathbf{1}_i$ satisfies $\mu(S) \leq \delta$ and $\Phi(S) \leq O(\sqrt{\eta/\beta})$.*

Theorem 5.1 (Subexponential Algorithm for SMALL-SET EXPANSION) follows by combining the previous lemmas.

Proof of Theorem 5.1. Let $\eta = O(\varepsilon/\gamma^9)$. If $\dim U_\eta > n^\beta/\delta$, then, using **Lemma 5.4**, we can efficiently find a vertex set S with $\mu(S) \leq \delta$ and $\Phi(S) \leq$

5. Subexponential Approximation Algorithms

$O(\sqrt{\eta/\beta}) = O(\sqrt{\varepsilon/\beta\gamma^9})$. Otherwise, if $\dim U_\eta < n^\beta/\delta$, we construct an $\varepsilon/100$ -net N of U_η of size $(1/\varepsilon)^{O(n^\beta/\delta)}$ and enumerate all level sets of functions in N . Since the given graph G contains a vertex set with volume at most δ and expansion at most ε , by [Lemma 5.3](#), one of enumerated level sets S satisfies $\mu(S) \leq (1 + \gamma)\delta$ and $\Phi(S) \leq O(\sqrt{\varepsilon/\gamma^9})$. \square

5.1.2. Unique Games

Related to the notion of threshold rank, we define the *soft-threshold rank* of a symmetric matrix A at τ as

$$\text{rank}_\tau^*(A) \stackrel{\text{def}}{=} \inf_{t \in \mathbb{N}} \text{Tr} A^{2t} / \tau^{2t}.$$

The soft-threshold rank upper bounds the threshold rank, $\text{rank}_\tau^*(A) \geq \text{rank}_\tau(A)$, because $\text{Tr} A^{2t} \geq \tau^{2t} \cdot \text{rank}_\tau(A)$ for all $t \in \mathbb{N}$.

The following theorem gives an algorithm for UNIQUE GAMES that achieves a good approximation in time exponential in the alphabet size and the soft-threshold rank of the constraint graph (for appropriate threshold value).

Theorem 5.5. *Given $\eta > 0$ and a regular unique game \mathfrak{U} with n vertices, alphabet size k , constraint graph G , and optimal value $\text{opt}(\mathfrak{U}) = 1 - \varepsilon$, we can compute in time $\exp(k \cdot \text{rank}_{1-\eta}^*(G))$ an assignment x of value $\mathfrak{U}(x) \geq 1 - O(\varepsilon/\eta)$.*

For a regular graph G and a vertex subset A , let $G[A]$ denote the “regularized subgraph” induced by A , i.e., we restrict G to the vertex set A and add self-loops to each vertices in A so as to restore its original degree.

The following theorem shows that by changing only a small fraction of its edges, any regular graph can be turned into a vertex disjoint union of regular graphs with few large eigenvalues.

Theorem 5.6 (Threshold rank decomposition). *There exists a polynomial-time algorithm that given parameters $\eta, \beta > 0$ and a regular graph G with n vertices, computes a partition \mathcal{P} of the vertices of G such that $\text{rank}_{1-\eta}^*(G[A]) \leq n^\beta$ for all $A \in \mathcal{P}$ and at most an $O(\sqrt{\eta/\beta^3})$ fraction of the edges of G does not respect the partition \mathcal{P} .*

[Theorem 5.2](#) (Subexponential algorithm for UNIQUE GAMES) follows by combining [Theorem 5.5](#) and [Theorem 5.6](#).

5. Subexponential Approximation Algorithms

Proof of Theorem 5.2 (special case). Let \mathfrak{U} be a regular unique game with vertex set $V = [n]$, alphabet size k , and optimal value $\text{opt}(\mathfrak{U}) = 1 - \varepsilon$. Let x be an optimal assignment for \mathfrak{U} so that $\mathfrak{U}(x) = 1 - \varepsilon$. Let $\eta > 0$ (we determine this parameter later). Apply Theorem 5.6 to the constraint graph G of \mathfrak{U} and $\mathcal{P} = \{A_1, \dots, A_r\}$ be the resulting partition of V . We replace every constraint that does not respect the partition \mathcal{P} by two self-loops with identity constraints (each of the new constraints has weight half). After replacing constraints that do not respect \mathcal{P} in this way, the unique game \mathfrak{U} decomposes into vertex disjoint unique games $\mathfrak{U}_1, \dots, \mathfrak{U}_r$ with vertex sets A_1, \dots, A_r , respectively. The constraint graph of \mathfrak{U}_i is $G[A_i]$ (the regularized subgraph of G induced by A_i). It holds that $\sum_i \mu(A_i) \mathfrak{U}_i(x) \geq \mathfrak{U}(x)$ (since self-loops with identity constraints are always satisfied) and therefore

$$\sum_i \mu(A_i) \text{opt}(\mathfrak{U}_i) \geq 1 - \varepsilon. \quad (5.1)$$

By Theorem 5.6, we have $\text{rank}_{1-\eta}^*(G[A_i]) \leq n^\beta$ for every $i \in [r]$. Hence, using Theorem 5.5, can compute in time $\text{poly}(n) \exp(n^\beta)$ assignments $x^{(1)}, \dots, x^{(r)}$ for $\mathfrak{U}_1, \dots, \mathfrak{U}_r$ such that $\mathfrak{U}_i(x^{(i)}) \geq 1 - \varepsilon_i/\eta$, where $\varepsilon_i \geq 0$ is such that $\text{opt}(\mathfrak{U}_i) = 1 - \varepsilon_i$. From (5.1) it follows that $\sum_i \mu(A_i) \varepsilon_i \leq \varepsilon$ and thus, $\sum_i \mu(A_i) \mathfrak{U}_i(x^{(i)}) \geq 1 - \varepsilon/\eta$. Let x' be the concatenation of the assignments $x^{(1)}, \dots, x^{(r)}$. Since at most $O(\sqrt{\eta/\beta^3})$ of the constraints of \mathfrak{U} do not respect the partition $\{A_1, \dots, A_r\}$, the value of the computed assignment x' is at least $\mathfrak{U}(x') \geq 1 - \varepsilon/\eta - O(\sqrt{\eta/\beta^3})$. Hence, in order to satisfy a constant fraction, say 0.9, of the constraints, we can choose $\eta = O(\varepsilon)$ and $\beta = O(\varepsilon^{1/3})$. In this case, the total running time of the algorithm is exponential in $n^{O(\varepsilon^{1/3})}$. \square

5.2. Subspace Enumeration

In this section, we prove Lemma 5.3 (see §5.2.1) and Theorem 5.5 (see §5.2.2). Before, we prove a general theorem (Theorem 5.7), which we use in the proof of Theorem 5.5 and which also implies a quantitatively weaker version of Lemma 5.3.

For a regular graph G (again identified with its stochastic adjacency matrix) and $\eta \in [0, 1)$, define $\text{rank}_{1-\eta}^+(G)$ as the number of eigenvalues of G larger than $1 - \eta$ (in contrast to $\text{rank}_{1-\eta}(G)$, we only consider positive eigenvalues). Clearly, $\text{rank}_{1-\eta}^+(G) \leq \text{rank}_{1-\eta}(G)$. We can express $\text{rank}_{1-\eta}^+(G)$ in terms of the

5. Subexponential Approximation Algorithms

usual threshold rank, for example

$$\text{rank}_{1-\eta}^+(G) = \text{rank}_{1-\eta/2}\left(\frac{1}{2}I + \frac{1}{2}G\right).$$

Most lemmas and theorems in this section use $\text{rank}_{1-\eta}^+(G)$ instead of $\text{rank}_{1-\eta}(G)$.

The next theorem shows that given a regular graph and a parameter η , we can compute a list of vertex sets of length exponential in the threshold rank of the graph at $1 - \eta$ such that every vertex set with expansion significantly smaller than η is close to one of the sets in the list.

Theorem 5.7. *Given a parameter $\eta > 0$ and a regular graph G with vertex set $V = [n]$, we can compute a list of vertex sets $S_1, \dots, S_r \subseteq V$ with $r \leq \exp(\text{rank}_{1-\eta}^+(G) \log n)$ such that for every vertex set $S \subseteq V$, there exists an index $i \in [r]$ with*

$$\frac{\mu(S \Delta S_i)}{\mu(S)} \leq 8\Phi(S)/\eta.$$

(The computation is efficient in the output size.)

Proof. Let U_η be the subspace of $L_2(V)$ spanned by the eigenfunctions of G with eigenvalue at least $1 - \eta$. Note $\dim U_\eta = \text{rank}_{1-\eta}^+(G)$. Construct a $1/99n$ -net N for (the unit ball of) U_η of size $\exp(\dim U_\eta \cdot \log n)$. Compute S_1, \dots, S_r by enumerating the level sets of the functions in N .

To verify the conclusion of the theorem, consider a vertex set S with expansion $\Phi(S) = \varepsilon$. Let $L = I - G$ be the Laplacian of G and let P_η be the projector onto U_η . On the one hand, $\langle \mathbf{1}_S, L\mathbf{1}_S \rangle = \varepsilon \|\mathbf{1}_S\|^2$. On the other hand,

$$\langle \mathbf{1}_S, L\mathbf{1}_S \rangle = \langle P_\eta \mathbf{1}_S, L(P_\eta \mathbf{1}_S) \rangle + \langle \mathbf{1}_S - P_\eta \mathbf{1}_S, L(\mathbf{1}_S - P_\eta \mathbf{1}_S) \rangle \geq \eta \|\mathbf{1}_S - P_\eta \mathbf{1}_S\|^2.$$

(The inequality uses that $\mathbf{1}_S - P_\eta \mathbf{1}_S$ is orthogonal to U_η and thus in the span of the eigenfunctions of L with eigenvalue larger than η .) It follows that $\|\mathbf{1}_S - P_\eta \mathbf{1}_S\|^2 \leq (\varepsilon/\eta) \|\mathbf{1}_S\|^2$. Since N is a $1/99n$ -net of U_η , it contains a function $f \in N$ such that $\|\mathbf{1}_S - f\|^2 < (\varepsilon/\eta) \|\mathbf{1}_S\|^2 + 1/8n$. Consider the index $i \in [r]$ such that S_i is the level set $S_i = \{j \in V \mid f(j) \geq 1/2\}$. We bound the volume of the symmetric difference of S and S_i ,

$$\mu(S \Delta S_i) \leq \mathbb{P}_{j \in V} \left\{ |\mathbf{1}_S(j) - f(j)| \geq 1/2 \right\} \leq 4 \|\mathbf{1}_S - f\|^2 < 4(\varepsilon/\eta) \|\mathbf{1}_S\|^2 + 1/2n.$$

It follows that $\mu(S \Delta S_i)$ is at most $8(\varepsilon/\eta)\mu(S)$. (Strictly speaking, we assume here $4(\varepsilon/\eta) \|\mathbf{1}_S\|^2 \geq 1/2n$. However, otherwise, we get $\mu(S \Delta S_i) < 1/n$, which means that $S = S_i$ and thus $\mu(S \Delta S_i) = 0$.) \square

5. Subexponential Approximation Algorithms

Remark 5.8. For simplicity, the proof of the theorem used a net of width $1/\text{poly}(n)$. For interesting range of the parameters, a net of constant, say η , width is enough. In this case, the length of the list (and essentially the running time) is bounded by $n \cdot \exp(\text{rank}_{1-\eta}^+(G) \log(1/\eta))$.

5.2.1. Small-Set Expansion

Let G be a regular graph with vertex set $V = [n]$. For $\eta \in [0, 1]$, let $U_\eta \subseteq L_2(V)$ be the subspace spanned by the eigenfunctions of G with eigenvalue larger than $1 - \eta$.

Recall the spectral profile of G (see §2.2),

$$\Lambda(\delta) = \min_{\substack{f \in L_2(V) \\ \|f\|_1^2 \leq \delta \|f\|^2}} \frac{\langle f, Lf \rangle}{\langle f, f \rangle},$$

where $L = I - G$ is the Laplacian of G .

The next lemma shows the following: Suppose $\Lambda(\delta) = \varepsilon$ (which is the case if the given graph G contains a vertex set with volume at most δ and expansion at most ε). Then, given a fine-enough net of the unit ball of U_η for $\eta \gg \varepsilon$, one can efficiently reconstruct a function h that demonstrates that $\Lambda(4/3\delta) = O(\varepsilon)$. The constant hidden by the $O(\cdot)$ -notation depends on the ratio of η and ε (denoted by β in the statement of the lemma). Here, the constant $4/3$ can be replaced by any constant larger than one (leading to larger constants in the $O(\cdot)$ -notation).

Lemma 5.9. *Let $\beta, \gamma > 0$ be sufficiently small with $\gamma \geq \sqrt{\beta}$. Suppose $\Lambda(\delta) = \varepsilon$ and let $\eta = \varepsilon/\beta$. Then for every $\varepsilon/100$ -net N of U_η , there exists a function $g \in N$ and a threshold value τ such that the function $h = \max\{g - \tau, 0\} \in L_2(V)$ satisfies*

$$\begin{aligned} \langle h, Lh \rangle &\leq (\varepsilon/\beta\gamma) \langle h, h \rangle, \\ \mu(\text{supp}(h)) &\leq (1 + 16\gamma^{1/3})\delta. \end{aligned}$$

Proof. Let f be a minimizer for the spectral profile $\Lambda(\delta)$, so that $\langle f, Lf \rangle = \varepsilon \|f\|^2$ and $\|f\|_1^2 \leq \delta \|f\|^2$. We may assume $f \geq 0$ and $\|f\|^2 = 1$. Let P_η be the projector onto U_η . Then, $\|f - P_\eta f\|^2 \leq \varepsilon/\eta$. (Same argument as in the proof of [Theorem 5.7](#).) Hence, N contains a function g such that $\|f - g\|^2 \leq 2\varepsilon/\eta$. Since $g \in U_\eta$, we have $\langle g, Lg \rangle \leq \eta \|g\|^2$. Therefore, the function h also satisfies $\langle h, Lh \rangle \leq \eta \|g\|^2$.

5. Subexponential Approximation Algorithms

For a threshold value τ , we can upper bound the volume of the support of h by roughly $1/\tau^2$. (Later we choose $\tau \approx 1/\sqrt{\delta}$, so that the support has volume roughly δ .) For $\alpha = (\varepsilon/\eta)^{1/3}/2 \leq 1/2$,

$$\begin{aligned} \mu(\text{supp}(h)) &= \mathbb{P}\{g_i > \tau\} \leq \mathbb{P}\{|f_i - g_i| > \alpha\tau\} + \mathbb{P}\{f_i > (1 - \alpha)\tau\} \\ &\leq \frac{1}{\alpha^2\tau^2}\|f - g\|^2 + \frac{1}{(1-\alpha)^2\tau^2}\|f\|^2 \leq \left(1 + 4\alpha + \frac{2\varepsilon/\eta}{\alpha^2}\right) \cdot \frac{1}{\tau^2} = \left(1 + 10(\varepsilon/\eta)^{1/3}\right) \frac{1}{\tau^2}. \end{aligned} \quad (5.2)$$

On the other hand, we can lower bound the L_2 -mass of h in terms of the L_2 -mass of the function $h' = \max\{f - \tau, 0\}$

$$\begin{aligned} \|h\|^2 &= \mathbb{E}_i \max\{g_i - \tau, 0\}^2 \geq \mathbb{E}_i \max\{f_i - \tau, 0\}^2 - 2\mathbb{E}_i \max\{f_i - \tau, 0\} |f_i - g_i| - \mathbb{E}_i (f_i - g_i)^2 \\ &\geq \|h'\|^2 - 2\|h'\| \|f - g\| - \|f - g\|^2 \geq \|h'\|^2 - 5\sqrt{\varepsilon/\eta}. \end{aligned} \quad (5.3)$$

(Here, the first inequality holds for every vertex $i \in V$. The second inequality uses Cauchy–Schwarz. For the third inequality, we use $\|h'\| \leq \|f\| = 1$ and $\|f - g\| \leq \sqrt{2\varepsilon/\eta}$.) If we choose $\tau = (1 - 2\gamma^{1/3})^2/\sqrt{\delta}$, then [Lemma 5.10](#) asserts that $\|h'\|^2 \geq 8\gamma$. Combined with (5.3), we get $\|h\|^2 \geq 3\gamma$. (Recall that $\gamma \geq \sqrt{\varepsilon/\eta}$.)

We can now verify that h has the desired properties. On the one hand,

$$\frac{\langle h, Lh \rangle}{\langle h, h \rangle} \leq \frac{\eta \|g\|^2}{3\gamma} \leq \eta/\gamma.$$

On the other hand, using (5.2),

$$\mu(\text{supp}(h)) \leq (1 + 10(\varepsilon/\eta)^{1/3}) \frac{1}{\tau^2} \leq (1 + 16\gamma^{1/3}) \delta.$$

(The second inequality uses that $1/\tau^2 \leq (1 + 6\gamma^{1/3})\delta$ for sufficiently small γ .) \square

The following technical lemma is used in the proof of the previous lemma.

Lemma 5.10. *Let (Ω, μ) be a finite probability space. For a nonnegative function $f \in L_2(\Omega)$ and a threshold value $\tau \leq (1 - \varepsilon)^2 \|f\|^2 / \|f\|_1$, consider $g = \max\{f - \tau, 0\}$. Then,*

$$\|g\|^2 \geq \varepsilon^3 \|f\|^2$$

Proof. For $u = \varepsilon^2$, let $A \subseteq \Omega$ be the set of points $\omega \in \Omega$ such that $g(\omega)^2 \geq u f(\omega)^2$ and let $B = \Omega \setminus A$ be its complement. For every point $\omega \in B$, the function f satisfies $(f(\omega) - \tau)^2 \leq u f(\omega)^2$ and thus $(1 - \varepsilon)f(\omega) \leq \tau$. Then,

$$\|g\|^2 \geq u \langle \mathbb{1}_A, f^2 \rangle = u (\|f\|^2 - \langle \mathbb{1}_B f, f \rangle) \geq u (\|f\|^2 - \|f\|_1 \cdot \tau / (1 - \varepsilon))$$

(The last inequality uses that $\mathbb{1}_B f \leq \tau / (1 - \varepsilon)$ pointwise.) The claimed lower bound on $\|g\|^2$ follows from the assumption about τ . \square

5. Subexponential Approximation Algorithms

The following lemma follows by combining [Lemma 5.9](#) and [Lemma 2.2](#) (local Cheeger bound).

Lemma 5.11. *Let $\beta > 0$ and $\gamma \geq 16\beta^{1/6}$ be small enough. Suppose $\Phi(\delta) = \varepsilon$. Then for $\eta = \varepsilon/\beta$, every $\varepsilon/100$ -net of the unit ball of U_η contains a function with a level set S that satisfies*

$$\begin{aligned}\mu(S) &\leq (1 + \gamma)\delta, \\ \Phi(S) &\leq O\left(\sqrt{\varepsilon/\beta\gamma^3}\right).\end{aligned}$$

In the above lemma, a convenient parameter choice is $\beta = (\gamma/16)^6$. In this case, the vertex set S has expansion $\Phi(S) \leq O(\sqrt{\varepsilon/\gamma^9})$.

5.2.2. Unique Games

Let \mathfrak{U} be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$. Let $G = G(\mathfrak{U})$ be the constraint graph of \mathfrak{U} (see [§2.3](#) for a definition).

The *label-extended graph* $\hat{G} = \hat{G}(\mathfrak{U})$ is defined as follows (also see [§2.3](#)): The vertex set \hat{V} of \hat{G} consists of all pairs of vertices $u \in V$ and labels $i \in \Sigma$. For ease of notation, we denote such a pair by u_i . The edge distribution of \hat{G} is generated as follows: Sample a constraint $(u, v, \pi) \sim \mathfrak{U}$. Sample a label $i \in \Sigma$. Output an edge between u_i and $v_{\pi(i)}$.

For $\eta \in [0, 1)$, let \hat{U}_η be the subspace of $L_2(\hat{V})$ spanned by the eigenfunctions of \hat{G} with eigenvalue larger than $1 - \eta$. The dimension of \hat{U}_η is equal to $\text{rank}_{1-\eta}^+(\hat{G})$.

The algorithms described in this subsection run in time exponential in the (positive) threshold rank of the label-extended graph. The following lemma allows us to upper bound (soft-)threshold rank of the label extended graph \hat{G} by k times the soft-threshold rank of the constraint G .

Lemma 5.12. *For every $t \in \mathbb{N}$,*

$$\text{Tr } \hat{G}^t \leq k \text{Tr } G^t.$$

Proof. Let $u_i \in \hat{V}$ be a label-extended vertex. The diagonal entry of \hat{G}^t corresponding to u_i is equal to the probability that the t -step random walk in \hat{G} started at u_i returns to u_i . The projection of this random walk corresponds

5. Subexponential Approximation Algorithms

to the t -step random walk in the constraint graph G started at u . Hence, the probability that the random walk in \hat{G} returns to u_i is at most the probability that the random walk in G return to u . It follows that

$$(\hat{G}^t)_{u_i, u_i} \leq \sum_{j \in \Sigma} (\hat{G}^t)_{u_i, u_j} = (G^t)_{u, u},$$

which implies $\text{Tr } \hat{G}^t \leq k \text{Tr } G^t$ as desired. \square

Lemma 5.13. *Let $\beta > 0$ be sufficiently small. Suppose $\text{opt}(\mathfrak{U}) = 1 - \varepsilon$ and let $\eta = \varepsilon/\beta$. Then every $\varepsilon/100$ -net N of the unit ball of \hat{U}_η , contains a function $g \in N$ such that the function $h \in L_2(\hat{V})$ defined by $h(u_i) = \max\{g(u_i) - \max_{j \neq i} g(u_j), 0\}$ has squared norm at least $1/2$ and satisfies*

$$\langle h, \hat{L}h \rangle \leq 16(\varepsilon/\beta)\langle h, h \rangle,$$

where $\hat{L} = I - \hat{G}$ is the Laplacian of the label-extended graph \hat{G} .

Proof. The proof is similar to the proof of [Lemma 5.9](#). The construction of h is related to the smooth nonnegative orthogonalization in [§4.3.2](#).

Let x be an assignment such that $\mathfrak{U}(x) = 1 - \varepsilon$. Consider $f \in L_2(\hat{V})$ such that $f(u_i) = k$ if $x_u = i$ and $f(u_j) = 0$ otherwise (i.e., f is k times the indicator function of the assignment x). The function f satisfies $\langle f, Lf \rangle = \varepsilon\langle f, f \rangle$ and $\|f\|^2 = 1$. Following the argument in [Lemma 5.9](#), the net N contains a function g such that $\|f - g\|^2 \leq 2\varepsilon/\eta$. Since g is in the unit ball of the subspace \hat{U}_η , it satisfies $\langle g, \hat{L}g \rangle \leq \eta = \varepsilon/\beta$. The analysis of the smooth nonnegative orthogonalization in [§4.3.2](#) shows that

$$\langle h, \hat{L}h \rangle \leq 8\langle g, \hat{L}g \rangle.$$

It remains to lower bound the squared norm of h . It is enough to show that h is close to the function f (which implies that the squared norm of h cannot be much smaller than the squared norm of f). Consider the (non-linear) operator Q on $L_2(\hat{V})$ that maps g to h (extended to all of $L_2(\hat{V})$ in the natural way). Notice that $Qf(u_i) = \max\{f(u_i) - \max_{i \neq j} f(u_j), 0\} = f(u_i)$ (the operator Q maps f to itself), since f has only one non-zero among the points u_1, \dots, u_k . Hence, we can estimate the distance h and f in terms of the Lipschitz constant L_Q of the operator Q ,

$$\|h - f\|^2 = \|Qg - Qf\|^2 \leq L_Q^2 \|g - f\|^2 \leq L_Q^2 \cdot 2\beta.$$

5. Subexponential Approximation Algorithms

It is straight-forward to verify that the Lipschitz constant L_Q is an absolute constant (see the analysis in §4.3.2 for details). It follows that $\|h\|^2 \geq \|f\|^2 - O(\|f - h\|) \geq 1 - O(\beta)^{1/2}$. Since β is assumed to be sufficiently small, we get $\|h\|^2 \geq 1/2$. \square

The next lemma follows by combining the previous lemma and [Lemma 2.2](#) (local Cheeger bound). If we apply the local Cheeger bound to the function h in the previous lemma, the resulting vertex set in the label-extended graph corresponds to a partial assignment. (The transformation used to construct h in fact ensures that the support of h is a partial assignment.)

Lemma 5.14. *Let $\beta > 0$ be sufficiently small. Suppose $\text{opt}(\mathfrak{U}) = 1 - \varepsilon$ and let $\eta = \varepsilon/\beta$. Then, in time $\text{poly}(n)(1/\varepsilon)^{O(\dim U_\eta)}$, we can compute a partial assignment $x \in (\Sigma \cup \{\perp\})^V$ with $\mathbb{P}_u\{x_u \neq \perp\} = \alpha > 0$ that satisfies*

$$\mathfrak{U}(x) \geq (1 - O(\sqrt{\varepsilon/\beta})) \cdot \alpha.$$

Here, one can also ensure $\alpha \geq 1/4$.

5.3. Threshold Rank vs. Small-Set Expansion

In this section we prove [Lemma 5.4](#) (restated below). The general goal is to establish a relation between the threshold rank and the expansion profile of graphs.

Let G be a regular graph with vertex set $V = [n]$. Let $L = I - G$ be the Laplacian of G . (Recall that we identify G with its stochastic adjacency matrix.) For $\eta \in [0, 1)$, let U_η denote the subspace of $L_2(V)$ spanned by the eigenfunctions of G with eigenvalue larger than $1 - \eta$.

Lemma (Restatement of [Lemma 5.4](#)). *Suppose $\dim U_\eta > n^\beta/\delta$ for n^β sufficiently large. Then, there exist a vertex $i \in V$ and a number $t \in \{1, \dots, \eta^{-1} \log n\}$ such that a level set S of $(\frac{1}{2}I + \frac{1}{2}G)^t \mathbf{1}_i$ satisfies $\mu(S) \leq \delta$ and $\Phi(S) \leq O(\sqrt{\eta/\beta})$.*

An important ingredient of the proof of [Lemma 5.4](#) (and other results in this section) is the following local variant of Cheeger's inequality. (The lemma follows directly from [Lemma 2.2](#) in [Chapter 2](#).)

Lemma 5.15 (Local Cheeger Bound). *Suppose $f \in L_2(V)$ satisfies $\|f\|_1^2 \leq \delta \|f\|^2$ and $\langle f, Lf \rangle \leq \varepsilon \|f\|^2$. Then, for every $\gamma > 0$, there exists a level set S of f such that $\mu(S) \leq \delta/(1 - \gamma)$ and $\Phi(S) \leq \sqrt{2\varepsilon/\gamma}$.*

5. Subexponential Approximation Algorithms

Since $\dim U_\eta = \text{rank}_{1-\eta/2}(\frac{1}{2}I + \frac{1}{2}G) \leq \text{rank}_{1-\eta/2}^*(\frac{1}{2}I + \frac{1}{2}G)$, the following lemma combined with [Lemma 5.15](#) (Local Cheeger Bound) implies [Lemma 5.4](#). Note that $\frac{1}{2}I + \frac{1}{2}G = I - L/2$.

Lemma 5.16. *Suppose $\text{rank}_{1-\eta/2}^*(I - L/2) > n^\beta/\delta$ for $\eta > 0$ small enough and n^β large enough. Then, there exist a vertex $i \in V$ and a number $t \in \{0, \dots, (\beta/\eta) \log n\}$ such that the function $f = (I - L/2)^t \mathbf{1}_i$ satisfies $\|f\|_1^2 \leq \delta \|f\|^2$ and $\langle f, Lf \rangle / \langle f, f \rangle \leq 5\eta/\beta$.*

Proof. Let $f_{t,i} = (I - L/2)^t \mathbf{1}_i$ and $T = \lfloor (\beta/2\eta) \ln n \rfloor$. Since $I - L/2$ is symmetric stochastic, the L_1 -norm of $f_{t,i}$ is decreasing in t , so that $\|f_{t,i}\|_1 \leq \|f_{0,i}\|_1 = 1/n$ (in fact, equality holds). The lower bound on the soft-threshold rank of $I - L/2$ implies for every $t \in \mathbb{N}$ (in particular, for $t = T$),

$$(1 - \eta/2)^{2t} \cdot n^\beta/\delta \leq \text{Tr}(I - L/2)^{2t} = n \sum_i \langle \mathbf{1}_i, (I - L/2)^{2t} \mathbf{1}_i \rangle = n \sum_i \|f_{t,i}\|^2.$$

By our choice of T , we have $(1 - \eta/2)^{2T} n^\beta \geq 1$ (using that η is small enough). Therefore, there exists a vertex $i \in V$ such that $\|f_{T,i}\|^2 \geq 1/\delta n^2$. The following claims imply the current lemma:

- (1) for every $t \in \{0, \dots, T\}$, it holds that $\|f_{t,i}\|_1^2 \leq \delta \|f_{t,i}\|^2$, and
- (2) there exists $t \in \{0, \dots, T-1\}$ such that $\langle f_{t,i}, Lf_{t,i} \rangle \leq \varepsilon$ for $\varepsilon = 5\eta/\beta$.

The first claim holds because on the one hand, $\|f_{t,i}\|_1 \leq \|f_{0,i}\|_1 = 1/n^2$ and on the other hand, $\|f_{t,i}\|^2 \geq \|f_{T,i}\|^2 \geq 1/\delta n^2$. (Here, we use that the matrix $I - L/2$, being symmetric stochastic, decreases both L_1 -norms and L_2 -norms.) It remains to verify the second claim. Since $\|f_{T,i}\|^2 \geq 1/\delta n^2 = \|f_{0,i}\|^2/\delta n$, there exists $t \in \{0, \dots, T-1\}$ such that $\|f_{t+1,i}\|^2 \geq (\delta n)^{-1/T} \|f_{t,i}\|^2$. By our choice of T , we have $(\delta n)^{-1/T} \geq 1 - \varepsilon/2$ for $\varepsilon = 5\eta/\beta$ (using that n is large enough). It follows that

$$(1 - \varepsilon/2) \|f_{t,i}\|^2 \leq \|f_{t+1,i}\|^2 = \langle f_{t,i}, (I - L/2)^2 f_{t,i} \rangle \leq \langle f_{t,i}, (I - L/2) f_{t,i} \rangle$$

(Here, we use that $(I - L/2)^2 - (I - L/2) = L/2 - L^2/4$ is positive semidefinite, because the eigenvalues of $L/2$ lies in $[0, 1]$.) \square

5.4. Low Threshold Rank Decomposition

The goal of this section is to prove [Theorem 5.6](#).

Theorem (Restatement of [Theorem 5.6](#)). *There exists a polynomial-time algorithm that given parameters $\eta, \beta > 0$ and a regular graph G with n vertices, computes a partition \mathcal{P} of the vertices of G such that $\text{rank}_{1-\eta}^*(G[A]) \leq n^\beta$ for all $A \in \mathcal{P}$ and at most an $O(\sqrt{\eta/\beta^3})$ fraction of the edges of G does not respect the partition \mathcal{P} .*

Recall that $G[A]$ denote the “regularized subgraph” induced by A , i.e., we restrict G to the vertex set A and add self-loops to each vertices in A so as to restore its original degree.

The main ingredient of the proof of [Theorem 5.6](#) is the following consequence of [Lemma 5.16](#) (and [Lemma 5.15](#), the local Cheeger bound).

Lemma 5.17. *There exists a polynomial-time algorithm that given parameters $\eta, \beta > 0$, a regular graph G , and a vertex set K of G such that $\text{rank}_{1-\eta}^*(G[A]) > n^\beta$, computes a vertex set $S \subseteq K$ with volume $\mu(S) \leq \mu(A)/n^{\Omega(\beta)}$ and expansion $\Phi_{G[K]}(S) \leq O(\sqrt{\eta/\beta})$ in the graph $G[K]$.*

We remark that the condition $\text{rank}_{1-\eta}^*(G[A]) > n^\beta$ in general does not imply the condition $\text{rank}_{1-\eta}^*(\frac{1}{2}I + \frac{1}{2}G[A]) > n^\beta$ required for [Lemma 5.16](#) (in case $G[A]$ has negative eigenvalues). However, it holds that $\text{rank}_{1-\eta}^*(G[A]) \leq \text{rank}_{1-\eta}^*(\frac{1}{2}I + \frac{1}{2}G[A]) + \text{rank}_{1-\eta}^*(\frac{1}{2}I - \frac{1}{2}G[A])$. Furthermore, [Lemma 5.16](#) also works under the condition that $\text{rank}_{1-\eta}^*(\frac{1}{2}I - \frac{1}{2}G[A])$ is large (same proof).

The proof of [Theorem 5.6](#) follows by iterating the previous lemma.

Proof of [Theorem 5.6](#). We decompose the given graph G using the following iterative algorithm. As long as G contains a connected component $K \subseteq V(G)$ with $\text{rank}_{1-\eta}^*(G[A]) > n^\beta$, do the following steps:

1. Apply [Lemma 5.17](#) to the component K to obtain a vertex set $S \subseteq K$ with volume $\mu(S) \leq \mu(A)/n^{\Omega(\beta)}$ and expansion $\Phi(S) \leq O(\sqrt{\eta/\beta})$. (Since K is a connected component of G , the vertex set $S \subseteq K$ has the same expansion in the graphs G and $G[K]$.)
2. Replace every edge leaving S by two self-loops on the endpoints of the edge. (Each self-loop has weight half.)

5. Subexponential Approximation Algorithms

The algorithm terminates after at most n iterations (in each iteration the number of connected components increases by at least one). When the algorithm terminates, every connected component satisfies the desired bound on the soft-threshold rank.

It remains to bound the fraction of edges changed by the algorithm. Consider the following charging argument. When we replace the edge leaving a vertex set S , we charge cost $\Phi(S)$ to every vertex in S . Let $\text{cost}(i)$ be total cost charged to vertex i until the algorithm terminates. The fraction of edges changed by the algorithm is no more than $\mathbb{E}_{i \sim V(G)} \text{cost}(i)$. How often can we charge the same vertex? If we charge a vertex t times, then the volume of its connected component K is at most $\mu(K) \leq n^{-\Omega(t\beta)}$. It follows that every vertex is charge at most $O(1/\beta)$ times. Hence, we can bound the total cost charged to a vertex i by $\text{cost}(i) \leq O(\sqrt{\eta/\beta^3})$, which implies the desired bound on the fraction of edges changed by the algorithm. \square

5.5. Putting things together

In this section, we prove the general case of [Theorem 5.2](#) (restated below).

Theorem (Restatement of [Theorem 5.2](#)). *There exists an algorithm that given a unique game \mathfrak{U} with n vertices, alphabet size k , and optimal value $\text{opt}(\mathfrak{U}) \geq 1 - \varepsilon$, computes in time $\exp(kn^\beta)$ an assignment x of value $\mathfrak{U}(x) \geq 1 - O(\sqrt{\varepsilon/\beta^3})$. Here, β is a parameter of the algorithm that can be chosen as small as $\log \log n / \log n$.*

Proof. To prove the theorem it is enough to show how to compute in the desired running a *partial* assignment x with value $\mathfrak{U}(x) \geq (1 - O(\sqrt{\varepsilon/\beta^3}))\alpha$, where $\alpha > 0$ is the fraction of vertices labeled by x . (We can then extend this partial assignment to a total one by iterating the algorithm for the remaining unlabeled vertices. After at most n iterations, all vertices are labeled. It is straight-forward to verify that the approximation guarantee does not suffer computing an assignment in this iterative fashion.)

To compute the desired partial assignment, we first decompose the constraint graph G of the unique game \mathfrak{U} using [Theorem 5.6](#). In this way, we obtain a partition $\mathcal{P} = \{A_1, \dots, A_r\}$ of the vertex set such that $\text{rank}_{1-\eta}^*(G[A_i]) \leq n^\beta$ for all $i \in [r]$ and at most a $1 - O(\sqrt{\eta/\beta^3})$ fraction of the constraints of \mathfrak{U} do not respect the partition \mathcal{P} . Here, we choose the parameter $\eta > 0$ as $\eta = O(\varepsilon)$ (for a sufficiently large constant in the $O(\cdot)$ -notation).

5. Subexponential Approximation Algorithms

From the unique game \mathfrak{U} , we construct unique games $\mathfrak{U}_1, \dots, \mathfrak{U}_r$ with vertex sets A_1, \dots, A_r , respectively. (We replace every constraint of \mathfrak{U} that does not respect the partition by two identity constraints on the endpoints of weight half.) In this way, the constraint graph of \mathfrak{U}_i is the graph $G[A_i]$ (the regularized subgraph of G induced by A_i).

An averaging argument shows that there exists $i \in [r]$ such that $\text{opt}(\mathfrak{U}_i) \geq 1 - 2\varepsilon$ and $\Phi_G(A_i) = O(\sqrt{\eta/\beta^3})$.

Using [Lemma 5.14](#), we can compute partial assignments $x^{(i)}$ for the unique game \mathfrak{U}_i that labels an α fraction of the vertices in A_i such that $\mathfrak{U}_i(x^{(i)}) \geq (1 - O(\sqrt{\varepsilon}))\alpha$ (using that ε/η is sufficiently small). Furthermore, $\alpha_i \geq 1/4$. How good a partial assignment is $x^{(i)}$ for the original unique game \mathfrak{U} ? Since $x^{(i)}$ is a partial assignment for the vertex set A_i , we have

$$\mathfrak{U}(x^{(i)}) \geq \mu(A_i)\mathfrak{U}_i(x^{(i)}) - \Phi(A_i)\mu(A_i)$$

(using the fact we replaced the constraints leaving A_i by identity constraints). Since $\alpha \geq 1/4$, it holds that

$$\mathfrak{U}(x^{(i)}) \geq (1 - O(\sqrt{\varepsilon}) - 4\Phi(A_i))\alpha\mu(A_i).$$

Since $x^{(i)}$ labels an $\alpha\mu(A_i)$ fraction of the vertices of \mathfrak{U} and $\Phi(A_i) \leq O(\sqrt{\varepsilon/\beta^3})$, the partial assignment $x^{(i)}$ has the required properties,

$$\mathfrak{U}(x^{(i)}) \geq (1 - O(\sqrt{\varepsilon/\beta^3})) \mathbb{P}_{u \sim \mathfrak{U}} \{x_u^{(i)} \neq \perp\}. \quad \square$$

5.6. Notes

The material presented in this chapter is based on the paper “Subexponential Algorithms for Unique Games and Related Problems” [[ABS10](#)], joint with Sanjeev Arora and Boaz Barak. A preliminary version of the paper appeared at FOCS 2010.

Subspace Enumeration

Some of the material presented in [Section 5.2](#) (Subspace Enumeration) is inspired by the works of Kolla and Tulsiani [[KT07](#)] and Kolla [[Kol10](#)]. In [[KT07](#)], the authors present an alternative proof of a quantitatively slightly weaker version of [Theorem 3.1](#), the main theorem of [Chapter 3](#) (Unique Games with Expanding Constraint Graph). To find a good assignment for

5. Subexponential Approximation Algorithms

a unique game, their algorithm enumerates a suitably discretized subspace of real-valued functions on the label-extended graph of the unique game (as in [Section 5.2.2](#)). They show that if a unique game with alphabet size k has an assignment with value close to 1 and its constraint graph is an expander, then the label-extended graph has at most k eigenvalues close to 1. (We omit the precise quantitative trade-offs for this discussion.) Furthermore, they show that any good assignment for the unique game corresponds to a function close to the subspace spanned by the eigenfunctions of the label-extended graph with eigenvalue close to 1. It follows that by enumerating this subspace (suitably discretized) one finds a function from which one can recover a good assignment for the unique game. (The running time is exponential in the dimension of this subspace, which is bounded by k in this case.) Their recovering procedure roughly corresponds to the construction in the proof of [Theorem 5.7](#) (though we establish this theorem in the more general context of graph expansion). In (the proof of) [Lemma 5.14](#) we give a more sophisticated recovering procedure which has quantitatively better approximation guarantees.

Kolla [[Kol10](#)] noted that several unique games considered in the literature (in particular, the ones in [[KV05](#)]) have label-extended graphs with only few (say $\text{poly log } n$) eigenvalues close to 1 (even though their constraint graphs are not good enough expanders). Since the algorithm in [[KT07](#)] runs in time exponential in the number of eigenvalues of the label-extended graph that are close to 1, it follows that for the unique games in [[KV05](#)] the algorithm of [[KT07](#)] runs in quasi-polynomial time. (Since the unique games in [[KV05](#)] have no good assignments, the algorithm of [[KT07](#)] will certify that indeed no good assignments exists.) Furthermore, Kolla [[Kol10](#)] showed that if the constraint graph of a unique game has only few large eigenvalues, then also the label-extended graph has few large eigenvalues assuming that the eigenfunctions of the constraint graph are “well-spread” (in the sense that the L_∞ and the ℓ_2 norm are proportional). It follows that the algorithm of [[KT07](#)] applies to all unique games whose constraint graphs satisfy these properties (which includes the unique games in [[KV05](#)]). In this chapter, we use [Lemma 5.12](#) to related the eigenvalues of the constraint graph to the eigenvalues of the label-extended graph. The advantage of [Lemma 5.12](#) is that it does not assume additional properties of eigenfunctions of the constraint graphs (which are hard to control in general). The proof of [Lemma 5.12](#) is not very involved. It follows by comparing the behavior of random walks on the label-extended graph to random walks on the constraint graph.

Decomposition Approach for Unique Games

Our subexponential algorithm for UNIQUE GAMES works by decomposing a general instance into a collection of (independent) “easy instances”. In the context of UNIQUE GAMES, this divide-and-conquer strategy has been investigated by two previous works, which we discuss in the following.

One of Trevisan’s algorithms [Tre05] removes a constant fraction of the constraints in order to decompose a general unique game into a collection of disjoint unique games whose constraint graphs have spectral gap at least $1/\text{poly log } n$. His algorithm gives a non-trivial approximation guarantee if the value of the unique game is at least $1 - 1/\text{poly log } n$.

Arora et al. [AIMS10] studied a variation of Trevisan’s approach (combined with the results in Chapter 3). Here, the idea is to remove a constant fraction of constraints in order to decompose a general unique game with n vertices into disjoint instances that are either small (containing at most $n/2^{\Omega(1/\varepsilon)}$ vertices) or their constraint graph has spectral gap at least ε . (A decomposition with precisely these guarantees is not achieved in [AIMS10]. However, the algorithm works *as if* these guarantees were achieved. We refer to [AIMS10] for details.) For small instances, one can find good assignments in time $2^{n/2^{\Omega(1/\varepsilon)}}$ (using brute-force enumeration) and for instances whose constraint graph has spectral gap ε , one can use the algorithm for UNIQUE GAMES presented in Chapter 3. The algorithm of [AIMS10], given a unique game with value $1 - \varepsilon$, find an assignment with value at least $1/2$ in time

Part II.

Reductions

6. Graph Expansion and the Unique Games Conjecture

A major shortcoming of the current knowledge about the Unique Games Conjecture (compared to many other conjectures in complexity theory) is that only few consequences of a refutation of the conjecture are known. For example, the following scenario is not ruled out: There is a polynomial time algorithm refuting the Unique Games Conjecture and, at same time, no polynomial time algorithm — or even subexponential time algorithm — has a better approximation ratio for MAX CUT than the Goemans–Williamson algorithm. (In contrast, a refutation of the conjecture that 3-SAT has no polynomial time algorithms implies that *every* problem in NP has a polynomial time algorithm. Similarly, a refutation of the Exponential Time Hypothesis [IPZ01] — 3-SAT has no $2^{o(n)}$ time algorithm — implies surprising algorithms for a host of other problems.) The lack of surprising consequences of a refutation of the UGC is one of main reasons why no consensus regarding the truth of the conjecture has been reached among researchers (whereas many other complexity conjectures have a strong consensus among researchers).

Prior to this work, the only known non-trivial consequence of an algorithmic refutation of the Unique Games Conjecture is an improved approximation algorithm for UNIQUE GAMES itself. Rao [Rao08] showed that if there exists a polynomial time algorithm refuting the Unique Games Conjecture, then there exists an algorithm that given a unique game with value $1 - \varepsilon$, finds in polynomial time an assignment of value $1 - C\sqrt{\varepsilon}$. (The constant $C \geq 1$ and the degree of the polynomial for the running time of this algorithm depend on the guarantees of the algorithm assumed to refute the UGC.)

In this work, we demonstrate the first non-trivial consequence of an algorithmic refutation of the Unique Games Conjecture for a problem different than UNIQUE GAMES. This problem, called SMALL-SET EXPANSION, is a natural generalization of SPARSEST CUT, one of the most fundamental optimization problems on graphs (see e.g. [LR99, ARV09]). We give a reduction from SMALL-SET EXPANSION to UNIQUE GAMES. An alternative formulation of our

6. Graph Expansion and the Unique Games Conjecture

result: If SMALL-SET EXPANSION is NP-hard to approximate (in a certain range of parameters), then the Unique Games Conjecture holds true.

6.1. Main Results

To state our results we introduce the following hypothesis about the approximability of SMALL-SET EXPANSION (restated from §1.4).

Hypothesis 6.1 (Small Set-Expansion (SSE) Hypothesis). *For every $\varepsilon > 0$, there exists $\delta > 0$ such that the following promise problem is NP-hard: Given a graph G , distinguish between the cases,*

YES: *some vertex set of volume δ has expansion at most ε .*

NO: *no vertex set of volume δ has expansion less than $1 - \varepsilon$.*

We remark the best known polynomial-time approximation algorithms for SMALL-SET EXPANSION [RST10a] (based on a basic SDP relaxation) fail to refute this conjecture. The conjecture also holds in certain hierarchies of relaxations (see Chapter 8).

In this work, we show relations of this hypothesis to the approximability of UNIQUE GAMES. First, we show that the Unique Games Conjecture is true if the Small-Set Expansion Hypothesis holds.

Theorem 6.2. *The Small-Set Expansion Hypothesis implies the Unique Games Conjecture.*

The proof of this theorem is based on a reduction from SMALL-SET EXPANSION to UNIQUE GAMES (the composition of Reduction 6.5 and Reduction 6.14). We prove the theorem at the end of this section.

Second, we show that the Small-Set Expansion Hypothesis is true if the following stronger variant of the Unique Games Conjecture holds. (The results in the following chapter (Chapter 7) imply that this variant of the Unique Games Conjecture is in fact equivalent to the Small-Set Expansion Hypothesis.)

Hypothesis 6.3 (Unique Games Conjecture on Small-Set Expanders). *For every $\eta > 0$, there exists $\delta_\eta > 0$ such that for every $\zeta > 0$, the following promise problem is NP-hard for some $R = R_{\eta,\zeta}$: Given a unique game \mathcal{U} with alphabet size R , distinguish between the cases,*

6. Graph Expansion and the Unique Games Conjecture

YES: the optimal value of \mathfrak{U} satisfies $\text{opt}(\mathfrak{U}) \geq 1 - \eta$,

NO: the optimal value of \mathfrak{U} satisfies $\text{opt}(\mathfrak{U}) \leq \zeta$ and every vertex set of volume δ_η in the constraint of \mathfrak{U} has expansion at least $1 - \eta$.

We remark that this variant of the Unique Games Conjecture holds in certain hierarchies of SDP relaxation (using results from [Chapter 7](#) and [Chapter 8](#)).

The following theorem is based on a reduction from UNIQUE GAMES to SMALL-SET EXPANSION ([Reduction 6.20](#) in [§6.4](#)). It follows directly from [Theorem 6.21](#) in [§6.4](#).

Theorem 6.4. *Hypothesis 6.3 (Unique Games Conjecture on Small-Set Expanders) implies the Small-Set Expansion Hypothesis.*

6.1.1. Proof of [Theorem 6.2](#)

Our reduction from SMALL-SET EXPANSION to UNIQUE GAMES (used to prove [Theorem 6.2](#)) naturally decomposes into two parts. First, we reduce SMALL-SET EXPANSION to PARTIAL UNIQUE GAMES (partial assignments are allowed). See [Reduction 6.5](#) and [Theorem 6.6](#) in [§6.2](#). Then, we show how to reduce PARTIAL UNIQUE GAMES to UNIQUE GAMES. See [Reduction 6.14](#) and [Theorem 6.15](#) in [§6.3](#). [Theorem 6.2](#) follows by instantiating [Theorem 6.6](#) and [Theorem 6.15](#) with an appropriate choice of parameters.

6.2. From SMALL-SET EXPANSION to PARTIAL UNIQUE GAMES

Reduction 6.5 (From SMALL-SET EXPANSION to PARTIAL UNIQUE GAMES).

Input: A regular graph G with vertex set V and parameters $\varepsilon > 0$ and $R \in \mathbb{N}$ (satisfying $\varepsilon R \in \mathbb{N}$).

Output: A unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ with vertex set $V^{R'}$ and alphabet $\Sigma = [R']$ for $R' = (1 + \varepsilon)R$.

The unique game \mathfrak{U} corresponds to the following probabilistic verifier for an assignment $F: V^{R'} \rightarrow [R']$:

1. Sample R random vertices $a_1, \dots, a_R \sim G$.

6. Graph Expansion and the Unique Games Conjecture

2. Sample two random neighbors $b_i, b'_i \sim G(a_i)$ for every $i \in [R]$. (Here, the notation $b \sim G(a)$ means that b is a random neighbor of a in G .)
3. Sample $2\varepsilon R$ random vertices $b_{R+1}, b'_{R+1}, \dots, b_{R+\varepsilon R}, b'_{R+\varepsilon R} \sim V$.
4. Let $A = (a_1, \dots, a_R) \in V^R$ and $B = (b_1, \dots, b_R)$, $B' = (b'_1, \dots, b'_{R'}) \in V^{R'}$.
5. Sample two random permutation $\pi, \pi' \in \mathbb{S}_{[R']}$.
6. Verify that $\pi^{-1}(F(\pi.B)) = (\pi')^{-1}F(\pi'.B')$. (Here, $\pi.B$ refers to the tuple obtained by permuting the coordinates of B according to the permutation π . See for a formal definition.) (End of [Reduction 6.5](#))

[Reduction 6.5](#) has the following approximation guarantees.

Theorem 6.6. *Given a regular graph G with n vertices and parameters $R \in \mathbb{N}$ and $\varepsilon > 0$, [Reduction 6.5](#) computes in time $\text{poly}(n^R)$ a unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ such that the following assertions hold (for all $\varepsilon' > \varepsilon$):*

Completeness: *If the graph G contains a vertex set with volume $\delta = 1/2R$ and expansion at most ε , then the unique game \mathfrak{U} has α -partial value $\text{opt}_\alpha(\mathfrak{U}) \geq 1 - 5\varepsilon$, where $\alpha \geq 0.1$.*

Soundness I: *If every vertex set of G with volume $\delta = 1/2R$ has expansion at least $1 - \varepsilon$, then for all $\alpha \geq 0.1$, the unique game \mathfrak{U} has α -partial value $\text{opt}_\alpha(\mathfrak{U}) \leq O(\varepsilon^{1/5})$.*

Soundness II: *If every vertex set of G with volume between $\Omega(\varepsilon^2/R)$ and $O(1/\varepsilon^2R)$ has expansion at least ε' and half of the edges of every vertex of G are self-loops, then for all $\alpha \geq 0.1$, the unique game \mathfrak{U} has α -partial value at most $\text{opt}_\alpha(\mathfrak{U}) \leq 1 - \varepsilon'/4$.*

Remark 6.7. All properties of the reduction would essentially be preserved if we would modify step 3 of the reduction as follows: For every $i \in [R]$, with probability ε , replace b_i by a random vertex in V (do the same with b'_1, \dots, b'_R , independently). In step 4, one would set $B = (b_1, \dots, b_R)$ and $B' = (b_1, \dots, b_R)$.

6.2.1. Completeness

Let G be a regular graph with vertex set $V = \{1, \dots, n\}$. Note that the size of the unique game $\mathfrak{U} = \mathfrak{U}_{\varepsilon,R}(G)$ produced by [Reduction 6.5](#) is $n^{O(R)}$ and thus

6. Graph Expansion and the Unique Games Conjecture

polynomial in the size of the graph G for every $R = O(1)$. (Also the running time of the reduction is $n^{O(R)}$.)

The following lemma shows that [Reduction 6.5](#) is complete, that is, if the graph G contains a set with a certain volume and expansion close to 0, then the unique game obtained from [Reduction 6.5](#) has a partial assignment that satisfies almost all constraints with labeled vertices.

Lemma 6.8 (Completeness). *For every set $S \subseteq V$ with $\mu(S) = \delta$ and $\Phi(S) = \eta$, there exists a partial assignment $F = F_S$ for the unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ (as defined in [Reduction 6.5](#)) satisfying*

$$\mathfrak{U}(F) \geq (1 - \varepsilon - 4\eta)\alpha,$$

where $\alpha \geq (1 - R'\delta)R'\delta$ is the fraction of vertices of \mathfrak{U} labeled by the partial assignment x .

Proof. We may assume that $R' \leq 1/\delta$, because the lemma is trivial otherwise. Consider the following partial assignment F for the unique game \mathfrak{U} ,

$$F: V^{R'} \rightarrow [R'] \cup \{\perp\},$$

$$X = (x_1, \dots, x_{R'}) \mapsto \begin{cases} i & \text{if } \{i\} = \{i' \in [R'] \mid x_i \in S\}, \\ \perp & \text{otherwise.} \end{cases}$$

We compute the fraction of vertices of \mathfrak{U} labeled by F ,

$$\alpha = \mathbb{P}_{X \sim V^{R'}} \{F(X) \neq \perp\} = \binom{R'}{1} (1 - \delta)^{R'-1} \delta = ((1 - \delta)^{R'-1}) R' \delta \geq (1 - R'\delta) R' \delta,$$

Next, we estimate the fraction of constraints satisfied by F in \mathfrak{U} . Sample tuples of vertices A, B, B' as specified by [Reduction 6.5](#). Note that the assignment F behaves nicely when permuting coordinates, $F(\pi.X) = \pi(F(X))$ for every permutation π of $[R']$ and every $X \in V^{R'}$ such that $F(X) \neq \perp$. Hence,

$$\mathfrak{U}(F) = \mathbb{P}_{A,B,B'} \{F(B) = F(B') \in [R']\}.$$

(Here, we also use that a constraint is considered to be satisfied by a partial assignment only if both vertices of the constraint are labeled.) We bound the fraction of satisfied constraints from below

$$\begin{aligned} \mathbb{P}_{A,B,B'} \{F(B) = F(B') \in [R']\} &\geq \mathbb{P}_{A,B,B'} \{F(B) = F(B') \in [R]\} \\ &= \mathbb{P}_{X \sim V^{R'}} \{F(X) \in [R]\} \mathbb{P}_{A,B,B'} \{F(B') = F(B) \mid F(B) \in [R]\}. \end{aligned}$$

6. Graph Expansion and the Unique Games Conjecture

The inequality is almost tight, because the event $F(X) \in [R'] \setminus [R]$ is very unlikely (even relative to the event $F(X) \in [R']$),

$$\mathbb{P}_{X \sim V^{R'}} \{F(X) \in [R]\} = \mathbb{P}_{X \sim V^{R'}} \{F(X) \in [R']\} / (1 + \varepsilon) \geq (1 - \varepsilon)\alpha.$$

Next, we relate the probability of the event $F(B) = F(B')$ conditioned on $F(B) \in [R]$ to the expansion of the set S . It turns out that this probability is most directly related to the expansion of S in the graph G^2 (instead of G). Let $\eta' = \mathbb{P}_{a \sim V, b, b' \sim G(a)} \{b' \notin S \mid b \in S\}$ be the expansion of S in G^2 . Then,

$$\begin{aligned} \mathbb{P}_{A, B, B'} \{F(B') = F(B) \mid F(B) \in [R]\} &= \mathbb{P}_{A, B, B'} \{F(B') = 1 \mid F(B) = 1\} \\ &= (1 - \eta')(1 - \eta'\delta / (1 - \delta))^{R-1} \\ &\geq 1 - \left(1 + \frac{1-1/R}{1-\delta} R\delta\right) \eta' \geq 1 - 2\eta'. \end{aligned}$$

Here, we again use the symmetry of F and we also use that $\mathbb{P}_{a \sim V, b, b' \sim G(a)} \{b' \in S \mid b \notin S\} = \eta'\delta / (1 - \delta)$ (the expansion of $V \setminus S$ in G^2). The last inequality uses that $R' \leq 1/\delta$ and therefore $1 - 1/R \leq 1 - \delta$ and $R\delta \leq 1$.

It remains to relate η' (the expansion of S in G^2) to η (the expansion of S in G).

$$\begin{aligned} \eta'\delta &= \mathbb{P}_{a \sim V, b, b' \sim G(a)} \{b \in S \wedge b' \notin S\} \\ &\leq \mathbb{P}_{a \sim V, b, b' \sim G(a)} \{(b \in S \wedge a \notin S) \vee (a \in S \wedge b' \notin S)\} \\ &\leq \mathbb{P}_{a \sim V, b \sim G(a)} \{b \in S \wedge a \notin S\} + \mathbb{P}_{a \sim V, b' \sim G(a)} \{a \in S \wedge b' \notin S\} = 2\delta\eta. \end{aligned}$$

Combining the previous bounds shows the desired lower bound on the fraction of constraints satisfied by F in \mathfrak{U} ,

$$\begin{aligned} \mathfrak{U}(F) &= \mathbb{P}_{A, B, B'} \{F(B) = F(B') \in [R']\} \\ &\geq \mathbb{P}_{X \sim V^{R'}} \{F(X) \in [R]\} \mathbb{P}_{A, B, B'} \{F(B') = 1 \mid F(B) = 1\} \\ &\geq (1 - \varepsilon)\alpha \cdot (1 - 4\eta) \geq (1 - \varepsilon - 4\eta)\alpha. \end{aligned} \quad \square$$

6.2.2. Soundness

Let $F: V^{R'} \rightarrow [R'] \cup \{\perp\}$ be a partial assignment for the unique game $\mathfrak{U} = \mathfrak{U}_{R, \varepsilon}(G)$ obtained by applying [Reduction 6.5](#) to a regular graph G with vertex set $V = \{1, \dots, n\}$.

6. Graph Expansion and the Unique Games Conjecture

For a tuple $U \in V^{R'-1}$ and a vertex $x \in V$, let $f(U, x)$ be the probability that F selects the coordinate of x after we place it a random position of U and permute the tuple randomly,

$$f(U, x) \stackrel{\text{def}}{=} \mathbb{P}_{i \in [R'], \pi \in \mathbb{S}_{R'}} \left\{ F(\pi.(U +_i x)) = \pi(i) \right\}.$$

Here, $U +_i x$ denotes the tuple obtained from U by inserting x as the i -th coordinate (and moving the original coordinates $i, \dots, R' - 1$ of U by one to the right). (The above experiment wouldn't change if we fixed $i = R'$, because the permutation π is random.)

For $U \in V^{R'-1}$, define the function $f_U: V \rightarrow [0, 1]$,

$$f_U(x) \stackrel{\text{def}}{=} \mathbb{E}_{W \sim G^{\otimes(R-1)}(U), Z \in V^{\varepsilon R}} f(W, Z, x). \quad (6.1)$$

(Here, $G^{\otimes(R-1)}$ refers to the $(R-1)$ -fold tensor product of the graph G , and $W \sim G^{\otimes(R-1)}(U)$ means that W is a random neighbor of the vertex U in $G^{\otimes(R-1)}$. Note that for $U = (u_1, \dots, u_{R-1})$, the distribution $G^{\otimes(R-1)}(U)$ is the product of the distributions $G(u_1), \dots, G(u_{R-1})$.)

The following properties of the functions $\{f_U\}$ are straightforward to verify.

Lemma 6.9. *Let $\alpha = \mathbb{P}_{X \sim V^{R'}} \{F(X) \neq \perp\}$ be the fraction of vertices of \mathfrak{U} labeled by the partial assignment F .*

1. *The typical L_1 -norm of f_U equals $\mathbb{E}_{U \sim V^{R-1}} \|f_U\|_1 = \frac{\alpha}{R'}$.*
2. *For every $U \in V^{R-1}$, the L_1 -norm of f_U satisfies $\|f_U\|_1 \leq \frac{1}{\varepsilon R}$.*
3. *The typical squared L_2 -norm of Gf_U relates to the fraction of constraints satisfied by F in \mathfrak{U} as follows,*

$$\mathbb{E}_{U \sim V^{R-1}} \|Gf_U\|^2 \geq \frac{1}{R'} (\mathfrak{U}(F) - \frac{1}{\varepsilon R}).$$

(Here, we identify the regular graph G with its stochastic adjacency matrix.)

Proof. Item 1: The typical L_1 -norm of f_U evaluates to

$$\begin{aligned} \mathbb{E}_{U \sim V^{R-1}} \|f_U\|_1 &= \mathbb{E}_{U \sim V^{R-1}} \mathbb{P}_{W \sim G^{\otimes(R-1)}(U), Z \in V^{\varepsilon R}, x \sim V, \pi \in \mathbb{S}_{R'}} \left\{ F\left(\pi.\left((W, Z) +_i x\right)\right) = \pi(i) \right\} \\ &= \mathbb{P}_{\substack{X \sim V^{R'} \\ i \in \{1, \dots, R'\}}} \left\{ F(\pi.X) = \pi(i) \right\} = \frac{1}{R'} \mathbb{P}_{X \sim V^{R'}} \{F(X) \neq \perp\}. \end{aligned}$$

6. Graph Expansion and the Unique Games Conjecture

The second step uses that the joint distribution $(W, Z) +_i x$ and i is the same as the joint distribution of X and i . The last step uses that the distribution of $\pi(i)$ is uniformly random in $[R']$ even after conditioning on X and π (and thus $F(\pi.X)$).

Item 2: For fixed $U \in V^{R-1}$, the L_1 -norm of f_U evaluates to

$$\begin{aligned}
 \|f_U\|_1 &= \mathbb{P}_{\substack{W \sim G^{\otimes(R-1)}(U), Z \sim V^{\varepsilon R}, x \sim V, \pi \in \mathbb{S}_{R'} \\ i \in \{R, R+1, \dots, R'\}}} \left\{ F\left(\pi.\left((W, Z) +_i x\right)\right) = \pi(i) \right\} \\
 &= \mathbb{P}_{\substack{W \sim G^{\otimes(R-1)}(U), Z' \sim V^{\varepsilon R+1}, \pi \in \mathbb{S}_{R'} \\ i \in \{R, R+1, \dots, R'\}}} \left\{ F\left(\pi.\left(W, Z'\right)\right) = \pi(i) \right\} \\
 &= \frac{1}{\varepsilon R+1} \mathbb{P}_{W \sim G^{\otimes(R-1)}(U), Z' \sim V^{\varepsilon R+1}, \pi \in \mathbb{S}_{R'}} \left\{ \pi^{-1}\left(F\left(\pi.\left(W, Z'\right)\right)\right) \in \{R, R+1, \dots, R'\} \right\} \\
 &\leq \frac{1}{\varepsilon R+1}.
 \end{aligned}$$

In contrast to the proof of item 1, we insert x in a random coordinate among $\{R, R+1, \dots, R'\}$ (as opposed to completely random coordinate). The experiment as a whole does not change because π is a random permutation. The second step uses that $(i, (U, Z) +_i x)$ has the same distribution as $(i, (U, Z'))$.

Item 3: Sample tuples A, B, B' as specified by [Reduction 6.5](#). Note that B and B' are distributed independent and identically conditioned on A . We denote this conditional distribution by $B|A$. The fraction of constraints satisfied by F in \mathfrak{U} evaluates to

$$\begin{aligned}
 \mathfrak{U}(F) &= \sum_{r \in [R']} \mathbb{P}_{\substack{A, B, B' \\ \pi, \pi' \in \mathbb{S}_{R'}}} \left\{ F(\pi.B) = \pi(r) \wedge F(\pi'.B') = \pi'(r) \right\} \\
 &= \sum_{r \in [R']} \mathbb{E}_{A \sim V^R} \left(\mathbb{P}_{\substack{B|A \\ \pi \in \mathbb{S}_{R'}}} \left\{ F(\pi.B) = \pi(r) \right\} \right)^2.
 \end{aligned}$$

For every $A \in V^R$ and any $r, r' \in \{R+1, \dots, R'\}$, it holds that

$$\mathbb{P}_{\substack{B|A \\ \pi \in \mathbb{S}_{R'}}} \left\{ F(\pi.B) = \pi(r) \right\} = \mathbb{P}_{\substack{B|A \\ \pi \in \mathbb{S}_{R'}}} \left\{ F(\pi.B) = \pi(r') \right\}.$$

Here, we use that all coordinates b_r of B with $r \in \{R+1, \dots, R'\}$ are distributed identically (even conditioned on A). It follows that for every $r \in \{R+1, \dots, R'\}$,

$$\mathbb{P}_{\substack{B|A \\ \pi \in \mathbb{S}_{R'}}} \left\{ F(\pi.B) = \pi(r) \right\} \leq \frac{1}{\varepsilon R}.$$

6. Graph Expansion and the Unique Games Conjecture

Next, consider $r \in [R]$. For $B = (b_1, \dots, b_R) \in V^{R'}$, let $B^{-r} \in V^{R'-1}$ denote the tuple obtained from B by removing the r^{th} coordinate b_r . Then,

$$\mathbb{P}_{\substack{B|A, \\ \pi \in \mathbb{S}_{R'}}} \left\{ F(\pi \cdot B) = \pi(r) \right\} = \mathbb{P}_{\substack{B|A, \\ \pi \in \mathbb{S}_{R'}}} \left\{ F(\pi \cdot (B^{-r} +_r b_r)) = \pi(r) \right\} = \mathbb{E}_{B|A} f(B^{-r}, b_r).$$

Since $r \in [R]$, the vertex b_r is distributed as a random neighbor of a_r (the corresponding coordinate in A). It follows that

$$\mathbb{E}_{B|A} f(B^{-r}, b_r) = \mathbb{E}_{b_r \sim G(a_r)} f_{A^{-r}}(b_r) = G f_{A^{-r}}(a_r).$$

Recall that we identify the graph G with its stochastic adjacency matrix. Hence, $G f_{A^{-r}}$ denotes the function on V obtained by applying the linear operator G to the function $f_{A^{-r}}$. (Here, $A^{-r} \in V^{R-1}$ is the tuple obtained from A by removing the r^{th} coordinate a_r .) Combining the previous bounds, we get

$$\mathfrak{U}(F) \leq \sum_{r \in [R]} \mathbb{E}_{A \sim V^R} \left(G f_{A^{-r}}(a_r) \right)^2 + \varepsilon R \cdot \left(\frac{1}{\varepsilon R} \right)^2 = R \mathbb{E}_{U \sim V^{R-1}} \|G f_U\|^2 + \frac{1}{\varepsilon R},$$

which implies the desired bound on the typical squared L_2 -norm of $G f_U$. \square

The following lemma is a consequence of the previous lemma ([Lemma 6.9](#)) and a simple Markov-type inequality. The lemma shows that given a good partial assignment for unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ (as specified in [Reduction 6.5](#)), one can extract a function $f: V \rightarrow [0, 1]$ such that $\|f\|_1 \approx 1/R$ and, at the same time, the squared L_2 -norm of Gf is comparable to the L_1 -norm of f . In a later lemma ([Lemma 6.11](#)), we will see that given such a function we can find a non-expanding set for the graph G^2 with volume roughly $1/R$.

Lemma 6.10. *Let F be a partial assignment for the unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ (as specified in [Reduction 6.5](#)). Suppose $\alpha = \mathbb{P}_{X \sim V^{R'}} \{F(X) \neq \perp\}$ is the fraction of vertices of \mathfrak{U} labeled by F . Then, for every $\beta > 0$, there exists $U \in V^{R-1}$ such that the function $f_U: V \rightarrow [0, 1]$ (as defined in [\(6.1\)](#)) satisfies*

$$\begin{aligned} \|G f_U\|^2 &\geq (\mathfrak{U}(F)/\alpha - \beta - \frac{1}{\alpha \varepsilon R}) \|f_U\|_1, \\ \frac{\alpha \beta}{R'} &\leq \|f_U\|_1 \leq \frac{1}{\varepsilon R}. \end{aligned}$$

Proof. Since $\|G f_U\|^2 \leq \|f_U\|_1$ for every $U \in V^{R-1}$ (using that G is regular and $0 \leq f_U \leq 1$), we can lower bound the expected square L_2 -norm of $G f_U$

6. Graph Expansion and the Unique Games Conjecture

conditioned on $\|f_U\|_1 \geq \alpha\beta/R'$,

$$\begin{aligned} \mathbb{E}_{U \sim V^{R-1}} \|Gf_U\|^2 \mathbb{1}_{\{\|f_U\|_1 \geq \alpha\beta/R'\}} &= \mathbb{E}_{U \sim V^{R-1}} \|Gf_U\|^2 - \mathbb{E}_{U \sim V^{R-1}} \|Gf_U\|^2 \mathbb{1}_{\{\|f_U\|_1 \leq \alpha\beta/R'\}} \\ &\geq \mathbb{E}_{U \sim V^{R-1}} \|Gf_U\|^2 - \alpha\beta/R' \\ &\geq \frac{\alpha}{R'} \left(\lambda(F)/\alpha - \frac{1}{\alpha\epsilon R} - \beta \right). \end{aligned}$$

In the last step, we use that $\mathbb{E}_{U \sim V^{R-1}} \|Gf_U\|^2 \geq (\lambda(F) - \frac{1}{\epsilon R})/R'$ (Lemma 6.9, item 3). Since $\mathbb{E}_{U \sim V^{R-1}} \|f_U\|_1 = \alpha/R'$ (Lemma 6.9, item 1), there exists a tuple $U \in V^{R-1}$ such that

$$\|Gf_U\|^2 \mathbb{1}_{\{\|f_U\|_1 \geq \alpha\beta/R'\}} \geq \left(\lambda(F)/\alpha - \frac{1}{\alpha\epsilon R} - \beta \right) \|f_U\|_1.$$

This function f_U satisfies both $\|f_U\|_1 \geq \alpha\beta/R'$ and $\|Gf_U\|^2 \geq (\lambda(F)/\alpha - \beta - \frac{1}{\alpha\epsilon R})\|f_U\|_1$. (Note that we may assume $1 - \eta - \beta - \frac{1}{\alpha\epsilon R}$ is nonnegative, because otherwise the lemma is trivial.) On the other hand, f_U also satisfies $\|f_U\|_1 \leq 1/\epsilon R$ (by Lemma 6.9, item 2). \square

The spirit of the following lemma is similar to Cheeger's inequality. Given a function f on the vertex set V with values between 0 and 1, we can find a vertex set S with volume roughly $\|f\|_1$ and expansion roughly $1 - \|Gf\|^2/\|f\|_1$ in the graph G^2 . The proof of the lemma is much simpler than the construction for Cheeger's inequality. It is enough to analyze the distribution over vertex sets S obtained by including x in S with probability $f(x)$ independently for every vertex $x \in V$. Later, we will apply this lemma to the function obtained by the previous lemma (Lemma 6.10).

Lemma 6.11. *Suppose $f: V \rightarrow \mathbb{R}$ satisfies $0 \leq f(x) \leq 1$ for every vertex $x \in V$. Then, for every $\beta > 0$, there exists a set $S \subseteq V$ such that*

$$\begin{aligned} \beta \|f\|_1 &\leq \mu(S) \leq \frac{1}{\beta} \|f\|_1, \\ \Phi_{G^2}(S) &\leq 1 - \frac{\|Gf\|^2}{\|f\|_1} + 2\beta + \beta/(n\|f\|_1). \end{aligned}$$

Proof. Consider the following distribution over level sets $S \subseteq V$ of f : For every vertex $x \in V$, include x in S with probability $f(x)$ (independently for every vertex).

The expected volume of S is $\mathbb{E}_S \mu(S) = \mathbb{E}_S \|\mathbb{1}_S\|_1 = \|f\|_1$. The expectation of the square of the volume satisfies $\mathbb{E}_S \mu(S)^2 \leq \|f\|_1^2 + 1/n$. On the other hand, we can lower bound typical squared L_2 -norm of $G\mathbb{1}_S$ by $\mathbb{E}_S \|G\mathbb{1}_S\|^2 \geq \|Gf\|^2$.

6. Graph Expansion and the Unique Games Conjecture

Next, we lower bound the expected squared L_2 -norm of $G\mathbb{1}_S$ conditioned on the event $\beta \leq \mu(S)/\|f\|_1 \leq 1/\beta$,

$$\begin{aligned} \mathbb{E}_S \|G\mathbb{1}_S\|^2 \mathbb{1}_{\{\beta \leq \mu(S)/\|f\|_1 \leq 1/\beta\}} &\geq \mathbb{E}_S \|G\mathbb{1}_S\|^2 - \mathbb{E}_S \mu(S) \mathbb{1}_{\mu(S) > \|f\|_1/\beta} - \mathbb{E}_S \mu(S) \mathbb{1}_{\mu(S) < \beta \|f\|_1} \\ &\geq \|Gf\|^2 - \mathbb{E}_S \beta \mu(S)^2 / \|f\|_1 - \beta \|f\|_1 \\ &\geq \|Gf\|^2 - 2\beta \|f\|_1 - \beta / (n \|f\|_1). \end{aligned}$$

It follows that there exists a set $S^* \subseteq V$ with $\beta \leq \mu(S^*)/\|f\|_1 \leq 1/\beta$ such that

$$\frac{\|G\mathbb{1}_{S^*}\|^2}{\|\mathbb{1}_{S^*}\|_1} \geq \frac{\mathbb{E}_S \|G\mathbb{1}_S\|^2 \mathbb{1}_{\{\beta \leq \mu(S)/\|f\|_1 \leq 1/\beta\}}}{\mathbb{E}_S \|\mathbb{1}_S\|_1} \geq \frac{\|Gf\|^2 - 2\beta \|f\|_1 - \beta / (n \|f\|_1)}{\|f\|_1}.$$

The quantity $1 - \|G\mathbb{1}_{S^*}\|^2 / \|\mathbb{1}_{S^*}\|_1$ is the expansion of S^* in the graph G^2 . \square

If we combine the two previous lemmas ([Lemma 6.10](#) and [Lemma 6.11](#)), we can show that a good partial assignment for the unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ (obtained by applying [Reduction 6.5](#) to the graph G) implies that the graph G^2 contains a set with low expansion and volume roughly $1/R$.

Eventually, we want a vertex set with low expansion in the graph G . The following lemmas describe how to reduce this problem to the problem of finding vertex sets with low expansion in G^2 .

We first consider the case that there exists a partial assignment for the unique game \mathfrak{U} that satisfies a fraction of constraints bounded away from 0. In this case, we show how to construct a set with volume $1/R$ and expansion bounded away from 1 in G .

Lemma 6.12 (Soundness close to 0). *Suppose there exists a partial assignment F for the unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(G)$ (as defined in [Reduction 6.5](#)) with $\mathfrak{U}(F) \geq \eta\alpha$, where α is the fraction of vertices of \mathfrak{U} labeled by F . Then, there exists a set $S^* \subseteq V$ with $\mu(S^*) = \alpha/R$ and $\Phi(S^*) \leq 1 - O(\varepsilon\eta^4)$. Here, we make the (mild) assumptions that $\forall \alpha\varepsilon\eta \ll R \ll \alpha\eta n$.*

Proof. Let $\beta > 0$. (We determine the best choice for β later in the proof.) By [Lemma 6.10](#), there exists $U \in V^{R-1}$ such that the function f_U (as defined in [\(6.1\)](#)) satisfies the following conditions

$$\begin{aligned} \|Gf_U\|^2 &\geq (\eta - \beta - \frac{1}{\alpha\varepsilon R}) \|f_U\|_1, \\ \frac{\alpha\beta}{R'} &\leq \|f_U\|_1 \leq \frac{1}{\varepsilon R}. \end{aligned}$$

6. Graph Expansion and the Unique Games Conjecture

Using [Lemma 6.11](#) we round the function f_U to a vertex set $S \subseteq V$ with the following properties,

$$\frac{\alpha\beta^2}{R'} \leq \mu(S) \leq \frac{1}{\beta\epsilon R},$$

$$\Phi_{G^2}(S) \leq 1 - \eta + \beta + \frac{1}{\alpha\epsilon R} + 2\beta + \frac{R'}{n\alpha}.$$

If we choose β significantly smaller than η (say $\beta = \eta/10$) and use our assumptions on R , the expansion of S in G^2 is at most $1 - \eta/2$.

Next, we want to construct a set with expansion bounded away from 1 in the graph G and with roughly the same volume as S . Consider the set S' of all vertices x with $\mathbb{P}_{y \sim G(x)}\{y \in S\} \geq \gamma$. (We determine a good choice for $\gamma > 0$ later.) The volume of S' cannot be much larger than the volume of S . Concretely, $\mu(S') \leq \mu(S)/\gamma$. On the other hand, we can relate the fraction of edges between S and S' to the expansion of S in G^2 ,

$$\begin{aligned} (1 - \Phi_{G^2}(S))\mu(S) &= \mathbb{P}_{x \sim V, y, y' \sim G(x)}\{y \in S, y' \in S\} \\ &= \mu(S') \mathbb{P}_{x \sim V, y, y' \sim G(x)}\{y \in S, y' \in S \mid x \in S'\} \\ &\quad + \mu(V \setminus S') \mathbb{P}_{x \sim V, y, y' \sim G(x)}\{y \in S, y' \in S \mid x \notin S'\} \\ &\leq \mu(S') \mathbb{P}_{x \sim V, y \sim G(x)}\{y \in S \mid x \in S'\} \\ &\quad + \gamma \cdot \mu(V \setminus S') \mathbb{P}_{x \sim V, y \sim G(x)}\{y \in S \mid x \notin S'\} \\ &= (1 - \gamma)\mu(S') \mathbb{P}_{x \sim V, y \sim G(x)}\{y \in S \mid x \in S'\} \\ &\quad + \gamma \cdot \mu(S). \end{aligned}$$

It follows that $G(S, S') \geq (1 - \Phi_{G^2}(S) - \gamma)\mu(S)$. Therefore, for $S'' = S \cup S'$,

$$1 - \Phi_G(S'') \geq \frac{G(S, S')}{\mu(S \cup S')} \geq \left(\frac{\eta}{2} - \gamma\right) \cdot \frac{\mu(S)}{\mu(S \cup S')} \geq \left(\frac{\eta}{2} - \gamma\right) \cdot \frac{\gamma}{2}.$$

Choosing $\gamma = \eta/4$, we obtain $\Phi_G(S'') \leq 1 - \eta^2/32$. On the other hand, the volume of S'' satisfies

$$\Omega\left(\frac{\alpha\eta^2}{R}\right) \leq \mu(S'') \leq O\left(\frac{1}{\eta^2\epsilon R}\right)$$

To obtain a set S^* with the desired volume α/R , we either pad the set S'' with the desired number of vertices or we take a random subset of S'' with the desired cardinality. In either case, we obtain a set S^* with volume α/R and expansion at most $1 - O(\epsilon\eta^4)$. \square

6. Graph Expansion and the Unique Games Conjecture

In the next lemma, the goal is to find vertex sets in G with certain volume and expansion close to 0. It turns out that in this case it is convenient to apply [Reduction 6.5](#) not to G itself but to the graph $\frac{1}{2}I + \frac{1}{2}G$ (obtained by adding a self-loop of weight half to every vertex). A random step in the graph $\frac{1}{2}I + \frac{1}{2}G$ stays at the same vertex with probability $1/2$ and goes to a random neighbor in G with the remaining probability. The following lemma shows that if the unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(\frac{1}{2}I + \frac{1}{2}G)$ (obtained by applying [Reduction 6.5](#) to the graph $\frac{1}{2}I + \frac{1}{2}G$) has a partial assignment that satisfies almost as many constraints as possible, then we can find a vertex set in G with volume roughly $1/R$ and expansion close to 0. The proof is similar to the proof of the previous lemma ([Lemma 6.12](#)).

Lemma 6.13 (Soundness close to 1). *Suppose there exists a partial assignment x for the unique game $\mathfrak{U} = \mathfrak{U}_{R,\varepsilon}(\frac{1}{2}I + \frac{1}{2}G)$ (as defined in [Reduction 6.5](#)) with $\mathfrak{U}(x) \geq (1 - \eta)\alpha$, where α is the fraction of vertices of \mathfrak{U} labeled by x . Then, there exists a set $S \subseteq V$ with $\Omega(\alpha\eta^2/R) \leq \mu(S) \leq O(1/\varepsilon\eta R)$ and $\Phi(S) \leq 4\eta$. Here, we make the (mild) assumptions that $1/\alpha\varepsilon\eta \ll R \ll \alpha\eta n$.*

Proof. Let $\beta > 0$. (We choose this parameter later in the proof.) Let $G_{\cup} = \frac{1}{2}I + \frac{1}{2}G$. Combining [Lemma 6.10](#) and [Lemma 6.11](#), we obtain a set S with the following properties,

$$\begin{aligned} \frac{\alpha\beta^2}{R'} &\leq \mu(S) \leq \frac{1}{\beta\varepsilon R}, \\ \Phi_{G_{\cup}^2}(S) &\leq \eta + \beta + \frac{1}{\alpha\varepsilon R} + 2\beta + \frac{R'}{n\alpha}. \end{aligned}$$

If we choose β significantly smaller than η and use our assumptions on R , the expansion of S in G_{\cup}^2 is at most 2η . We compare the expansion of S in G_{\cup}^2 to its expansion in G ,

$$\langle \mathbf{1}_S, G_{\cup}^2 \mathbf{1}_S \rangle = \frac{1}{4} \langle \mathbf{1}_S, I \mathbf{1}_S \rangle + \frac{1}{2} \langle \mathbf{1}_S, G \mathbf{1}_S \rangle + \frac{1}{4} \langle \mathbf{1}_S, G^2 \mathbf{1}_S \rangle \leq \frac{1}{2} \|\mathbf{1}_S\|^2 + \frac{1}{2} \langle \mathbf{1}_S, G \mathbf{1}_S \rangle.$$

It follows that $\Phi_{G_{\cup}^2}(S) \geq \Phi_G(S)/2$, as desired. □

6.3. From PARTIAL UNIQUE GAMES to UNIQUE GAMES

Reduction 6.14 (From PARTIAL UNIQUE GAMES to UNIQUE GAMES).

Input: A unique game \mathfrak{U} with vertex set V and alphabet Σ , and a parameter $c \in \mathbb{N}$.

6. Graph Expansion and the Unique Games Conjecture

Output: A unique game $\mathfrak{U}' = \Psi_c(\mathfrak{U})$ with vertex set $V' = V^c$ and alphabet $\Sigma' = [c] \times \Sigma$.

The unique game \mathfrak{U}' corresponds to the following probabilistic verifier for an assignment $F: V' \rightarrow \Sigma'$:

1. Sample c random vertices $u_1, \dots, u_c \sim V$.
2. Sample two random constraints $(u_r, v_r, \pi_r), (u_r, v'_r, \pi'_r) \sim \mathfrak{U} \mid u_r$ for every $r \in [c]$. (Recall that $\mathfrak{U} \mid u_r$ denotes the uniform distribution over constraints of \mathfrak{U} that contain vertex u_r .)
3. Let $(r, j) = F(v_1, \dots, v_c)$ and $(r', j') = F(v'_1, \dots, v'_c)$.
4. Verify that $r = r'$ and that $j = \pi_r(i)$ and $j' = \pi'_{r'}(i)$ for some label $i \in \Sigma$. (Note that there can be at most one label $i \in \Sigma$ satisfying this condition.) (End of [Reduction 6.14](#))

[Reduction 6.14](#) has the following approximation guarantees.

Theorem 6.15. *Given a parameter $c \in \mathbb{N}$ and a unique game \mathfrak{U} with n vertices, [Reduction 6.14](#) computes in time $\text{poly}(n^c)$ a unique game $\mathfrak{U}' = \Psi_c(\mathfrak{U})$ such that the following assertions hold (for all $\alpha, \eta, \eta', \zeta > 0$):*

Completeness: *If the unique game \mathfrak{U} has α -partial value at least $\text{opt}_\alpha(\mathfrak{U}) \geq 1 - \eta$, then the unique game \mathfrak{U}' has value at least $\text{opt}(\mathfrak{U}') \geq 1 - 4\eta - 2e^{-\alpha c}$.*

Soundness I: *If the unique game \mathfrak{U} has $1/2c$ -partial value less than $\text{opt}_{1/2c}(\mathfrak{U}) < \zeta$, then the unique game \mathfrak{U}' has value less than $\text{opt}(\mathfrak{U}') < 8\zeta$.*

Soundness II: *If the unique game \mathfrak{U} has $1/2c$ -partial value less than $\text{opt}_{1/2c}(\mathfrak{U}) < 1 - \eta'$ and half of the constraints of every vertex are trivial identity constraints, then the unique game \mathfrak{U}' has value less than $\text{opt}(\mathfrak{U}') < 1 - \eta'/32$.*

The α -partial value of a unique game is the maximum value of an α -partial assignment normalized by the fraction of labeled vertices. See (2.1) in [Section 2.3](#) for the formal definition. We restate (2.1) below (out of context).

$$\text{opt}_\alpha(\mathfrak{U}) \stackrel{\text{def}}{=} \max \left\{ \frac{1}{\mathbb{P}_{u \sim V} \{x_u \neq \perp\}} \mathfrak{U}(x) \mid x \in (\Sigma \cup \{\perp\})^V, \mathbb{P}_{(u,v,\pi) \sim \mathfrak{U}} \{x_u \neq \perp\} \geq \alpha \right\}.$$

6. Graph Expansion and the Unique Games Conjecture

6.3.1. Completeness

Let \mathfrak{U} be a unique game with vertex set $V = \{1, \dots, n\}$ and alphabet $\Sigma = \{1, \dots, k\}$.

Recall that $\text{opt}_\alpha(\mathfrak{U})$ is the optimal value of an α -partial assignment for the unique game \mathfrak{U} . (See (2.1) in Section 2.3 for more details.)

The following lemma shows that Reduction 6.14 is complete, that is, given a unique game \mathfrak{U} with $\text{opt}_\alpha(\mathfrak{U}) \geq 1 - \eta$ for some constant α , then for $c = O(\log(1/\eta))$, the unique game $\mathfrak{U}' = \Psi_c(\mathfrak{U})$ obtained by applying Reduction 6.14 has value $\text{opt}(\mathfrak{U}') \geq 1 - O(\eta)$.

Lemma 6.16. *If the unique game \mathfrak{U} has α -partial value at least $\text{opt}_\alpha(\mathfrak{U}) \geq 1 - \eta$, then the unique game \mathfrak{U}' has value at least $\text{opt}(\mathfrak{U}') \geq 1 - 4\eta - 2e^{-\alpha c}$.*

Proof. Let $f: V \rightarrow \Sigma$ be an optimal α -partial assignment for \mathfrak{U} . We may assume $\mathbb{P}_{u \sim V} \{f(u) \neq \perp\} = \alpha$ and thus, $\mathfrak{U}(f) \geq (1 - \eta)\alpha$. To lower bound the value of the unique game \mathfrak{U}' , we consider the following partial assignment $F: V^c \rightarrow [c] \times \Sigma \cup \{\perp\}$,

$$F(u_1, \dots, u_c) = \begin{cases} (r, i) & \text{if } f(u_1) = \dots = f(u_{r-1}) = \perp \text{ and } f(u_r) = i \text{ for } r \in [c], \\ (1, 1) & \text{if } f(u_1) = \dots = f(u_c) = \perp. \end{cases}$$

In words, we determine the label of the vertex tuple (u_1, \dots, u_c) according to the first vertex u_r that is labeled in the assignment f . If none of the vertices u_r are labeled, we assign the dummy label \perp to the tuple.

We claim that this partial assignment F satisfies at least $1 - 4\eta - 2e^{-\alpha c}$ of the constraints of \mathfrak{U}' , which proves the lemma. (We could make this partial assignment into a total assignment by replacing \perp by an arbitrary label, say $(1, 1)$. This change can only increase the number of satisfied constraints.) To establish this claim, it is enough to show that

$$\mathbb{P}_{\substack{(u_1, v_1, \pi_1) \sim \mathfrak{U}, \\ \dots \\ (u_c, v_c, \pi_c) \sim \mathfrak{U}}} \left\{ \exists r \in [c], i \in \Sigma. F(u_1, \dots, u_c) = (r, i), F(v_1, \dots, v_c) = (r, \pi_r(i)) \right\} \geq 1 - 2\eta - e^{-\alpha c}. \quad (6.2)$$

(In fact, the unique game \mathfrak{U}' is the square of the unique game corresponding to the verifier above.)

For our assignment F , the probability in (6.2) simplifies to the product of

6. Graph Expansion and the Unique Games Conjecture

two probabilities p_1 and p_2 ,

$$\begin{aligned}
& \mathbb{P}_{\substack{(u_1, v_1, \pi_1) \sim \mathfrak{U}, \\ \dots \\ (u_c, v_c, \pi_c) \sim \mathfrak{U}}} \left\{ \exists r \in [c], i \in \Sigma. F(u_1, \dots, u_c) = (r, i), F(v_1, \dots, v_c) = (r, \pi_r(i)) \right\} \\
&= \mathbb{P}_{\substack{(u_1, v_1, \pi_1) \sim \mathfrak{U}, \\ \dots \\ (u_c, v_c, \pi_c) \sim \mathfrak{U}}} \left\{ \exists r \in [c]. f(u_r) \neq \perp \vee f(v_r) \neq \perp \right\} \\
&\quad \cdot \mathbb{P}_{(u, v, \pi) \sim \mathfrak{U}} \left\{ f(v) = \pi(f(u)) \mid f(u) \neq \perp \vee f(v) \neq \perp \right\} \\
&= p_1 \cdot p_2.
\end{aligned}$$

Since f is α -partial, the event $f(u_1) = \dots = f(u_c) = \perp$ has probability $(1 - \alpha)^c \leq e^{-\alpha c}$. Hence, with probability at least $1 - e^{-\alpha c}$, one of the vertices u_1, \dots, u_c is labeled. Therefore, $p_1 \geq 1 - e^{-\alpha c}$.

Furthermore, we can lower bound the second probability p_2 as follows

$$\begin{aligned}
p_2 &= \mathbb{P}_{(u, v, \pi) \sim \mathfrak{U}} \left\{ f(v) = \pi(f(u)) \mid f(u) \neq \perp \vee f(v) \neq \perp \right\} \\
&= \frac{\mathbb{P}_{(u, v, \pi) \sim \mathfrak{U}} \{f(v) = \pi(f(u))\}}{\mathbb{P}_{(u, v, \pi) \sim \mathfrak{U}} \{f(u) \neq \perp \vee f(v) \neq \perp\}} \\
&= \frac{\mathbb{P}_{(u, v, \pi) \sim \mathfrak{U}} \{f(v) = \pi(f(u))\}}{2\mathbb{P}_{u \sim V} \{f(u) \neq \perp\} - \mathbb{P}_{(u, v, \pi) \sim \mathfrak{U}} \{f(u) \in \Sigma, f(v) \in \Sigma\}} \\
&\geq \frac{(1 - \eta)\alpha}{2\alpha - (1 - \eta)\alpha} \geq 1 - 2\eta.
\end{aligned}$$

We conclude that the probability in (6.2) is at least $p_1 \cdot p_2 \geq 1 - 2\eta - e^{-\alpha c}$, as claimed. \square

6.3.2. Soundness

In the following, we show that [Reduction 6.14](#) (from PARTIAL UNIQUE GAMES to UNIQUE GAMES) is sound, that is, given a good assignment for the unique game $\mathfrak{U}' = \Psi_c(\mathfrak{U})$ (obtained by applying [Reduction 6.14](#) to the unique game \mathfrak{U}), one can construct a good partial assignment for the original unique game \mathfrak{U} .

Let $F: V^c \rightarrow [c] \times \Sigma$ be an assignment for the unique game $\mathfrak{U}' = \Psi_c(\mathfrak{U})$. (Recall that \mathfrak{U} is a unique game with vertex set $V = \{1, \dots, n\}$ and alphabet $\Sigma = \{1, \dots, k\}$.)

6. Graph Expansion and the Unique Games Conjecture

Based on the assignment F , we construct a collection $\{f_{U,r}\}$ of partial assignments for \mathfrak{U} . For $U \in V^{c-1}$ and $r \in \{1, \dots, c\}$, define $f_{U,r}: V \rightarrow \Sigma \cup \{\perp\}$ by

$$f_{U,r}(x) \stackrel{\text{def}}{=} \begin{cases} i & \text{if } F(U +_r x) = (r, i), \\ \perp & \text{otherwise.} \end{cases}$$

(Here, $U +_r x$ denotes the tuple $(u_1, \dots, u_r, x, u_{r+1}, \dots, u_{c-1})$ in V^c obtained by inserting u in the r -th coordinate of $U = (u_1, \dots, u_{c-1})$.)

We first show that one of the partial assignments $f_{U,r}$ has good value in the unique game \mathfrak{U}^2 (the square of \mathfrak{U} , obtained by sampling two constraints of \mathfrak{U} with a common vertex and composing the permutations).

Lemma 6.17. *Let $\beta > 0$. Then, there exists a partial assignment $f_{U,r}$ that labels at least β/c of the vertices and has value at least*

$$\mathfrak{U}^2(f_{U,r}) \geq \frac{\mathfrak{U}'(F) - \beta}{1 - \beta} \cdot \mathbb{P}_{x \sim V} \{f_{U,r}(x) \neq \perp\},$$

where \mathfrak{U}^2 denotes the square of the unique game \mathfrak{U} (obtained by sampling two constraints of \mathfrak{U} with a common vertex and composing the permutations). In particular, $\text{opt}_{\beta/c}(\mathfrak{U}^2) \geq (\mathfrak{U}(F) - \beta)/(1 - \beta)$.

We compute the typical value of $f_{U,r}$ for the square of the unique game \mathfrak{U} (obtained by sampling two constraints from \mathfrak{U} with a common vertex and composing the permutations),

$$\begin{aligned} & \mathbb{E}_{U \sim V^{c-1}, r \in [c]} \mathfrak{U}^2(f_{U,r}) \\ &= \mathbb{P}_{U \sim V^{c-1}, r \in [c], u \sim V, (u,v,\pi), (u,v',\pi') \sim \mathfrak{U}|u} \left\{ \pi^{-1}(f_{U,r}(v)) = (\pi')^{-1}(f_{U,r}(v')) \right\} \\ &= \frac{1}{c} \mathfrak{U}'(F). \end{aligned}$$

For $U \in V^{c-1}$ and $r \in [c]$, define $\alpha_{U,r}$ as the fraction of vertices labeled by the partial assignment $f_{U,r}$,

$$\alpha_{U,r} \stackrel{\text{def}}{=} \mathbb{P}_{x \sim V} \{f_{U,r}(x) \neq \perp\}.$$

The typical value of $\alpha_{U,r}$ is $1/c$,

$$\mathbb{E}_{U \sim V^{c-1}, r \in [c]} \alpha_{U,r} = \mathbb{P}_{U \sim V^{c-1}, r \in [c], x \sim V} \{f_{U,r}(x) \neq \perp\} = \frac{1}{c}.$$

6. Graph Expansion and the Unique Games Conjecture

Since $\mathfrak{U}^2(f_{U,r}) \leq \alpha_{U,r}$, the contribution of assignments $f_{U,r}$ with $\alpha_{U,r} \leq \beta/c$ to the expected value $\mathfrak{U}(f_{U,r})$ is at most

$$\mathbb{E}_{U \in V^{c-1}, r \in [c]} \mathfrak{U}^2(f_{U,r}) \cdot \mathbb{1}_{\alpha_{U,r} \leq \beta/c} \leq \beta/c.$$

Therefore, we can lower bound the β/c -partial value of the unique game \mathfrak{U}^2 as follows

$$\text{opt}_{\beta/c}(\mathfrak{U}) \geq \frac{\mathbb{E}_{U \sim V^{c-1}, r \in [c]} \mathfrak{U}^2(f_{U,r}) \cdot \mathbb{1}_{\alpha_{U,r} \geq \beta/c}}{\mathbb{E}_{U \sim V^{c-1}, r \in [c]} \alpha_{U,r} \cdot \mathbb{1}_{\alpha_{U,r} \geq \beta/c}} \geq \frac{\mathfrak{U}'(F)/c - \beta/c}{1/c - \beta/c} = \frac{\mathfrak{U}'(F) - \beta}{1 - \beta}.$$

Using the previous lemma ([Lemma 6.17](#)) and basic relations between the optimal value of a unique game and its square, we can show the following lemmas.

Lemma 6.18 (Soundness close to 0). *Let $\mathfrak{U}' = \Psi_c(\mathfrak{U})$ be the unique game obtained by applying [Reduction 6.14](#) with parameter $c \in \mathbb{N}$ to the unique game \mathfrak{U} . Suppose the value of \mathfrak{U}' is at least $\text{opt}(\mathfrak{U}') \geq \zeta$. Then, the $1/2c$ -partial value of \mathfrak{U} satisfies $\text{opt}_{1/2c}(\mathfrak{U}) \geq \zeta/8$.*

Let \mathfrak{U}_\cup be the unique game obtained by sampling with probability $1/2$ a random constraint from \mathfrak{U} and sampling with the remaining probability a trivial constraint (u, u, id) . Note that $\mathfrak{U}_\cup(f) = 1/2(\alpha + \mathfrak{U}(f))$ for every assignment $f: V \rightarrow \Sigma \cup \{\perp\}$ that labels an α fraction of the vertices.

Lemma 6.19 (Soundness close to 1). *Let $\mathfrak{U}' = \Psi_c(\mathfrak{U}_\cup)$ be the unique game obtained by applying [Reduction 6.14](#) with parameter $c \in \mathbb{N}$ to the unique game \mathfrak{U}_\cup . Suppose the value of \mathfrak{U}' is at least $\text{opt}(\mathfrak{U}') \geq 1 - \eta$. Then, the $1/2c$ -partial value of \mathfrak{U} satisfies $\text{opt}_{1/2c}(\mathfrak{U}) \geq 1 - \eta/16$.*

6.4. From UNIQUE GAMES to SMALL-SET EXPANSION

Reduction 6.20 (Reduction from UNIQUE GAMES to SMALL-SET EXPANSION).

Input: A unique game \mathfrak{U} with vertex set V and alphabet $\Sigma = \{1, \dots, R\}$, and parameters $q \in \mathbb{N}$ and $\varepsilon > 0$.

Output: A graph $H = H_{q,\varepsilon}(\mathfrak{U})$ with vertex set $V \times [q]^R$.

The edge distribution of H is obtained from the unique game \mathfrak{U} in the following way:

6. Graph Expansion and the Unique Games Conjecture

1. Sample a random vertex $u \sim V$.
2. Sample a random point $x \in [q]^R$.
3. Sample two random constraints $(u, v, \pi), (u, v', \pi') \sim \mathfrak{U} \mid u$. (Here, $\mathfrak{U} \mid u$ denotes the uniform distribution over constraints of \mathfrak{U} containing vertex u .)
4. Sample two random points $y, y' \sim T_{1-\varepsilon}(x)$. (Here, $T_{1-\varepsilon}$ is the usual noise graph on $[q]^R$ with noise parameter $1 - \varepsilon$. See for a detailed definition.)
5. Output an edge between (v, y) and (v', y') .

(End of [Reduction 6.20](#))

The running time of [Reduction 6.20](#) is $\text{poly}(|V|, q^R)$. In particular, the running time is polynomial if $q = O(1)$ and $R = O(\log n)$.

The reduction [Reduction 6.20](#) has the following approximation guarantee.

Theorem 6.21. *For every $\varepsilon > 0$ and $q \in \mathbb{N}$, there exist $\eta_\varepsilon > 0$ and $\zeta_{q,\varepsilon} > 0$ such that if $H = H_{q,\varepsilon}(\mathfrak{U})$ is the graph obtained by applying [Reduction 6.20](#) to a unique game \mathfrak{U} , then the following two assertions hold:*

Completeness: *If the unique game \mathfrak{U} has value $\text{opt}(\mathfrak{U}) \geq 1 - \varepsilon$, then the graph H contains a vertex set with volume $1/q$ and expansion at most 4ε .*

Soundness: *If the unique game \mathfrak{U} has value $\text{opt}(\mathfrak{U}) \leq \zeta_{q,\varepsilon}$ and every vertex set of volume $1/q$ in the constraint graph of \mathfrak{U} has expansion at least $1 - \eta_\varepsilon$, then every vertex set of volume $1/q$ has expansion $1 - \varepsilon$ in H .*

Remark 6.22. The completeness assertion of [Reduction 6.20](#) is in fact stronger than stated in [Theorem 6.21](#). If the unique game \mathfrak{U} has value $\text{opt}(\mathfrak{U}) \geq 1 - \varepsilon$, then the graph $H = H_{q,\varepsilon}$ allows a partition into q vertex sets, each of volume $1/q$ and with expansion at most 4ε .

6.4.1. Completeness

In the following, we show that [Reduction 6.20](#) is complete, that is, we show that if the unique game \mathfrak{U} has a good assignment (value $1 - \eta$), then the graph $H = H_{q,\varepsilon}(\mathfrak{U})$ obtained by [Reduction 6.20](#) contains a set with volume $1/q$ and expansion $O(\varepsilon + \eta)$.

6. Graph Expansion and the Unique Games Conjecture

Lemma 6.23 (Completeness of [Reduction 6.20](#)). *Suppose the unique game \mathfrak{U} has optimal value $\text{opt}(\mathfrak{U}) \geq 1 - \eta$. Then, the graph $H = H_{q,\varepsilon}(\mathfrak{U})$ obtained from [Reduction 6.20](#) contains a vertex set S with volume $\mu(S) = 1/q$ and expansion $\Phi(S) \leq 2(\eta + \varepsilon)$.*

Proof. Let $F: V \rightarrow [R]$ be an assignment for the unique game \mathfrak{U} . Suppose the assignment F has value $\mathfrak{U}(F) = 1 - \eta$. Consider the vertex subset $S \subseteq V \times [q]^R$ in the graph $H = H_{q,\varepsilon}(\mathfrak{U})$ (as specified in [Reduction 6.20](#)),

$$S \stackrel{\text{def}}{=} \{(u, x) \in V \times [q]^R \mid x_{F(u)} = 1\}.$$

(We remark that the choice of 1 is arbitrary. The following proof would work for any element of $[q]$.)

The volume of S in the graph H equals $\mu(S) = 1/q$. To determine the expansion of S , we compute $H(S, S)$, the fraction of edge of H inside of S . Let $(u, x), (v, y), \pi, (v', y'), \pi'$ be distributed as in [Reduction 6.20](#). Then,

$$\begin{aligned} H(S, S) &= \mathbb{P}_{(u,x), (v,y), \pi, (v',y'), \pi'} \left\{ (\pi \cdot y)_{F(v)} = 1 \wedge (\pi' \cdot y')_{F(v')} = 1 \right\} \\ &\geq (1 - \varepsilon)^2 \mathbb{P}_{(u,x), (v,y), \pi, (v',y'), \pi'} \left\{ (\pi \cdot x)_{F(v)} = 1 \wedge (\pi' \cdot x)_{F(v')} = 1 \right\} \\ &\geq (1 - \varepsilon)^2 \mathbb{P}_{(u,x), (v,y), \pi, (v',y'), \pi'} \left\{ x_{F(u)} = 1 \wedge F(v) = \pi(F(u)) \wedge F(v') = \pi'(F(u)) \right\} \\ &= (1 - \varepsilon)^2 \mathfrak{U}^2(F)/q \geq (1 - \varepsilon)^2 (1 - 2\eta)/q. \end{aligned}$$

The first inequality uses that $\mathbb{P}_{y \sim T_\rho(x)} \{y_r = x_r\} \geq \rho$ for every coordinate $r \in [R]$. For the second inequality, we use that one event logically implies the other event. The notation $\mathfrak{U}^2(F)$ denotes the value of the assignment in the square of the unique game \mathfrak{U} . (The square of a unique game is obtained by sampling two random constraints with a common vertex and composing the permutations.)

We conclude that, as desired, S satisfies

$$\begin{aligned} \mu(S) &= \frac{1}{q}, \\ \Phi(S) &\leq 2\varepsilon + 2\eta. \end{aligned} \quad \square$$

6.4.2. Soundness

In the following, we show that [Reduction 6.20](#) is sound, that is, if the graph $H = H_{q,\varepsilon}(\mathfrak{U})$ contains a vertex set S with small volume and expansion bounded

6. Graph Expansion and the Unique Games Conjecture

away from 1, then either there exists a good assignment for the unique game \mathfrak{U} or the constraint graph of \mathfrak{U} contains a vertex with small volume and expansion bounded away from 1. (See [Lemma 6.24](#) at the end of this section for the formal statement.)

Let $S \subseteq V \times [q]^R$ be a vertex subset of the graph $H = H_{q,\varepsilon}(\mathfrak{U})$ obtained by applying [Reduction 6.20](#) to a unique game \mathfrak{U} with vertex set V and alphabet $[R]$. Let $f = \mathbb{1}_S$ be the indicator function of S .

For a vertex $u \in V$, we define $f_u: [q]^R \rightarrow \{0, 1\}$ by $f_u(x) := f(u, x)$. We define $g_u: [q]^R \rightarrow [0, 1]$ by averaging the functions f_v over the constraints $(u, v, \pi) \sim \mathfrak{U} \mid u$,

$$g_u(x) := \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U} \mid u} f_v(\pi \cdot x).$$

(For a constraint (u, v, π) , we identify the i^{th} input coordinate of g_u with the $\pi(i)^{\text{th}}$ input coordinate of f_v .)

The typical noise stability of the functions g_u equals $H(S, S)$, the fraction of edges staying inside of S ,

$$\begin{aligned} H(S, S) &= \mathbb{E}_{u \sim V, x \in [q]^R} \left(\mathbb{P}_{\substack{(u,v,\pi) \sim \mathfrak{U} \mid u \\ y \sim T_{1-\varepsilon}(x)}} \{(v, y) \in S\} \right)^2 \\ &= \mathbb{E}_{u \sim V, x \in [q]^R} T_{1-\varepsilon} g_u(x)^2 = \mathbb{E}_{u \sim V} \|T_{1-\varepsilon} g_u\|^2. \end{aligned}$$

The invariance principle ([Theorem 2.10](#)) allows us to estimate the noise stability $\|T_{1-\varepsilon} g_u\|^2 = \langle g_u, T_{(1-\varepsilon)^2} g_u \rangle$ in terms of Gaussian noise stability bounds, whenever the function g_u has no influential coordinates. For the reader's convenience, we restate the invariance principle.

Theorem (Restatement of [Theorem 2.10](#)). *For every finite probability space Ω and constants $\rho \in [0, 1)$, $\eta > 0$, there exists constants $\tau, \gamma > 0$ such that for every function $f \in L_2(\Omega^R)$ with $0 \leq f \leq 1$, either*

$$\langle f, T_\rho f \rangle \leq \Gamma_\rho(\mathbb{E} f) + \eta,$$

or $\text{Inf}_i T_{1-\gamma} f > \tau$ for some coordinate $i \in [R]$.

Let $\eta > 0$ be sufficiently small (we determine this parameter later more precisely). By the invariance principle, there exists constants $\tau, \gamma > 0$ (depending on q, ε , and η) such that

$$H(S, S) = \mathbb{E}_{u \sim V} \|T_{1-\varepsilon} g_u\|^2 \leq \mathbb{E}_{u \sim V} \Gamma_{(1-\varepsilon)^2}(\mathbb{E} g_u) + \eta + \mathbb{P}_{v \sim V} \left\{ \max_{i \in [R]} \text{Inf}_i T_{1-\gamma} g_u > \tau \right\}.$$

6. Graph Expansion and the Unique Games Conjecture

(Here, we also use that $\|T_{1-\varepsilon}g_u\|^2 \leq 1$ since $0 \leq g_u \leq 1$.)

We use the following basic bound for the Gaussian noise stability. Let $\delta > 0$ be sufficiently small (we determine this parameter later). Then,

$$\Gamma_{(1-\varepsilon)^2}(\mathbb{E}g_u) \leq \mathbb{E}g_u \cdot (\delta^{\varepsilon/2} + \mathbb{1}_{\mathbb{E}g_u > \delta}).$$

(If $\mathbb{E}g_u \leq \delta$, then we can bound the noise stability by $\delta^{\varepsilon/2} \mathbb{E}g_u$. Otherwise, we can use the trivial bound $\Gamma_{(1-\varepsilon)^2}(\mathbb{E}g_u) \leq \mathbb{E}g_u$.)

Influence decoding ([Lemma 2.11](#)) allows us to estimate the fraction of vertices u such that $T_{1-\gamma}g_u$ has a coordinate with influence larger than τ .

Lemma (Restatement of [Lemma 2.11](#)). *Let \mathfrak{U} be a unique game with vertex set V and alphabet $[R]$. For some probability space Ω , let $\{f_u\}_{u \in V}$ be a collection of normalized functions in $L_2(\Omega^R)$. Consider functions g_u in $L_2(\Omega^R)$ defined by*

$$g_u(x) = \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}|u} f_v(\pi.x).$$

Then, for all $\gamma, \tau > 0$, there exists $c_{\gamma,\tau} > 0$ (in fact, $c_{\gamma,\tau} = \text{poly}(\gamma, \tau)$) such that

$$\text{opt}(\mathfrak{U}) \geq c_{\gamma,\tau} \cdot \mathbb{P}_{u \sim V} \left\{ \max_{i \in [R]} \text{Inf}_i T_{1-\gamma}g_u > \tau \right\}.$$

Combining Gaussian noise stability bounds and influence decoding, our previous upper bound on $H(S, S)$ simplifies to

$$H(S, S) \leq \delta^{\varepsilon/2} \cdot \mu(S) + \eta + \text{opt}(\mathfrak{U})/c_{\gamma,\tau} + \mathbb{E}_{u \sim V} \mathbb{1}_{\mathbb{E}g_u > \delta} \mathbb{E}g_u.$$

This bound asserts the following: If $H(S, S)$ is bounded away from $\delta^{\varepsilon/2} \cdot \mu(S) + \eta$ (meaning the expansion S is bounded away from 1), then either $\text{opt}(\mathfrak{U})$ is bounded away from 0 or $\mathbb{E}_{u \sim V} \mathbb{1}_{\mathbb{E}g_u > \delta} \mathbb{E}g_u$ is bounded away from 0. To show the soundness of [Reduction 6.20](#), it remains to argue that the latter condition implies that the constraint graph of the unique game \mathfrak{U} contains a vertex set with small volume and expansion bounded away from 1.

To this end, we consider functions $\tilde{f}, \tilde{g}: V \rightarrow [0, 1]$, obtained by averaging the functions f_u and g_u over the hypercube $[q]^R$,

$$\begin{aligned} \tilde{f}(u) &= \mathbb{E}_x f_u(x), \\ \tilde{g}(u) &= \mathbb{E}_x g_u(x). \end{aligned}$$

6. Graph Expansion and the Unique Games Conjecture

The functions \bar{f} and \bar{g} satisfy the relation $\bar{g} = G\bar{f}$, where G is the constraint graph of the unique game \mathfrak{U} , because

$$\begin{aligned} G\bar{f}(u) &= \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}|u} \bar{f}(v) \\ &= \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}|u} \mathbb{E}_{x \in [q]^R} f_v(x) \\ &= \mathbb{E}_{x \in [q]^R} \mathbb{E}_{(u,v,\pi) \sim \mathfrak{U}|u} f_v(\pi.x) \\ &= \mathbb{E}_{x \in [q]^R} g_u(x) = \bar{g}(u). \end{aligned}$$

We can estimate $\mathbb{E}_{u \sim V} \mathbb{1}_{\mathbb{E}g_u > \delta} \mathbb{E}g_u$ in terms of the L_2 -norm of $\bar{g} = G\bar{f}$,

$$\mathbb{E}_{u \sim V} \mathbb{1}_{\mathbb{E}g_u > \delta} \mathbb{E}g_u \leq \mathbb{E}_u \bar{g}(u) \mathbb{1}_{\bar{g}(u) > \delta} \leq \mathbb{E}_u \bar{g}(u)^2 / \delta = \|G\bar{f}\|^2 / \delta.$$

At this point, we can use the following lemma, which we used in the soundness analysis of [Reduction 6.20](#) earlier in this chapter.

Lemma (Restatement of [Lemma 6.11](#)). *Suppose $f : V \rightarrow \mathbb{R}$ satisfies $0 \leq f(x) \leq 1$ for every vertex $x \in V$. Then, for every $\beta > 0$, there exists a set $S \subseteq V$ such that*

$$\begin{aligned} \beta \|f\|_1 &\leq \mu(S) \leq \frac{1}{\beta} \|f\|_1, \\ \Phi_{G^2}(S) &\leq 1 - \frac{\|Gf\|^2}{\|f\|_1^2} + 2\beta + \beta/(n\|f\|_1). \end{aligned}$$

Let $\beta > 0$ be sufficiently small (we determine this parameter later) and let $S' \subseteq V$ be the vertex set obtained by applying [Lemma 6.11](#) to the function \bar{f} . The set S' satisfies

$$\begin{aligned} \beta \mu(S) &\leq \mu_G(S') \leq \mu(S) / \beta \\ \Phi_{G^2}(S') &\leq 1 - \|G\bar{f}\|^2 / \mu(S) + 2\beta + \beta / (n\mu(S)). \end{aligned}$$

We arrive at the following upper bound on the fraction of edges staying inside the vertex set S (normalized by its volume),

$$\frac{H(S, S)}{\mu(S)} \leq \delta^{\varepsilon/2} + \frac{\eta}{\mu(S)} + \frac{1}{c_{\gamma, \tau} \mu(S)} \text{opt}(\mathfrak{U}) + \frac{1}{\delta} \cdot \frac{G^2(S', S')}{\mu(S')} + 2\beta + \frac{\beta}{n\mu(S)}$$

To simplify this bound further, we eliminate a few parameters. Choose δ such that $\delta^{\varepsilon/2} \ll \beta$ and choose η such that $\eta \ll \beta \mu(S)$. Furthermore, we assume

6. Graph Expansion and the Unique Games Conjecture

that $\mu(S) \gg 1/n$ (a very mild assumption). With these parameter choices, the bound on $H(S, S)$ simplifies to

$$\frac{H(S, S)}{\mu(S)} \leq 4\beta + C_{\beta, q, \varepsilon, \mu(S)} \text{opt}(\mathfrak{U}) + \beta^{-2/\varepsilon} \cdot \frac{G^2(S', S')}{\mu(S')},$$

where $C_{\beta, q, \varepsilon}$ is some constant depending only on β , q , $\mu(S)$, and ε .

The last bound on $H(S, S)$ implies the following lemma asserting the soundness of [Reduction 6.20](#). (Here, we also use the fact that a vertex set with expansion bounded away from 1 in G^2 implies a vertex set with roughly the same volume and expansion bounded away from 1 in G .)

Lemma 6.24. *For every $q \in \mathbb{N}$ and $\varepsilon, \eta > 0$, there exists $\zeta_{q, \varepsilon, \eta}, \zeta'_{\varepsilon, \eta} > 0$ such that the following holds: Let $H = H_{q, \varepsilon}(\mathfrak{U})$ be the graph obtained by applying [Reduction 6.20](#) with parameters q and ε to the unique game \mathfrak{U} . Then, if H contains a vertex set S with volume $\mu_H(S) = 1/q$ and expansion at most $\Phi_H(S) \leq 1 - \eta$, then either $\text{opt}(\mathfrak{U}) > \zeta_{q, \varepsilon, \eta}$ or the constraint graph G of \mathfrak{U} contains a vertex set S' with volume $\mu_G(S') = 1/q$ and expansion at most $\Phi_G(S') \leq 1 - \zeta'_{\varepsilon, \eta}$.*

6.5. Notes

The material presented in this chapter is based on the paper “Graph Expansion and the Unique Games Conjecture” [[RS10](#)] joint with Prasad Raghavendra. A preliminary version appeared at STOC 2010.

7. Reductions between Expansion Problems

We show that the SSE hypothesis implies quantitatively tight inapproximability results for many graph expansion problems, in particular `SPARSEST CUT` and `BALANCED SEPARATOR`. Concretely, our results imply that it is SSE-hard to beat Cheeger’s inequality and achieve an $(\varepsilon, o(\sqrt{\varepsilon}))$ -approximation for `SPARSEST CUT`. (We say a problem is SSE-hard if an efficient algorithm for the problem implies that the SSE hypothesis is false.)

These are the first strong hardness of approximation results for `SPARSEST CUT` and `BALANCED SEPARATOR`. (Even assuming the Unique Games Conjecture, no strong hardness result for these problems were known. The best known result was that one cannot achieve approximation ratios arbitrarily close to 1 unless 3-SAT has subexponential algorithms [AMS07].)

There is a strong parallel between this result and the known consequences of the Unique Games Conjecture. The Unique Games Conjecture asserts a qualitative inapproximability for `UNIQUE GAMES`, which is generalization of `MAX CUT`. In turn, the conjecture implies tight quantitative inapproximability results for `MAX CUT` and similar basic problems. Similarly, the SSE hypothesis asserts a qualitative inapproximability for `SMALL-SET EXPANSION`, which generalizes `SPARSEST CUT`. In turn, the hypothesis implies tight quantitative inapproximability results for `SPARSEST CUT` and other graph expansion problems. The value of results of this kind is that they unify the question of improving known algorithms for a class of problems to a question about the qualitative approximability of a single problem.

These significant consequences of a confirmation of the SSE hypothesis make it an interesting open question to prove or refute the hypothesis.

7.1. Main Results

We show that following general graph expansion problem is SSE-hard.

7. Reductions between Expansion Problems

Theorem 7.1. *The following problem is SSE-hard for all constants $\varepsilon, \delta, \eta > 0$: Given a graph G distinguish between the cases,*

YES: *the graph G contains a vertex set with volume at most δ and expansion at most $\varepsilon + \eta$,*

NO: *every vertex set S in G satisfies*

$$G(S, S) \leq \Gamma_{1-\varepsilon}(\mu(S)) + \eta.$$

A corollary of this theorem is an essentially tight SSE-hardness for SPARSEST CUT. Concretely, given a graph G , it is SSE-hard for every $\varepsilon > 0$ to distinguish between the cases $\Phi_G \leq \varepsilon$ and $\Phi_G > \Omega(\sqrt{\varepsilon})$. This hardness is essentially tight because Cheeger's inequality allows us to distinguish between $\Phi_G \leq \varepsilon$ and $\Phi_G > O(\sqrt{\varepsilon})$. (However, the constants hidden in the $\Omega(\cdot)$ - and $O(\cdot)$ -notations are not the same.)

The YES case in the above theorem can be strengthened in various ways (leading to stronger hardness results). For example, we can assert that there exists a distribution over vertex sets S such that

- with probability 1, the vertex set S has volume at most δ and expansion at most ε ,
- the distribution over vertex sets S uniformly covers essentially the whole graph, in the sense that,

$$\mathbb{E}_{v \sim G} \left| \mathbb{P}_S \{v \in S\} - \mathbb{E}_S \mu(S) \right| \leq \eta \mathbb{E}_S \mu(S).$$

(The volume of S is also lower bounded by a function of η and δ .)

Such a distribution over vertex sets implies the existence of SDP solutions with certain symmetry properties, leading to quantitatively tight integrality gaps for BALANCED SEPARATOR. (Using techniques in [Chapter 8](#).)

Another consequence of such distributions is that one can extract vertex sets with desired volume while maintaining the expansion up to a factor of 2. (To extract the vertex set with the desired volume we use correlated sampling, similar to [§4.2](#) in [Chapter 4](#).) This observation implies the corollary of the previous theorem.

Theorem 7.2. *The following problem is SSE-hard for all constants $\varepsilon, \delta, \eta > 0$: Given a graph G distinguish between the cases,*

7. Reductions between Expansion Problems

YES: the graph G allows a partition of the vertex set into sets with volume equal to δ and expansion at most $2\varepsilon + \eta$,

NO: every vertex set S in G satisfies

$$G(S, S) \leq \Gamma_{1-\varepsilon}(\mu(S)) + \eta.$$

7.2. Reduction

The reduction in this section composes the input graph with certain other graphs, in particular, noise graphs (see [Section 2.4](#) for details about noise graphs). Beyond noise graphs, the reduction uses the following kind of graphs (which are related to noise graphs).

Let Ω be a finite probability space and let $R \in \mathbb{N}$. For $w \in \{\perp, \top\}^R$, let M_w be the Markov operator on $L_2(\Omega^R)$ corresponding to the graph on Ω^R with the following edge distribution:

1. Sample $x \sim \Omega^R$.
2. For every coordinate $r \in [R]$, put $y_r := x_r$ if $w_r = \top$ and sample $y_r \sim \Omega$ if $w_r = \perp$.
3. Output the pair xy as a random edge of M_w .

(The graph M_w also has a simple combinatorial description. It is a disjoint union of cliques and two vertices $x, y \in \Omega^R$ are in the same clique if they agree on the coordinates r with $w_r = \top$.)

Let $\{\perp, \top\}_\beta^R$ be R -fold product of the β -biased distribution on $\{\perp, \top\}$. In other words, $\{\perp, \top\}_\beta^R$ is the distribution over $\{\perp, \top\}^R$ in which each coordinate equals \top with probability β (independently for each coordinate).

Averaging the graph M_w over $w \sim \{\perp, \top\}_\beta^R$ yields the noise graph T_β on Ω^R ,

$$\mathbb{E}_{w \sim \{\perp, \top\}_\beta^R} M_w = T_\beta. \tag{7.1}$$

Reduction 7.3.

Input: A regular graph G with vertex set V , and parameters $R, q \in \mathbb{N}$ and $\varepsilon, \rho, \beta \in (0, 1)$.

7. Reductions between Expansion Problems

Output: A regular graph $H = H_{R,q,\varepsilon,\rho,\beta}(G)$ with vertex set $V(H) = V^R \times [q]^R \times \{\perp, \top\}_\beta^R$.

The edge distribution of the graph H is obtained from G in the following way:

1. Sample a random R -tuple $A \sim V^R$ of vertices.
2. Sample two random neighbors $B, B' \sim (T_{1-\varepsilon}G)^{\otimes R}(A)$ of A in the graph $(T_{1-\varepsilon}G)^{\otimes R}$. (Here, the graph $T_{1-\varepsilon}G$ corresponds to first taking a random step in G and then taking a random step in $T_{1-\varepsilon}$, the usual noise graph on V with parameter $1 - \varepsilon$. See for further details.)
3. Sample two random permutations $\pi, \pi' \in \mathbb{S}_R$.
4. Sample a random point $(x, z) \in [q]^R \times \{\perp, \top\}_\beta^R$.
5. Sample a random neighbors $(y, w), (y', w') \sim T_\rho^{\otimes R}(x, z)$ of x in the graph $T_\rho^{\otimes R}$. (Here, T_ρ is the usual noise graph on graph on $[q] \times \{\perp, \top\}_\beta$. See .)
6. Sample random neighbors $(\tilde{B}, \tilde{y}) \sim M_w(B, y)$ and $(\tilde{B}', \tilde{y}') \sim M_{w'}(B', y')$. (Here, M_w is the operator that corresponds to resampling the coordinates $i \in [R]$ with $w_i = \perp$. See .)
7. Output an edge between $\pi.(\tilde{B}, \tilde{y}, w)$ and $\pi'.(\tilde{B}', \tilde{y}', w')$.

(End of [Reduction 7.3](#))

[Reduction 7.3](#) has the following approximation guarantees.

Theorem 7.4. *Given a graph G with n vertices and parameters $R \in \mathbb{N}$, $q \in \mathbb{N}$, $\varepsilon, \rho, \beta \in (0, 1)$ with $\rho \geq 0.1$, [Reduction 7.3](#) computes in time $\text{poly}(n^R, q^R)$ a graph $H = H_{R,q,\varepsilon,\rho,\beta}(G)$ such that the following assertions hold (for some constants $\zeta = \zeta(q, \varepsilon, \rho, \beta) > 0$ and $\alpha \geq (1 - \varepsilon)\varepsilon$):*

Completeness I: *If G contains a vertex set with volume $\varepsilon/\beta R$ and expansion at most ε , then H contains q disjoint vertex sets T_1, \dots, T_q , each with volume $\mu(T_i) = \alpha/q$ and expansion at most $\Phi(T_i) \leq 1 - \rho^2 + O(\varepsilon + \beta)$.*

Completeness II: *If G contains $\alpha\beta R/\varepsilon$ disjoint vertex sets, each with volume $\varepsilon/\beta R$ and expansion at most ε , then there exists a distribution over q disjoint vertex sets T_1, \dots, T_q such that*

- *with probability 1, each set T_i has volume and expansion as above,*

7. Reductions between Expansion Problems

- a typical vertex v in H satisfies $|\mathbb{P}_{T_i}\{v \in T_i\} - \alpha/q| \leq O(\varepsilon)\alpha/q + 2^{-\Omega(\varepsilon^3\beta R)}$ for every $i \in [q]$ (which means that the sets T_i uniformly cover all but a small fraction of the graph H).

Soundness: If every vertex set in G with volume $\varepsilon/\beta R$ has expansion less than $1 - \zeta$, then every vertex set T in H satisfies

$$H(T, T) \leq \Gamma_{\rho^2}(\mu(T)) + 3\beta^{1/3}.$$

Remark 7.5. The graph H obtained from [Reduction 7.3](#) is the label-extended graph of unique game. The second completeness assertion shows that this unique game has a uniform cover of $\Omega(1)$ -partial assignments with value close to 1 if G contains a vertex set with volume roughly $\varepsilon/\beta R$ and small expansion. Using correlated sampling (similar to the rounding in [Chapter 4, §4.2](#), such a distribution over partial assignments implies that the value of the unique game is close to 1 (only a factor of two is lost in the closeness to 1). This observation implies that the Small-Set Expansion Hypothesis implies [Hypothesis 6.3](#) (Unique Games Conjecture on Small-Set Expanders), which establishes the equivalence of the two hypotheses (we showed the other direction already in [Chapter 6, §6.4](#)).

7.2.1. Completeness

Let G be a regular graph with vertex set $V = \{1, \dots, n\}$. Let $H = H_{R,q,\varepsilon,\rho,\beta}(G)$ be the graph with vertex set $V(H) = V^R \times [q]^R \times \{\perp, \top\}_\beta^R$ obtained by applying [Reduction 7.3](#) to G . The number of vertices of H is $n^R \cdot O(q)^R$ (and the running time of the reduction is polynomial in the size of H). In particular, [Reduction 7.3](#) is polynomial-time for $R = O(1)$ and $q = O(1)$.

The completeness of [Reduction 7.3](#) asserts that if G contains a set of volume roughly $1/\beta R$ and with small expansion, then a constant fraction of the graph H can be partitioned into q sets, each with the same volume and expansion roughly $1 - \rho^2$.

Lemma 7.6 (Completeness I). *If the graph G contains a vertex set of volume δ and expansion at most η , then the graph $H = H_{R,q,\varepsilon,\rho,\beta}(G)$ obtained by [Reduction 7.3](#) contains disjoint vertex set T_1, \dots, T_q , each of volume α/q and with expansion at most $1 - \rho^2 + O(\varepsilon + \eta + \beta/\rho + \beta R\delta)$, where $\alpha \geq (1 - \beta R\delta)\beta R\delta$.*

7. Reductions between Expansion Problems

Proof. Let $S \subseteq V$ be a vertex subset of G with $\mu(S) = \delta$ and $\Phi(S) = \eta$. Similar to the completeness proof of [Reduction 6.5](#), we construct a partial assignment $F: V^R \times \{\perp, \top\}_\beta^R \rightarrow [R] \cup \{\perp\}$,

$$F(A, z) \stackrel{\text{def}}{=} \begin{cases} r & \text{if } \{r\} = \{r' \in [R] \mid a_{r'} \in S \wedge z_{r'} = \top\}, \\ \perp & \text{otherwise.} \end{cases}$$

(The difference to the construction in the completeness proof for [Reduction 6.5](#) is the additional argument $z \in \{\perp, \top\}_\beta^R$ of F . Effectively, the partial assignment F ignores all coordinates a_r of A with $z_r = \perp$.) We compute the fraction α of vertices labeled by F ,

$$\alpha := \mathbb{P}_{(A, z) \sim V^R \times \{\perp, \top\}_\beta^R} \{F(A, z) \neq \perp\} = \binom{R}{1} (1 - \beta\delta)^{R-1} \beta\delta \geq \beta R\delta(1 - \beta R\delta).$$

Let $S' = S \times \{\top\} \subseteq V \times \{\perp, \top\}_\beta$. Consider the graph $G' = (T_{1-\varepsilon}G) \otimes T_\rho$ on $V \times \{\perp, \top\}_\beta$. The set S' has volume $\beta\delta$ in the graph G' . Furthermore, the fraction of edges of G' staying inside of S' is

$$\begin{aligned} G'(S', S') &= \mu(S') \cdot \mathbb{P}_{a_r \sim V, b_r \sim T_{1-\varepsilon}G} \{b_r \in S \mid a_r \in S\} \cdot \mathbb{P}_{z_r \sim \{\perp, \top\}_\beta, w_r \sim T_\rho(z_r)} \{w_r = \top \mid z_r = \top\} \\ &= \beta\delta \left((1 - \varepsilon)(1 - \eta) + \varepsilon\delta \right) \cdot (\rho + (1 - \rho)\beta). \end{aligned}$$

(Here, we use the notation a_r, b_r, w_r, z_r only for consistency with the description of [Reduction 7.3](#). The index r plays no role at this point.) Hence, the expansion of S' in the graph G' equals

$$\phi' := 1 - \left((1 - \varepsilon)(1 - \eta) + \varepsilon\delta \right) \cdot (\rho + (1 - \rho)\beta) \leq (1 - \rho) + \varepsilon + \eta + \beta.$$

For every $r \in [R]$, we have

$$\begin{aligned} &\mathbb{P}_{A \sim V^R, z \sim \{\perp, \top\}_\beta^R, B \sim (T_{1-\varepsilon}G)^{\otimes R}(A), w \sim T_\rho^{\otimes R}(z)} \{F(B, w) = r \mid F(A, z) = r\} \\ &= (1 - \phi') \left(1 - \phi' \frac{\mu(S')}{1 - \mu(S')} \right)^{R-1} \\ &\geq 1 - \phi' - R\beta\delta. \end{aligned}$$

Furthermore, for every $r \in [R]$,

$$\begin{aligned} \mathbb{P}_{(y, w) \sim T_\rho^{\otimes R}(x, z)} \{x_r = y_r \mid w_r = \top, z_r = \top\} &\geq \rho / \mathbb{P}_{w_r \sim T_\rho(z_r)} \{w_r = \top \mid z_r = \top\} \\ &\geq \rho / (\rho + (1 - \rho)\beta) \geq 1 - \frac{1 - \rho}{\rho} \beta \end{aligned}$$

7. Reductions between Expansion Problems

Let A, B, w, x, y, z be distributed as before (and as in [Reduction 7.3](#)), i.e., $A \sim V^R$, $B \sim (T_{1-\varepsilon}G)^{\otimes R}$, $(x, z) \sim [q]^R \times \{\perp, \top\}_\beta^R$, and $(y, w) \sim T_\rho^{\otimes R}(x, z)$. In addition, sample $\tilde{B} \sim M_w(B)$ and $\tilde{y} \sim M_w(y)$ as in [Reduction 7.3](#). Then, for every $r \in [R]$,

$$\begin{aligned}
\rho' &:= \mathbb{P}\{x_{F(A,z)} = \tilde{y}_{F(\tilde{B},w)} \mid F(A,z) = r\} \\
&\geq \mathbb{P}\{x_r = \tilde{y}_r \wedge F(\tilde{B}, w) = r \mid F(A, z) = r\} \\
&= \mathbb{P}\{x_r = y_r \wedge F(B, w) = r \mid F(A, z) = r\} \\
&= \mathbb{P}\{F(B, w) = r \mid F(A, z) = r\} \cdot \mathbb{P}\{x_r = y_r \mid F(B, w) = r, F(A, z) = r\} \\
&\geq (1 - \phi' - R\beta\delta) \cdot \left(1 - \frac{1-\rho}{\rho}\beta\right) \geq 1 - \phi' - R\beta\delta - \frac{1-\rho}{\rho}\beta.
\end{aligned}$$

The second step uses that (\tilde{B}, \tilde{y}) and (B, y) agree on all coordinates r with $w_r = \top$ (and that $F(\cdot, w)$ ignores coordinates with $w_r \neq \top$).

For our final choice of parameters, ρ' will be arbitrarily close to ρ (also note that by symmetry of F , the expression used to define ρ' does not depend on r .)

We define vertex subsets $T_1, \dots, T_q \subseteq V(H)$ as follows

$$T_i = \{(A, x, z) \in V^R \times [q]^R \times \{\perp, \top\}_\beta^R \mid x_{F(A,z)} = i\}.$$

By symmetry, each set T_i has volume $\mu(T_i) = \alpha/q$ in H . (Recall that α is the fraction of $V^R \times \{\perp, \top\}_\beta^R$ labeled by the partial assignment F .) We claim that the sets T_i have expansion at most $1 - (\rho')^2$ in the graph H . Note that the sets T_1, \dots, T_i are invariant under permuting coordinates (using that $F(\pi.(A, z)) = \pi(F(A, z))$ whenever $F(A, z) \neq \perp$). Hence, in order to compute the expansion of the sets, we can ignore the permutations π and π' in the construction of H (see steps 3 and 7 in [Reduction 7.3](#)). We can compute the fraction of edges of

7. Reductions between Expansion Problems

H staying in T_i as

$$\begin{aligned}
H(T_i, T_i) &= \mathbb{P}_{((\tilde{B}, \tilde{y}, w), (\tilde{B}', \tilde{y}', w')) \sim H} \left\{ y_{F(B, y)} = i \wedge y'_{F(B', w')} = i \right\} \\
&= \mathbb{E}_{(A, x, z)} \left(\mathbb{P}_{(\tilde{B}, \tilde{y}, w) | (A, x, z)} \left\{ \tilde{y}_{F(\tilde{B}, w)} = i \right\} \right)^2 \\
&\geq \mathbb{E}_{(A, x, z)} \left(\mathbb{P}_{(\tilde{B}, \tilde{y}, w) | (A, x, z)} \left\{ x_{F(A, z)} = i \wedge \tilde{y}_{F(\tilde{B}, w)} = x_{F(A, z)} \right\} \right)^2 \\
&= \frac{\alpha}{q} \mathbb{E}_{(A, x, z) | F(A, z) \neq \perp} \left(\mathbb{P}_{(\tilde{B}, \tilde{y}, w) | (A, x, z)} \left\{ \tilde{y}_{F(\tilde{B}, w)} = x_{F(A, z)} \right\} \right)^2 \\
&\geq \frac{\alpha}{q} \left(\mathbb{E}_{(A, x, z) | F(A, z) \neq \perp} \mathbb{P}_{(\tilde{B}, \tilde{y}, w) | (A, x, z)} \left\{ \tilde{y}_{F(\tilde{B}, w)} = x_{F(A, z)} \right\} \right)^2 \\
&= \frac{\alpha}{q} \cdot (\rho')^2 \geq \frac{\alpha}{q} \cdot (\rho^2 - O(\varepsilon + \eta + \beta/\rho + \beta R\delta)). \quad \square
\end{aligned}$$

Remark 7.7 (Completeness II). The second completeness assertion for [Reduction 7.3](#) follows from the proof of the previous lemma (which establishes the first completeness assertion). If the graph G contains $k = \beta R$ disjoint vertex sets S_1, \dots, S_k , each with volume $\delta = \varepsilon/k$ and expansion $\eta \leq \varepsilon$, then we can sample a random vertex set $S \in \{S_1, \dots, S_k\}$ and apply the construction in previous the proof. In this way, we obtain a distribution over disjoint vertex sets T_1, \dots, T_q satisfying the conditions of our second completeness assertion. To verify the condition that the sets T_i uniformly cover all but a small fraction of the graph H , it is enough to show the following the claim: Consider k random vertices $x_1, \dots, x_k \sim G$. Then, in expectation, at least a $(1 - O(\varepsilon))\varepsilon - 2^{-\Omega(\varepsilon^2 k)}$ fraction of the vertex sets S_j contain exactly one of the vertices x_1, \dots, x_k . (Note that in expectation, at most an ε fraction of the vertex sets S_j contain *at least* one of the vertices x_1, \dots, x_k .) To verify the claim, let us condition on the event that at least $k' = (1 - \varepsilon)\varepsilon k$ of the vertices x_1, \dots, x_k fall into one of the vertex sets S_j . We may assume that the first k' vertices $x_1, \dots, x_{k'}$ fall into one of the vertex sets S_j . Each of these vertices x_i (with $i \leq k'$) falls into a random set among S_1, \dots, S_k . Let us imagine we distribute these vertices one-by-one among sets S_j (in the order $x_1, \dots, x_{k'}$). Then, the probability that x_i falls into a set that already contains one of the vertices x_1, \dots, x_{i-1} is at most $k'/k \leq \varepsilon$. Hence, we expect that a $1 - O(\varepsilon)$ fraction of the vertices $x_1, \dots, x_{k'}$ fall into a unique set S_j , which demonstrates the claim. (End of [Remark 7.7](#))

7. Reductions between Expansion Problems

7.2.2. Soundness

Let G be a regular graph with vertex set $V = \{1, \dots, n\}$. Let $H = H_{R, q, \varepsilon, \rho, \beta}(G)$ be the graph with vertex set $V(H) = V^R \times [q]^R \times \{\perp, \top\}_\beta^R$ obtained by applying [Reduction 7.3](#) to G .

In the following, we show that [Reduction 7.3](#) is sound, that is, if H contains a vertex set with volume $\delta = \Omega(1/q)$ and expansion bounded away from $1 - \Gamma_{\rho^2}(\delta)/\delta$ (the minimum expansion of vertex sets of volume δ in the Gaussian noise graph U_{ρ^2}), then the graph H contains a vertex set of volume $\varepsilon/\beta R$ and expansion bounded away from 1.

Let Ω denote the product probability space $[q] \times \{\perp, \top\}_\beta$ of the uniform distribution on $[q]$ and the β -biased distribution on $\{\perp, \top\}$. We can identify the vertex set of H with $V^R \times \Omega^R$.

Let $T \subseteq V(H)$ be a vertex set in H . Let $f: V^R \times \Omega^R$ be its $\{0, 1\}$ -indicator function. We consider the following two symmetrizations of f ,

$$\bar{f}(A, x, z) := \mathbb{E}_{\pi \in \mathbb{S}_R} f(\pi.(A, x, z)), \quad (7.2)$$

$$\tilde{f}(A, x, z) := \mathbb{E}_{(\tilde{A}, \tilde{x}) \sim M_z(A, x)} \tilde{f}(\tilde{A}, \tilde{x}, z). \quad (7.3)$$

By the symmetries of the graph H , we have

$$\langle f, Hf \rangle = \langle \bar{f}, H\bar{f} \rangle = \langle \tilde{f}, H\tilde{f} \rangle.$$

We write $\tilde{f}_A(x, z) = \tilde{f}(A, x, z)$. For all $A \in V^R$, we define functions $g_A: \Omega^R \rightarrow [0, 1]$ as the average of \tilde{f}_B over the neighbors B of the vertex A in the product graph $G^{\otimes R}$ (with additional noise),

$$g_A := \mathbb{E}_{B \sim (T_{1-\varepsilon}G)^{\otimes R}(A)} \tilde{f}_B.$$

The fraction of edges staying inside of T equals the typical noise stability of g_A in the graph T_ρ on Ω^R ,

$$\begin{aligned} H(T, T) &= \mathbb{E}_{A \sim V^R} \mathbb{E}_{(x, z) \sim \Omega^R} \left(\mathbb{E}_{B \sim (T_{1-\varepsilon}G)^{\otimes R}(A)} \mathbb{E}_{(y, w) \sim T_\rho(x, z)} \tilde{f}_B(y, w) \right)^2 \\ &= \mathbb{E}_{A \sim V^R} \mathbb{E}_{(x, z) \sim \Omega^R} \left(T_\rho g_A(x, z) \right)^2 = \mathbb{E}_{A \sim V^R} \|T_\rho g_A\|^2. \end{aligned}$$

We can further upper bound $H(T, T)$ by combining the soundness analysis of [Reduction 6.20](#) (in particular, the Invariance Principle and Influence Decoding) with the soundness analysis of [Reduction 6.5](#). Combining these tools (which is non-trivial, but standard), establishes the following claim.

7. Reductions between Expansion Problems

Claim 7.8. *There exists a constant $\zeta = \zeta(q, \varepsilon, \rho, \beta) > 0$ such that either*

$$H(T, T) \leq \mathbb{E}_{A \sim V^R} \Gamma_{\rho^2}(\mathbb{E} g_A) + \beta.$$

or G contains a vertex set with volume $\varepsilon/\beta R$ and expansion less than $1 - \zeta$.

Next, we show that for most vertex tuples $A \in V^R$, the expectation of the function g_A (over Ω^R) roughly equals the expectation of the function f . (Since f indicates the vertex set T , its expectation is the volume of T .) To this end, we compare $\mathbb{E}_A(\mathbb{E} g_A)^2$ to $(\mathbb{E} f)^2$ in the following claim. (We remark that one could derive upper bounds on $\mathbb{E}_A(\mathbb{E} g_A)^2$ using properties of the graph $G^{\otimes R}$. In fact, we followed this approach in the soundness proof of [Reduction 6.20](#). The key novelty of [Reduction 7.3](#) is that we can control $\mathbb{E}_A(\mathbb{E} g_A)^2$ without further assumptions on the graph $G^{\otimes R}$.)

Claim 7.9.

$$\mathbb{E}_A(\mathbb{E} g_A)^2 \leq (\mathbb{E} f)^2 + \beta \|f\|^2.$$

Proof. Recall that g_A is defined by averaging the function \bar{f} ,

$$g_A(x, z) = \mathbb{E}_{B \sim (T_{1-\varepsilon} G)^{\otimes R}(A)} \mathbb{E}_{(\tilde{B}, \tilde{x}) \sim M_z(B, x)} \bar{f}(\tilde{B}, \tilde{x}, z).$$

Using Cauchy–Schwarz, the typical value of the square of $\mathbb{E} g_A$ is at most

$$\begin{aligned} \mathbb{E}_A(\mathbb{E} g_A)^2 &= \mathbb{E}_{A \sim V^R} \left(\mathbb{E}_{(x, z) \sim \Omega^R} \mathbb{E}_{B \sim (T_{1-\varepsilon} G)^{\otimes R}(A)} \mathbb{E}_{(\tilde{A}, \tilde{x}) \sim M_z(A, x)} \bar{f}(\tilde{A}, \tilde{x}, z) \right)^2 \\ &\leq \mathbb{E}_{A \sim V^R, x \sim [q]^R} \underbrace{\left(\mathbb{E}_{z \sim \{\perp, \top\}_\beta^R} \mathbb{E}_{(\tilde{A}, \tilde{x}) \sim M_z(A, x)} \bar{f}(\tilde{A}, \tilde{x}, z) \right)^2}_{=: M\bar{f}(A, x)} = \|M\bar{f}\|^2. \end{aligned}$$

Notice that M (as defined above) is a linear operator from the space $L_2(V^R \times [q]^R \times \{\perp, \top\}_\beta^R)$ to the space $L_2(V^R \times [q]^R)$.

To upper bound the norm $\|M\bar{f}\|$, we consider the spectral decomposition of the linear operator M^*M on $L_2(V^R \times [q]^R \times \{\perp, \top\}_\beta^R)$. Here, M^* is the adjoint of M , so that $\langle Mh, g \rangle = \langle h, M^*g \rangle$ for any two functions $h \in L_2(V^R \times [q]^R \times \{\perp, \top\}_\beta^R)$ and $g \in L_2(V^R \times [q]^R)$. One can verify that M^*M is a Markov operator. Its largest eigenvalue is 1 and the constant functions form the corresponding eigenspace.

7. Reductions between Expansion Problems

To show the claim, it is enough to verify that the second largest eigenvalue of M^*M equals β , because

$$\|M\bar{f}\|^2 = \langle \bar{f}, M^*M\bar{f} \rangle \leq (\mathbb{E}\bar{f})^2 + \beta\|\bar{f} - \mathbb{E}\bar{f}\|^2 \leq (\mathbb{E}f)^2 + \beta\|f\|^2.$$

(The first inequality uses that $\bar{f} - \mathbb{E}\bar{f}$ is orthogonal to the first eigenspace and thus contributes at most $\beta\|\bar{f} - \mathbb{E}\bar{f}\|^2$ to the inner product $\langle \bar{f}, M^*M\bar{f} \rangle$. The second inequality uses that \bar{f} is obtained from f by averaging, which implies that $\mathbb{E}\bar{f} = \mathbb{E}f$ and $\|\bar{f}\|^2 \leq \|f\|^2$.)

To compute the second largest eigenvalue of M^*M , we note that this operator has the same non-zero eigenvalues as the operator MM^* (which acts on $L_2(V^R \times [q]^R)$). One can verify that MM^* corresponds to the noise graph T_β on $V^R \times [q]^R$. (See also (7.1).) Hence, the second largest eigenvalue of M^*M equals β . \square

We conclude the soundness analysis of [Reduction 7.3](#) by combining the previous claims. Let $\gamma > 0$ be sufficiently small. [Claim 7.9](#) allows us to upper bound the fraction of vertex tuples $A \in V^R$ with $\mathbb{E}g_A > \mathbb{E}f + \gamma$. Concretely, Chebyshev's inequality implies

$$\mathbb{P}_{A \sim V^R} \left\{ \mathbb{E}g_A > \mathbb{E}f + \gamma \right\} \leq \mathbb{E}_{A \sim V^R} \left(\mathbb{E}g_A - \mathbb{E}f \right)^2 / \gamma^2 \leq \beta / \gamma^2.$$

Let us assume that G contains no vertex set with volume $\varepsilon/\beta R$ and expansion less than $1 - \zeta$. Then, using [Claim 7.8](#), the following upper bound on $H(T, T)$ holds

$$H(T, T) \leq \mathbb{E}_{A \sim V^R} \Gamma_{\rho^2}(\mathbb{E}g_A) + \beta \leq \Gamma_{\rho^2}(\mathbb{E}f + \gamma) + \beta/\gamma^2 + \beta \leq \Gamma_{\rho^2}(\mathbb{E}f) + 2\gamma + 2\beta/\gamma^2.$$

(In the last step, we use that the function Γ_{ρ^2} is 2-Lipschitz.) To achieve the best upper bound on $H(T, T)$, we choose $\gamma = \beta^{1/3}$, which shows that $H(T, T) \leq \Gamma_{\rho^2}(\mathbb{E}f) + 3\beta^{1/3}$.

To summarize, we established the following soundness guarantee of [Reduction 7.3](#).

Lemma 7.10 (Soundness). *There exists a constant $\zeta = \zeta(q, \varepsilon, \rho, \beta) > 0$ such that if all vertex sets with volume $\varepsilon/\beta R$ in G have expansion at least $1 - \zeta$, then every vertex set T in $H = H_{R, q, \varepsilon, \rho, \beta}(G)$ satisfies*

$$H(T, T) \leq \Gamma_{\rho^2}(\mu(T)) + 3\beta^{1/3}$$

7.3. Notes

The material of this chapter is based on an unpublished manuscript with Prasad Raghavendra and Madhur Tulsiani, titled “Reductions between Expansion Problems” [RST10b].

Part III.
Lower Bounds

8. Limits of Semidefinite Relaxations

Non-trivial consequences of a refutation of the Unique Games Conjecture (as we showed in the [Chapter 6](#)) are some form of evidence for the truth of the conjecture. Another form of evidence for the truth of the conjecture are lower bounds for UNIQUE GAMES in restricted models of computation.

In the context of approximation problems, it is natural to study models of computations defined by hierarchies of relaxations (typically linear or semidefinite relaxations). Such hierarchies contain relaxations with gradually increasing complexity — from linear complexity to exponential complexity. Relaxations with higher complexity provide better approximations for the optimal value of optimization problems. (The relaxations with highest, i.e., exponential, complexity typically compute the optimal value exactly.)

Starting with the seminal work of Arora, Bollobás, and Lovász [[ABL02](#)], lower bounds on the complexity of many approximation problems were obtained in various hierarchies. In this chapter, we show the first super-polynomial lower bound for UNIQUE GAMES in a hierarchy that captures the best known algorithms for all constraint satisfaction problems (see [Chapter 8](#)). Previously known lower bounds for UNIQUE GAMES, considered only relaxations with fixed polynomial complexity [[KV05](#)], or considered hierarchies that do not capture current algorithms (in particular, the hierarchies provide only trivial approximation for MAX CUT) [[CMM09](#)].

Our lower bounds translate to corresponding lower bounds for all UG-hard problems (via the known UG-hardness reductions). Even for specific UG-hard problem like MAX CUT such lower bounds were not known before.

8.1. Overview

In this chapter, we exhibit an integrality gap for certain strong SDP relaxations of UNIQUE GAMES. More precisely, we consider two strong hierarchies of SDP relaxations $\{LH_r\}_{r \in \mathbb{N}}$ and $\{SA_r\}_{r \in \mathbb{N}}$ (see [Section 8.3](#) for the definitions). We give

8. Limits of Semidefinite Relaxations

rough definitions of these relaxations here for the convenience of the reader. The r^{th} level relaxation LH_r consists of the following: 1) SDP vectors for every vertex of the unique game, 2) All valid constraints on vectors corresponding to at most r vertices. Equivalently, the LH_r relaxation consists of SDP vectors and local distributions μ_S over integral assignments to sets S of at most r variables, such that the second moments of local distributions μ_S match the corresponding inner products of SDP vectors.

The SA_r relaxation is a strengthening of LH_r with the additional constraint that for two sets S, T of size at most r , the corresponding local distribution over integral assignments μ_S, μ_T must have the same marginal distribution over $S \cap T$. The SA_r relaxation corresponds to simple SDP relaxation strengthened by r^{th} round of Sherali-Adams hierarchy [SA90]. Let $\text{LH}_r(\Phi)$ and $\text{SA}_r(\Phi)$ denote the optimum value of the corresponding SDP relaxations on the instance Φ . Further, let $\text{opt}(\Phi)$ denote the value of the optimum labeling for Φ . For the LH and SA hierarchies, we show:

Theorem 8.1. *For all constants $\eta > 0$, there exists a UNIQUE GAMES instance Φ on N vertices such that $\text{LH}_r(\Phi) \geq 1 - \eta$ and $\text{opt}(\Phi) \leq \eta$ for $r = O(2^{(\log \log N)^{\frac{1}{4}}})$*

Theorem 8.2. *For all constants $\eta > 0$, there exists a UNIQUE GAMES instance Φ on N vertices such that $\text{SA}_r(\Phi) \geq 1 - \eta$ and $\text{opt}(\Phi) \leq \eta$ for $r = O((\log \log N)^{\frac{1}{4}})$*

Demonstrated for the first time in [KV05], and used in numerous later works [CMM09, STT07b, Tul09, Rag08, GMR08, MNRS08], it is by now well known that integrality gaps can be composed with hardness reductions.

In particular, given a reduction Red from UNIQUE GAMES to a certain problem Λ , on starting the reduction with a integrality gap instance Φ for UNIQUE GAMES , the resulting instance $\text{Red}(\Phi)$ is a corresponding integrality gap for Λ . Composing the integrality gap instance for LH_r or SA_r relaxation of UNIQUE GAMES , along with UG reductions in [KKMO07, Aus07, Rag08, GMR08, MNRS08, RS09b], one can obtain integrality gaps for LH_r and SA_r relaxations for several important problems. For the sake of succinctness, we will state the following general theorem:

Theorem 8.3. *Let Λ denote a problem in one of the following classes:*

- *A Generalized Constraint Satisfaction Problem*
- *An Ordering Constraint Satisfaction Problem*

8. Limits of Semidefinite Relaxations

Let SDP denote the SDP relaxation that yields the optimal approximation ratio for Λ under UGC. Then the following holds: Given an instance \mathfrak{I} of the problem Λ , with $\text{SDP}(\mathfrak{I}) \geq c$ and $\text{opt}(\mathfrak{I}) \leq s$, for every constant $\eta > 0$, there exists an instance Ψ_η over N variables such that:

- $\text{LH}_r(\Psi_\eta) \geq c - \eta$ and $\text{opt}(\Psi_\eta) \leq s + \eta$ with $r = O(2^{(\log \log N)^{1/4}})$.
- $\text{SA}_r(\Psi_\eta) \geq c - \eta$ and $\text{opt}(\Psi_\eta) \leq s + \eta$ with $r = O((\log \log N)^{1/4})$.

The O notation in the number of rounds hides a constant depending on η .

The classes of problems for which the above result holds include MAX CUT [KKMO07], MAX 2-SAT [Aus07], $\text{GROTHENDIECK PROBLEM}$ [RS09b] k - WAY CUT [MNRS08] and $\text{MAXIMUM ACYCLIC SUBGRAPH}$ [GMR08]. Notable exceptions that do not directly fall under this framework are VERTEX COVER and SPARSEST CUT .

Reductions from UNIQUE GAMES to SPARSEST CUT have been exhibited in [KV05] and [CKK⁺06]. With the integrality gap for LH_r relaxation of UNIQUE GAMES (Theorem 8.1), these reductions imply a corresponding LH_r integrality gap for SPARSEST CUT . Integrality gaps for SPARSEST CUT are directly related to lower bounds for distortion required to embed a given metric into L_1 metric. Here the L_1 metric consists of points in \mathbb{R}^d for arbitrarily large d , and the distance between two points (x, y) is $\|x - y\|_1$. An L_2^2 metric consists of a set of points in \mathbb{R}^d such that the squares of the distances between them also form a metric (satisfy triangle inequality). Restated in this language, the SDP vectors of the SPARSEST CUT integrality gap that we construct, yield the following result:

Theorem 8.4. *For some absolute constants $\gamma, \delta > 0$, there exists an N -point L_2^2 metric that requires distortion at least $\Omega(\log \log N)^\delta$ to embed into L_1 , while every set of size at most $O(2^{(\log \log N)^\gamma})$ embeds isometrically into L_1 .*

The $\text{UNIFORM SPARSEST CUT}$ problem is among the many important problems for which no UNIQUE GAMES reduction is known. In [DKSV06], the techniques of [KV05] were extended to obtain an integrality gap for $\text{UNIFORM SPARSEST CUT}$ for the SDP with triangle inequalities. Roughly speaking, the SDP gap construction in [DKSV06] consists of the hypercube with its vertices identified by certain symmetries such as cyclic shift of the coordinates. Using the techniques from this chapter, the following SDP integrality gap for the $\text{BALANCED SEPARATOR}$ problem can be exhibited. The details of the proof of this theorem are omitted from the thesis.

8. Limits of Semidefinite Relaxations

Theorem 8.5. *For some absolute constants $\gamma, \delta > 0$, there exists an instance G on N vertices of BALANCED SEPARATOR such that the ratio $\text{opt}(G)/\text{LH}_r(G) \geq \Omega(\log \log N)^\delta$ for $r = O(\log \log N)^\gamma$.*

8.1.1. Techniques

In this section, we will present a brief overview of the techniques and a roadmap for the rest of the chapter.

The overall strategy in this work to construct SDP integrality gaps is along the lines of Khot–Vishnoi [KV05]. Let us suppose we wish to construct a SDP integrality gap for a problem Λ (say MAX CUT). Let Red_Λ be a reduction from UNIQUE GAMES to the problem Λ . The idea is to construct a SDP integrality gap Φ for UNIQUE GAMES, and then execute the reduction Red_Λ on the instance Φ , to obtain the SDP gap construction $\text{Red}_\Lambda(\Phi)$. Surprisingly, as demonstrated in [KV05], the SDP vector solution for Φ can be transformed through the reduction to obtain the SDP solution for $\text{Red}_\Lambda(\Phi)$.

Although this technique has been used extensively in numerous works [CMM09, STT07a, Tul09, Rag08, GMR08, MNRS08] since [KV05], there is a crucial distinction between [KV05] and later works. In all other works, starting with an SDP gap Φ for UNIQUE GAMES, one obtains an integrality gap for an SDP relaxation that is no stronger. For instance, starting with a integrality gap for 10-rounds of a SDP hierarchy, the resulting SDP gap instance satisfies at most 10 rounds of the same hierarchy.

The surprising aspect of [KV05], is that it harnesses the UG reduction Red_Λ to obtain an integrality gap for a “stronger” SDP relaxation than the one which it started with. Specifically, starting with an integrality gap Φ for a simple SDP relaxation of UNIQUE GAMES, [KV05] exhibit an SDP gap for MAX CUT which obeys all valid constraints on 3 variables. The proof of this fact (the triangle inequality) is perhaps the most technical and least understood aspect about [KV05]. One of the main contributions of this chapter is to conceptualize and simplify this aspect of [KV05]. Armed with the understanding of [KV05], we then develop the requisite machinery to extend it to a strong SDP integrality gap for UNIQUE GAMES.

To obtain strong SDP gaps for UNIQUE GAMES, we will apply the above strategy on the reduction from UNIQUE GAMES to E2Lin_q obtained in [KKMO07]. Note that E2Lin_q is a special case of UNIQUE GAMES. Formally, we show the following reduction from a *weak gap* instance for UNIQUE GAMES over a large alphabet to a integrality gap for a strong SDP relaxation of E2Lin_q .

8. Limits of Semidefinite Relaxations

Theorem 8.6. (*Weak Gaps for UNIQUE GAMES \implies Strong gaps for E2Lin $_q$*)
 For a positive integer q , let $\text{Red}_{\text{E2Lin}_q}$ denote the reduction from UNIQUE GAMES to E2Lin $_q$. Given a $(1 - \eta, \delta)$ -weak gap instance Φ for UNIQUE GAMES, the E2Lin $_q$ instance $\text{Red}_{\text{E2Lin}_q}(\Phi)$ is a $(1 - 2\gamma, 1/q^{\gamma/2} + o_\delta(1))$ SDP gap for the relaxation LH_r for $r = 2^{O(1/\eta^{1/4})}$. Further, $\text{Red}_{\text{E2Lin}_q}(\Phi)$ is a $(1 - \gamma, \delta)$ SDP gap for the relaxation SA_r for $r = O(1/\eta^{1/4})$.

Using the weak gap for UNIQUE GAMES constructed in [KV05], along with the above theorem, implies Theorems 8.1 and 8.2. As already pointed out, by now it is fairly straightforward to compose an r -round integrality gap for UNIQUE GAMES, with reductions to obtain a r round integrality gaps for other problems. Hence, Theorem 8.3 is a fairly straightforward consequence of Theorems 8.1 and 8.2.

8.1.2. Organization

In the next section, we present a detailed proof overview that describes the entire integrality gap construction restricted to the case of MAX CUT. The formal definitions of the SDP hierarchies LH_r, SA_r and their robustness are presented in Section 8.3. We formally define weak gap instances for UNIQUE GAMES in Section 8.4. We also outline an alternate integrality gap for a very minimal SDP relaxation of UNIQUE GAMES in the same section. This section is followed by the description of the integrality gap instance for E2Lin $_q$ obtained by reduction of Khot et al. [KKMO07]. In the rest of the chapter, we construct SDP vectors and local distributions to show that this is an integrality gap for the strong SDP relaxations – LH_r and SA_r . The two subsequent sections are devoted to developing the requisite machinery of integral vectors, their tensor products and local distributions for UNIQUE GAMES. The SDP vectors and local distributions for the integrality gap instance described in Section 8.5 are exhibited in Sections Section 8.8 and Section 8.8.2.

8.2. Proof Sketch

For the sake of exposition, we will describe the construction of an SDP integrality gap for MAX CUT. To further simplify matters, we will exhibit an integrality gap for the basic Goemans-Williamson relaxation, augmented with the triangle inequalities on every three vectors. While an integrality gap of

8. Limits of Semidefinite Relaxations

this nature is already part of the work of Khot–Vishnoi [KV05], our proof will be conceptual and amenable to generalization.

Let Φ be a SDP integrality gap for UNIQUE GAMES on an alphabet $[R]$. For each vertex B in Φ , the SDP solution associates R orthogonal unit vectors $B = \{b_1, \dots, b_R\}$. For the sake of clarity, we will refer to a vertex B in Φ and the set of vectors $B = \{b_1, \dots, b_R\}$ associated with it as a “cloud”. The clouds satisfy the following properties:

- (Matching Property) For every two clouds A, B , there is a unique matching $\pi_{B \leftarrow A}$ along which the inner product of vectors between A and B is maximized. Specifically, if $\rho(A, B) = \max_{a \in A, b \in B} \langle a, b \rangle$, then for each vector a in A , we have $\langle a, \pi_{B \leftarrow A}(a) \rangle = \rho(A, B)$.
- (High objective value) For most edges $e = (A, B)$ in the UNIQUE GAMES instance Φ , the maximal matching $\pi_{A \leftarrow B}$ is the same as the permutation π_e corresponding to the edge, and $\rho(A, B) \approx 1$.

Let $\text{Red}_{\text{MAX CUT}}(\Phi)$ be the MAX CUT instance obtained by executing the reduction in [KKMO07] on Φ . The reduction $\text{Red}_{\text{MAX CUT}}$ in [KKMO07] introduces a long code (2^R vertices indexed by $\{-1, 1\}^R$) for every cloud in Φ . Hence the vertices of $\text{Red}_{\text{MAX CUT}}(\Phi)$ are given by pairs (B, x) where B is a cloud in Φ and $x \in \{-1, 1\}^R$.

The SDP vectors we construct for the integrality gap instance resemble (somewhat simpler in this work) the vectors in [KV05]. Roughly speaking, for a vertex (B, x) , we associate an SDP vector $\mathbf{V}^{B,x}$ defined as follows:

$$\mathbf{V}^{B,x} = \frac{1}{\sqrt{R}} \sum_{i \in [R]} x_i b_i^{\otimes t}$$

The point of departure from [KV05] is the proof that the vectors form a feasible solution for the stronger SDP. Instead of directly showing that the inequalities hold for the vectors, we exhibit a distribution over integral assignments whose second moments match the inner products. Specifically, to show that triangle inequality holds for three vertices $S = \{(A, x), (B, y), (C, z)\}$, we will exhibit a μ_S distribution over $\{\pm 1\}$ assignments to the three vertices, such that

$$\mathbb{E}_{\{Y^{A,x}, Y^{B,y}, Y^{C,z}\} \sim \mu_S} [Y^{A,x} Y^{B,y}] = \langle \mathbf{V}^{A,x}, \mathbf{V}^{B,y} \rangle$$

The existence of an integral distribution matching the inner products shows that the vectors satisfy all valid inequalities on the three variables, including

8. Limits of Semidefinite Relaxations

the triangle inequality. We shall construct the distribution $\mu_{\mathcal{S}}$ over local assignments in three steps,

Local Distributions over Labelings for Unique Games. For a subset of clouds \mathcal{S} within the UNIQUE GAMES instance Φ , we will construct a distribution $\mu_{\mathcal{S}}$ over labelings to the set \mathcal{S} . The distribution $\mu_{\mathcal{S}}$ over $[R]^{\mathcal{S}}$ will be “consistent” with the SDP solution to Φ . More precisely, if two clouds A and B are *highly correlated* ($\rho(A, B) \approx 1$), then when the distribution $\mu_{\mathcal{S}}$ assigns label ℓ to A , with high probability it assigns the corresponding label $\pi_{B \leftarrow A}(\ell)$ to B . Recall that $\rho(A, B)$ was defined as $\max_{a \in A, b \in B} \langle a, b \rangle$.

Consider a set \mathcal{S} where every pair of clouds A, B are *highly correlated* ($\rho(A, B) \geq 0.9$). We will refer to such a set of clouds as *Consistent*. For a *Consistent* set \mathcal{S} , assigning a label ℓ for a cloud A in \mathcal{S} , forces the label of every other cloud B to $\pi_{B \leftarrow A}(\ell)$. Furthermore, it is easy to check that the resulting labeling satisfies consistency for every pair of clouds in \mathcal{S} . (see Lemma 8.39 for details) Hence, in this case, the distribution $\mu_{\mathcal{S}}$ could be simply obtained by picking the label ℓ for an arbitrary cloud in \mathcal{S} uniformly at random, and assigning every other cloud the induced label.

Now consider a set \mathcal{S} which is not consistent. Here the idea is to decompose the set of clouds \mathcal{S} into clusters, such that each cluster is consistent. Given a decomposition, for each cluster the labeling can be independently generated as described earlier. In this chapter, we will use a geometric decomposition to decompose the set of clouds \mathcal{S} into clusters. The crucial observation is that the correlations $\rho(A, B)$ for clouds $A, B \in \mathcal{S}$, can be approximated well by a certain L_2^2 metric. More precisely, for each cloud A , we can associate a unit vector $\mathbf{v}_A = \sum_{a \in A} a^{\otimes s}$ such that the L_2^2 distance between $\mathbf{v}_A, \mathbf{v}_B$ is a good approximation of the quantity $1 - \rho(A, B)$.

By using t random halfspace cuts on this geometric representation, we obtain a partition into 2^t clusters. A pair of clouds A, B that are not highly correlated ($\rho(A, B) < 1 - 1/16$), are separated by the halfspaces with probability at least $1 - (3/4)^t$. Hence for a large enough t , all resulting clusters are consistent with high probability. (see Lemma 8.42).

A useful feature of the geometric clustering is that for two subsets $\mathcal{T} \subset \mathcal{S}$, the distribution over labelings $\mu_{\mathcal{T}}$ is equal to the marginal of the distribution $\mu_{\mathcal{S}}$ on \mathcal{T} . To see this, observe that the distribution over clusterings depends solely on the geometry of the associated vectors. On the downside, the geometric clustering produces inconsistent clusters with a very small but non-zero probability. (see Corollary 8.44).

8. Limits of Semidefinite Relaxations

The details of the construction of local distributions to UNIQUE GAMES are presented in 8.7.

Constructing Approximate Distributions. Fix a set $S \subseteq \mathcal{S} \times \{\pm 1\}^R$ of vertices in the MAX CUT instance $\text{Red}_{\text{MAX CUT}}(\Phi)$. We will now describe the construction of the local integral distribution μ_S .

In the reduction $\text{Red}_{\text{MAX CUT}}$, the labeling ℓ to a cloud B in the UNIQUE GAMES instance is encoded as choosing the ℓ th dictator cut in the long code corresponding to cloud B . Specifically, assigning the label ℓ to a cloud B should translate into assigning x_ℓ for every vertex (B, x) in the long code of B . Hence, a straightforward approach to define the distribution μ_S would be the following:

- Sample a labeling $\ell : \mathcal{S} \rightarrow [R]$ from the distribution μ_S ,
- For every vertex $(B, x) \in \mathcal{S}$, assign $x_{\ell(B)}$.

Although inspired by this, our actual construction of μ_S is slightly more involved. First, we make the following additional assumption regarding the UNIQUE GAMES instance Φ :

Assumption: All the SDP vectors for the integrality gap instance Φ are $\{\pm 1\}$ -vectors (have all their coordinates from $\{\pm 1\}$).

The SDP gap instance for UNIQUE GAMES constructed in [KV05] satisfies this additional requirement. Furthermore, we outline a generic transformation to convert an arbitrary UNIQUE GAMES SDP gap into one that satisfies the above property (see Observation 8.25). A $\{\pm 1\}$ -vector is to be thought of as a distribution over $\{\pm 1\}$ assignments. It is easy to see that tensored powers of $\{\pm 1\}$ -vectors yield $\{\pm 1\}$ -vectors. Let T denote the number of coordinates in the vectors $\mathbf{V}^{B,x}$. The distribution μ_S is defined as follows,

- Sample a labeling $\ell : \mathcal{S} \rightarrow [R]$ from the distribution μ_S , and a coordinate $i \in [T]$ uniformly at random.
- For every vertex $(B, x) \in \mathcal{S}$, assign $Y^{B,x}$ to be the i th coordinate of the vector $x_{\ell(B)} b_{\ell(B)}^{\otimes t}$.

We will now argue that the first two moments of the local distributions μ_S defined above, approximately match the corresponding inner products between SDP vectors.

8. Limits of Semidefinite Relaxations

Consider the inner product $\langle \mathbf{V}^{A,x}, \mathbf{V}^{B,y} \rangle$ of the SDP vectors corresponding to some pair of vertices (A, x) and (B, y) in S . The inner product consists of R^2 terms of the form $\langle x_i a_i^{\otimes t}, y_j b_j^{\otimes t} \rangle$. The crucial observation we will utilize is that the inner product $\langle \mathbf{V}^{A,x}, \mathbf{V}^{B,y} \rangle$ is approximately determined by the R terms corresponding to the matching $\pi_{B \leftarrow A}$. In other words, we have

$$\langle \mathbf{V}^{A,x}, \mathbf{V}^{B,y} \rangle \approx \frac{1}{n} \sum_{\ell \in [R]} x_\ell y_{\pi_{B \leftarrow A}(\ell)} \langle a_\ell^{\otimes t}, b_{\pi_{B \leftarrow A}(\ell)}^{\otimes t} \rangle \leq \rho(A, B)^t$$

(see Section 8.4.1 for details)

If $\rho(A, B) < 0.9$, then with high probability the clustering would place the clouds A, B in different clusters. Hence the labels assigned to A, B would be completely independent of each other, and so would the assignments to (A, x) and (B, y) . Hence, we would have $\mathbb{E}[Y^{A,x} Y^{B,y}] = 0$. On the other hand, by the above inequality the inner product $\langle \mathbf{V}^{A,x}, \mathbf{V}^{B,y} \rangle \leq 0.9^t \approx 0$. Therefore, for clouds A, B that are not *highly correlated*, the inner product of vectors $\mathbf{V}^{A,x}, \mathbf{V}^{B,y}$ agree approximately with the distribution over local assignments.

At the other extreme, if $\rho(A, B) \approx 1$, then with high probability the clustering would not separate A from B . If A, B are not separated, then the distribution μ_S over labelings will respect the matching between A and B . Specifically, whenever A is assigned label ℓ by μ_S , with high probability B is assigned the label $\pi_{B \leftarrow A}(\ell)$. Consequently, in this case we have

$$\mathbb{E}_{\mu_S}[Y^{A,x} Y^{B,y}] = \frac{1}{n} \sum_{\ell \in [R]} \langle x_\ell a_\ell^{\otimes t}, y_{\pi_{B \leftarrow A}(\ell)} b_{\pi_{B \leftarrow A}(\ell)}^{\otimes t} \rangle \approx \langle \mathbf{V}^{A,x}, \mathbf{V}^{B,y} \rangle$$

Smoothing. In Section 8.9, we show a robustness property for the LH_r and SA_r relaxations by which approximately feasible solutions to these hierarchies can be converted (*smoothed*) into perfectly feasible solutions with a small loss in the objective value.

To illustrate the idea behind the robustness, consider a set of unit vectors $\{v_i\}_{i=1}^R$ that satisfy all triangle inequalities up to an additive error of ε , i.e.,

$$\|v_i - v_j\|^2 + \|v_j - v_k\|^2 - \|v_i - v_k\|^2 \geq -\varepsilon$$

We include formal statements of the claims about the robustness of solutions to LH_r and SA_r (Theorems 8.11, 8.10) in Section 8.3. We refer the reader to Section 8.9 for the proofs of these claims.

8. Limits of Semidefinite Relaxations

Extending to $E2Lin_q$. The above argument for MAX CUT can be made precise. However, to obtain an SDP gap for larger number of rounds, we use a slightly more involved construction of SDP vectors.

$\{\pm 1\}$ -vectors were natural in the above discussion, since MAX CUT is a CSP over $\{0, 1\}$. For $E2Lin_q$, it is necessary to work with vectors whose coordinates are from \mathbb{F}_q , as opposed to $\{\pm 1\}$. The tensoring operation for \mathbb{F}_q -integral vectors is to be appropriately defined to ensure that while the behaviour of the inner products resemble traditional tensoring, the tensored vectors are \mathbb{F}_q -integral themselves (see Section 8.6 for details).

For the case of MAX CUT, we used a gap instance Φ for UNIQUE GAMES all of whose SDP vectors were $\{\pm 1\}$ -vectors. In case of $E2Lin_q$, the SDP vectors corresponding to the UNIQUE GAMES instance Φ would have to be \mathbb{F}_q -integral vectors. We outline a generic transformation to convert an arbitrary UNIQUE GAMES SDP gap into one that satisfies this property (see Observation 8.35).

8.3. Hierarchies of SDP Relaxations

In this section, we define the LH_r and SA_r hierarchies and formally state their robustness properties.

8.3.1. LH_r -Relaxation

Let \mathfrak{F} be a CSP (say UNIQUE GAMES) instance over a set of variables \mathcal{V} , alphabet size q and arity k . A feasible solution to the LH_r relaxation consists of the following:

1. A collection of (local) distributions $\{\mu_S\}_{S \subseteq \mathcal{V}, |S| \leq r}$, where $\mu_S: [q]^S \rightarrow \mathbb{R}_+$ is a distribution over $[q]$ -assignments to S , that is, $\mu_S \in \Delta([q]^S)$.
2. A (global) vector solution $\{\mathbf{b}_{i,a}\}_{i \in \mathcal{V}, a \in [q]}$, where $\mathbf{b}_{i,a} \in \mathbb{R}^d$ for every $i \in \mathcal{V}$ and $a \in [q]$.

8. Limits of Semidefinite Relaxations

LH_r-Relaxation.

$$\begin{aligned} \text{maximize} \quad & \mathbb{E}_{P \sim \mathcal{P}} \mathbb{E}_{x \sim \mu_P} P(x) && (\text{LH}_r) \\ \text{subject to} \quad & \langle \mathbf{b}_{i,a}, \mathbf{b}_{j,b} \rangle = \mathbb{P}_{x \sim \mu_S} \{x_i = a, x_j = b\} \quad S \subseteq \mathcal{V}, |S| \leq r, i, j \in S, a, b \in [q], && (8.1) \\ & \mu_S \in \Delta([q]^S) && (8.2) \end{aligned}$$

Here, $\Delta([q]^S)$ denotes probability distributions over $[q]^S$. As usual, we denote by $\text{LH}_r(\mathfrak{F})$ the value of an optimal solution to this relaxation.

The above relaxation succinctly encodes all possible inequalities on up to r vectors. The next remark makes this observation precise.

Remark 8.7. A linear inequality on the inner products of a subset of vectors $\{\mathbf{b}_{i,a}\}_{i \in S, a \in [q]}$ for $S \subseteq \mathcal{V}$ is *valid* if it holds for all distributions over $[q]$ -assignments to the variables S . A feasible solution to the LH_r-relaxation satisfies all valid inequalities on sets of up to r vectors.

8.3.2. SA_r-Relaxation

Enforcing consistency between the marginals of the local distributions yields the SA_r-relaxation.

SA_r-Relaxation:

$$\begin{aligned} \text{maximize} \quad & \mathbb{E}_{P \sim \mathcal{P}} \mathbb{E}_{x \sim \mu_P} P(x) && (\text{SA}_r) \\ \text{subject to} \quad & \langle \mathbf{b}_{i,a}, \mathbf{b}_{j,b} \rangle = \mathbb{P}_{x \sim \mu_S} \{x_i = a, x_j = b\} \quad S \subseteq \mathcal{V}, |S| \leq r, i, j \in S, a, b \in [q], && (8.3) \\ & \left\| \text{margin}_{A \cap B} \mu_A - \text{margin}_{A \cap B} \mu_B \right\|_1 = 0 \quad A, B \subseteq \mathcal{V}, |A|, |B| \leq r. && (8.4) \\ & \mu_S \in \Delta([q]^S) && (8.5) \end{aligned}$$

Remark 8.8. The SA_r relaxation is closely related to the r^{th} level of the Sherali–Adams hierarchy. In fact, SA_r is obtained from the basic SDP relaxation by r -rounds Sherali–Adams lift-and-project.

8. Limits of Semidefinite Relaxations

8.3.3. Robustness

Here we state the formal claims regarding robustness for the hierarchies LH_r and SA_r . The proofs are in [Section 8.9](#).

Definition 8.9. An SDP solution $\{\mathbf{v}_{i,a}\}_{i \in \mathcal{V}, a \in \mathbb{F}_q}, \{\mu_S\}_{S \subseteq \mathcal{V}, |S| \leq r}$ is said to be ε -infeasible for LH_r (or SA_r) if it satisfies all the constraints of the program up to an additive error of ε .

Theorem 8.10. Given an ε -infeasible solution $\{\mathbf{b}_{i,a}\}_{i \in \mathcal{V}, a \in \mathbb{F}_q}, \{\mu_S\}_{S \subseteq \mathcal{V}, |S| \leq r}$ to the LH_r relaxation, there exists a feasible solution $\{\mathbf{b}'_{i,a}\}, \{\mu'_S\}_{S \subseteq \mathcal{V}, |S| \leq r}$ for LH_r such that for all subsets $S \subseteq \mathcal{V}, |S| \leq r$, $\|\mu_S - \mu'_S\|_1 \leq \text{poly}(q) \cdot r^2 \varepsilon$.

Theorem 8.11. Given an ε -infeasible solution $\{\mathbf{b}_{i,a}\}_{i \in \mathcal{V}, a \in \mathbb{F}_q}, \{\mu_S\}_{S \subseteq \mathcal{V}, |S| \leq r}$ to the SA_r relaxation, there exists a feasible solution $\{\mathbf{b}'_{i,a}\}, \{\mu'_S\}_{S \subseteq \mathcal{V}, |S| \leq r}$ for SA_r such that for all subsets $S \subseteq \mathcal{V}, |S| \leq r$, $\|\mu_S - \mu'_S\|_1 \leq \text{poly}(q) \cdot \varepsilon \cdot q^r$.

8.4. Weak Gaps for UNIQUE GAMES

We refer to an integrality gap instance for a fairly simple SDP relaxation of **UNIQUE GAMES** as a *weak gap* instance. Formally, a *weak gap* instance for Unique games is defined as follows.

Definition 8.12. (*Weak SDP solutions and weak gap instances*) Let $\Upsilon = (V, E, \{\pi_e: [n] \rightarrow [n]\}_{e \in E})$. We say a collection $\mathcal{B} = \{B_u\}_{u \in V}$ is a *weak SDP solution* of value $1 - \eta$ for Υ if the following conditions hold:

1. (**Orthonormality**) For every vertex $u \in V$, the collection \mathcal{B} contains an ordered set $B_u = \{b_{u,1}, \dots, b_{u,n}\}$ of n orthonormal vectors in \mathbb{R}^d .
2. (**ℓ_2^2 -triangle inequality**) Any two vectors in $\bigcup \mathcal{B}$ have non-negative inner product and any three vectors in $\bigcup \mathcal{B}$ satisfy the ℓ_2^2 -triangle inequality ($\|x - y\|^2 \leq \|x - z\|^2 + \|z - y\|^2$).
3. (**Strong Matching Property**) For every pair of vertices $u, v \in V$, the sets B_u and B_v satisfy the following *strong matching property*: There exists n disjoint matchings between B_u, B_v given by bijections $\pi^{(1)}, \dots, \pi^{(n)}: B_u \rightarrow B_v$ such that for all $i \in [n], b, b' \in B_u$, we have $\langle b, \pi^{(i)}(b) \rangle = \langle b', \pi^{(i)}(b') \rangle$.
4. (**High SDP value**) For every edge $e = (u, v) \in E$, the vector sets B_u and B_v have significant correlation under the permutation $\pi = \pi_e$. Specifically, $\langle b_{u,\ell}, b_{v,\pi(\ell)} \rangle^2 \geq 0.99$ for all $\ell \in [n]$.

8. Limits of Semidefinite Relaxations

5. The collection \mathcal{B} of orthonormal sets is a good SDP solution for Υ , in the sense that

$$\mathbb{E}_{v \in V} \mathbb{E}_{\substack{w, w' \in N(v) \\ \pi = \pi_{w,v}, \pi' = \pi_{w',v}}} \frac{1}{n} \sum_{\ell \in [n]} \langle b_{w, \pi(\ell)}, b_{w', \pi'(\ell)} \rangle \geq 1 - \eta.$$

We say that Υ is a *weak* $(1 - \eta, \delta)$ -gap instance of UNIQUE GAMES if Υ has a weak SDP solution of value $1 - \eta$ and no labeling for Υ satisfies more than a δ fraction of the constraints.

Remark 8.13. The weak gap instances defined here are fairly natural objects. In fact, if \mathfrak{F} is an instance of Γ -Max-2Lin(R) with $\text{sdp}(\mathfrak{F}) \geq 1 - \eta$ and $\text{opt}(\mathfrak{F}) \leq \delta$, it is easy to construct a corresponding weak gap instance \mathfrak{F}' . The idea is to start with an optimal SDP solution for \mathfrak{F} , symmetrize it (with respect to the group Φ), and delete all edges of \mathfrak{F} that contribute less than $\sqrt{3/4}$ to the SDP objective.

We observe the following consequence of Fact 8.18 and item 4 of Definition 8.12.

Observation 8.14. *If $\mathcal{B} = \{B_u\}_{u \in V}$ is a weak SDP solution for $\Phi = (V, E, \{\pi_e\}_{e \in E})$, then for any two edges $(w, v), (w', v) \in E$, the two bijections $\pi = \pi_{(w',v)}^{-1} \circ \pi_{(w,v)}$ and $\pi_{B_{w'} \leftarrow B_w}$ (see Def. 8.17) give rise to the same matching between the vector sets B_w and $B_{w'}$,*

$$\pi(i) = j \iff \pi_{B_{w'} \leftarrow B_w}(b_{w,i}) = b_{w',j}.$$

The previous observation implies that in a weak gap instance Φ the collection of permutations $\{\pi_e\}_{e \in E}$ is already determined by the geometry of the vector sets in a weak SDP solution \mathcal{B} .

There are a few explicit constructions of weak gap instances of UNIQUE GAMES, most prominently the Khot–Vishnoi instance [KV05]. In particular, the following observation is a restatement of Theorem 9.2 and Theorem 9.3 in [KV05].

Observation 8.15. *For all $\eta, \delta > 0$, there exists a weak $(1 - \eta, \delta)$ -gap instance with $2^{2^{O(\log(1/\delta)/\eta)}}$ vertices.*

8.4.1. Properties of Weak Gap Instances

Observation 8.35 implies that without much loss we can assume that a weak SDP solution is \mathbb{F}_q -integral, that is, all vectors are \mathbb{F}_q -integral. Here we use again $\langle \cdot, \cdot \rangle_\psi := \langle \psi(\cdot), \psi(\cdot) \rangle$ as inner product for \mathbb{F}_q -integral vectors.

8. Limits of Semidefinite Relaxations

Lemma 8.16. *Let $\Phi = (V, E, \{\pi_e\}_{e \in E})$ be a weak $(1 - \eta, \delta)$ -gap instance. Then, for every $q \in \mathbb{N}$, we can find a weak \mathbb{F}_q -integral SDP solution of value $1 - O(\sqrt{\eta \log q})$ for a UNIQUE GAMES instance Φ' which is obtained from Φ by deleting $O(\sqrt{\eta \log q})$ edges.*

Proof. Let \mathcal{B} be a weak SDP solution for Φ of value $1 - \eta$. By applying the transformation from Observation 8.35 to the vectors in \mathcal{B} , we obtain a collection $\mathcal{B}' = \{B'_u\}_{u \in V}$ of sets of \mathbb{F}_q -integral vectors. For every $u \in V$, the vectors in B'_u are orthonormal. Furthermore, any two sets B'_u, B'_v in \mathcal{B}' satisfy the strong matching property (using the facts that the original sets B_u, B_v satisfy this property and that $\langle b'_{u,i}, b'_{v,j} \rangle_\psi$ is a function of $\langle b_{u,i}, b_{v,j} \rangle$).

Let $\eta_{v,w,w',\ell} = 1 - \langle b_{w,\pi(\ell)}, b_{w',\pi'(\ell)} \rangle$. Using Jensen's inequality, we can verify that the value of the SDP solution \mathcal{B}' is high,

$$\begin{aligned} & \mathbb{E}_{v \in V} \mathbb{E}_{\substack{w, w' \in N(v) \\ \pi = \pi_{w,v}, \pi' = \pi_{w',v}}} \frac{1}{R} \sum_{\ell \in [R]} \langle b'_{w,\pi(\ell)}, b'_{w',\pi'(\ell)} \rangle_\psi \\ & \geq \mathbb{E}_{v \in V} \mathbb{E}_{\substack{w, w' \in N(v) \\ \pi = \pi_{w,v}, \pi' = \pi_{w',v}}} \frac{1}{R} \sum_{\ell \in [R]} 1 - O\left(\sqrt{\eta_{v,w,w',\ell} \log q}\right) \quad (\text{by Obs. 8.35}) \\ & \geq 1 - O(\sqrt{\eta \log q}) \quad (\text{using Jensen's inequality}). \end{aligned}$$

So far, we verified that \mathcal{B}' satisfies all requirements of a weak SDP solution besides item 4 of Definition 8.12. We can ensure that this condition is also satisfied by deleting all edges from E where the condition is violated. Using standard averaging arguments, it is easy to see that the matching property and the high SDP value imply that this condition is satisfied for all but at most an $O(\sqrt{\eta \log q})$ fraction of edges. \square

We will refer to the set of orthonormal vectors associated with a vertex B as a *cloud*. In what follows, we identify the vertices B in a *weak* gap instance with their corresponding clouds, and thus refer to vertices/clouds interchangeably.

Definition 8.17. For $A, B \in \mathcal{B}$, we denote

$$\rho(A, B) \stackrel{\text{def}}{=} \max_{a \in A, b \in B} |\langle a, b \rangle|.$$

We define $\pi_{B \leftarrow A}: A \rightarrow B$ to be any¹ bijection from A to B such that $|\langle a, \pi_{B \leftarrow A}(a) \rangle| = \rho(A, B)$ for all $a \in A$.

¹The matching property asserts that such a matching exists. If it is not unique, we pick an arbitrary one. We will assume $\pi_{A \rightarrow B} = \pi_{B \rightarrow A}^{-1}$.

8. Limits of Semidefinite Relaxations

As a direct consequence of the orthogonality of the clouds in \mathcal{B} , we have the following fact about the uniqueness of $\pi_{B \leftarrow A}$ for highly correlated clouds $A, B \in \mathcal{B}$.

Fact 8.18. *Let $A, B \in \mathcal{B}$. If $\rho(A, B)^2 > 3/4$, then there exists exactly one bijection $\pi: A \rightarrow B$ such that $|\langle a, \pi(a) \rangle| = \rho(A, B)$ for all $a \in A$.*

Remark 8.19. The collection \mathcal{B} succinctly encodes a UNIQUE GAMES instance. For a graph $G = (\mathcal{B}, E)$ on \mathcal{B} , the goal is to find a labeling $\{\ell_A \in A\}_{A \in \mathcal{B}}$ (a labeling can be seen as a system of representatives for the clouds in \mathcal{B}) so as to maximize the probability

$$\mathbb{P}_{(A, B) \in E} \left\{ \ell_A = \pi_{A \leftarrow B}(\ell_B) \right\}.$$

Tensoring

Lemma 8.20. *For $t \in \mathbb{N}$ and every pair of clouds $A, B \in \mathcal{B}$,*

$$\frac{1}{n} \sum_{\substack{a \in A, b \in B \\ a \neq \pi_{A \leftarrow B}(b)}} |\langle a, b \rangle|^t \leq 2 \cdot (3/4)^{t/2}.$$

Proof. By near-orthogonality, $\sum_{a \in B} \langle a, b \rangle^2 \leq 3/2$ for every $b \in B$. Hence, $\langle a, b \rangle^2 \leq 3/4$ for all $a \neq \pi_{A \leftarrow B}(b)$. Thus,

$$\frac{1}{n} \sum_{\substack{a \in A, b \in B \\ a \neq \pi_{A \leftarrow B}(b)}} |\langle a, b \rangle|^t \leq (3/4)^{t/2} \cdot \frac{1}{n} \sum_{a \in A, b \in B} |\langle a, b \rangle|^2 \leq (3/4)^{t/2} \cdot 3/2.$$

□

The notation $X = Y \pm Z$ means that $|X - Y| \leq Z$.

Corollary 8.21. *For $t \in \mathbb{N}$ and every pair of clouds $A, B \in \mathcal{B}$,*

$$\frac{1}{n} \sum_{a \in A, b \in B} \langle a, b \rangle^t = \frac{1}{n} \sum_{a \in A} \langle a, \pi_{B \leftarrow A}(a) \rangle^t \pm 2 \cdot (3/4)^{t/2}.$$

Remark 8.22. The left-hand side in the corollary is the inner product of the vectors $1/\sqrt{n} \sum_{u \in A} u^{\otimes t}$ and $1/\sqrt{n} \sum_{v \in B} v^{\otimes t}$. If t is even, then we can replace the right-hand side by $\rho(A, B)^t$. This fact that the functional $\rho(A, B)^t$ is closely approximated by inner products averaged-tensored vectors has implicitly been used in [KV05] and was explicitly noted in [AKK⁺08, Lemma 2.2].

8.5. Integrality Gap Instance for UNIQUE GAMES

In this section, we will exhibit the construction strong SDP integrality gap for the E2Lin_q problem. Recall that the E2Lin_q problem is a special case of UNIQUE GAMES . To this end, we follow the approach of Khot–Vishnoi [KV05] to construct the gap instance.

Khot et al. [KKMO07] show a UGC-based hardness result for the E2Lin_q problem. Specifically, they exhibit a reduction $\Phi_{\gamma,q}$ that maps a UNIQUE GAMES instance Φ to an E2Lin_q instance $\Phi_{\gamma,q}(\Phi)$ such that the following holds: For every $\gamma > 0$ and all $q \geq q_0(\gamma)$,

Completeness: If Φ is $1 - \eta$ -satisfiable then $\Phi_{\gamma,q}(\Phi)$ is $1 - \gamma - o_{\eta,\delta}(1)$ satisfiable.

Soundness: If Φ has no labeling satisfying more than δ -fraction of the constraints, then no assignment satisfies more than $q^{-\eta/2} + o_{\eta,\delta}(1)$ -fraction of equations in $\Phi_{\gamma,q}(\Phi)$.

Here the notation $o_{\eta,\delta}(1)$ refers to any function that tends to 0 whenever η and δ go to naught. The details of the $\Phi_{\gamma,q}$ reduction are included in Figure 8.1 for the sake of completeness.

The rest of the chapter is devoted to the proof of the following theorem.

Theorem 8.23. *Let Φ be a weak $(1 - \eta, \delta)$ -gap instance of UNIQUE GAMES . Then, for every q of order unity, there exists an SDP solution for the E2Lin_q instance $\Phi_{\gamma,q}(\Phi)$ such that*

- the SDP solution is feasible for LH_r with $r = 2^{\Omega(1/\eta^{1/4})}$,
- the SDP solution is feasible for SA_r with $r = \Omega(\eta^{1/4})$,
- the SDP solution has value $1 - \gamma - o_{\eta,\delta}(1)$ for $\Phi_{\gamma,q}(\Phi)$.

In particular, the E2Lin_q instance $\Phi_{\gamma,q}(\Phi)$ is a $(1 - \gamma - o_{\eta,\delta}(1), q^{-\eta/2} + o_{\eta,\delta}(1))$ integrality gap instance for the relaxation LH_r for $r = 2^{\Omega(1/\eta^{1/4})}$. Further, $\Phi_{\gamma,q}(\Phi)$ is a $(1 - \gamma - o_{\eta,\delta}(1), q^{-\eta/2} + o_{\eta,\delta}(1))$ integrality gap instance for the relaxation SA_r for $r = \Omega(1/\eta^{1/4})$.

8.6. Integral Vectors

In this section, we will develop tools to create and manipulate vectors all of whose coordinates are “integral”.

8. Limits of Semidefinite Relaxations

E2Lin_q Hardness Reduction [KKMO07]

Input A UNIQUE GAMES instance Φ with vertex set V , edge set $E \subseteq V \times V$ (we assume the graph (V, E) to be regular), and permutations $\{\pi_e: [R] \rightarrow [R]\}_{e \in E}$.

Output An E2Lin_q instance $\Phi_{\gamma, q}(\Phi)$ with vertex set $\mathcal{V} = V \times \mathbb{F}_q^R$. Let $\{\mathcal{F}_v: \mathbb{F}_q^R \rightarrow \mathbb{F}_q\}_{v \in V}$ denote an \mathbb{F}_q -assignment to \mathcal{V} . The constraints of $\Phi_{\gamma, q}(\Phi)$ are given by the tests performed by the following probabilistic verifier:

- Pick a random vertex $v \in V$. Choose two random neighbours $w, w' \in N(v) \subseteq V$. Let π, π' denote the permutations on the edges (w, v) and (w', v) .
- Sample $x \in \mathbb{F}_q^R$ uniformly at random. Generate $y \in \mathbb{F}_q^R$ as follows:

$$y_i = \begin{cases} x_i & \text{with probability } 1 - \gamma \\ \text{uniform random element from } \mathbb{F}_q & \text{with probability } \gamma \end{cases}$$
- Generate a uniform random element $c \in \mathbb{F}_q$.
- Test if $\mathcal{F}_w(y \circ \pi + c \cdot \mathbf{1}) = \mathcal{F}_{w'}(x \circ \pi') + c$. (Here, $x \circ \pi$ denotes the vector $(x_{\pi(i)})_{i \in [R]}$.)

Figure 8.1.: Reduction from UNIQUE GAMES to E2Lin_q

8. Limits of Semidefinite Relaxations

$\{\pm 1\}$ -integral vectors. We begin by defining our notion of a $\{\pm 1\}$ -integral vector.

Definition 8.24. Let $\mathcal{R} = (\Omega, \mu)$ be a probability space. A function $u: \mathcal{R} \rightarrow \{\pm 1\}$ is called an $\{\pm 1\}$ -integral vector. In other words, u is a $\{\pm 1\}$ -valued random variable defined on the probability space \mathcal{R} . We define an inner product of functions $u, v: \mathcal{R} \rightarrow \{\pm 1\}$ by

$$\langle u, v \rangle = \mathbb{E}_{r \sim \mathcal{R}} u(r)v(r).$$

In our construction, we often start with $\{\pm 1\}$ -integral vectors given by the hypercube $\{\pm 1\}^R$. In the terminology of $\{\pm 1\}$ -integral vectors, we can think of the hypercube $\{\pm 1\}^R$ as the set of $\{\pm 1\}$ -integral vectors where \mathcal{R} is the uniform distribution over $\{1, \dots, R\}$.

The following lemma shows how the Goemans–Williamson [GW95] rounding scheme can be thought of as a procedure to “round” arbitrary real vectors to $\{\pm 1\}$ -integral vectors.

Observation 8.25. Given a family of unit vectors $\{v_1, \dots, v_R\} \in \mathbb{R}^d$, define the set of $\{\pm 1\}$ -valued functions $v_1^*, \dots, v_R^*: \mathcal{R} \rightarrow \{\pm 1\}$ with $\mathcal{R} = \mathcal{G}^d$ - the Gaussian space of appropriate dimension as follows:

$$v_i^*(g) = \text{sign}(\langle v_i, g \rangle)$$

for $g \in \mathcal{G}^d$. The $\{\pm 1\}$ -valued functions $\{v_i^*\}$ satisfy $\langle v_1^*, v_2^* \rangle = 2 \arccos(\langle v_1, v_2 \rangle) / \pi$. Specifically, this operation obeys the following properties:

$$\langle u, v \rangle = 0 \iff \langle u^*, v^* \rangle = 0 \quad \langle u, v \rangle = 1 - \varepsilon \implies \langle u^*, v^* \rangle \geq 1 - O(\sqrt{\varepsilon})$$

The tensor product operation on $\{\pm 1\}$ -integral vectors, yields a $\{\pm 1\}$ -integral vector.

Definition 8.26. Given two $\{\pm 1\}$ -valued functions $u: \mathcal{R}_1 \rightarrow \{\pm 1\}$ and $v: \mathcal{R}_2 \rightarrow \{\pm 1\}$, the tensor product $u \otimes v: \mathcal{R}_1 \times \mathcal{R}_2 \rightarrow \{\pm 1\}$ is defined as $u \otimes v(r_1, r_2) = u(r_1)v(r_2)$.

Observation 8.27. For $u, u': \mathcal{R}_1 \rightarrow \{\pm 1\}$ and $v, v': \mathcal{R}_2 \rightarrow \{\pm 1\}$, we have

$$\begin{aligned} \langle u \otimes v, u' \otimes v' \rangle &= \mathbb{E}_{r_1, r_2} [u \otimes v(r_1, r_2) u' \otimes v'(r_1, r_2)] \\ &= \mathbb{E}_{r_1} [u(r_1) u'(r_1)] \mathbb{E}_{r_2} [v(r_2) v'(r_2)] = \langle u, u' \rangle \langle v, v' \rangle \end{aligned}$$

8. Limits of Semidefinite Relaxations

\mathbb{F}_q -integral vectors. Let q be a prime. Now, we will define \mathbb{F}_q -integral vectors and their tensor products.

Definition 8.28. A \mathbb{F}_q -integral vector $v: \mathcal{R} \rightarrow \mathbb{F}_q$ is a function from a measure space \mathcal{R} to \mathbb{F}_q . For a \mathbb{F}_q -integral vector $v: \mathcal{R} \rightarrow \mathbb{F}_q$, its symmetrization $\tilde{v}: \mathcal{R} \times \mathbb{F}_q^* \rightarrow \mathbb{F}_q$ is defined by $\tilde{v}(r, t) = t \cdot v(r)$.

Given a map $f: \mathbb{F}_q \rightarrow \mathbb{C}^d$, we denote by $f(v) := f \circ v$ the composition of functions f and v . Here are few examples of functions that will be relevant to us:

1. The function $\chi: \mathbb{F}_q \rightarrow \mathbb{C}^{q-1}$ given by

$$\chi(i) \stackrel{\text{def}}{=} \frac{1}{\sqrt{q-1}}(\omega^{1 \cdot i}, \dots, \omega^{j \cdot i}, \dots, \omega^{(q-1)i}),$$

where ω is a primitive q^{th} root of unity. The vector $\chi(i) \in \mathbb{C}^{q-1}$ is the restriction of the i^{th} character function of the group \mathbb{Z}_q to the set \mathbb{F}_q^* . It is easy to see that

$$\langle \chi(a), \chi(b) \rangle = \mathbb{E}_{t \in \mathbb{F}_q^*} \left[\omega^{ta} \cdot \omega^{-tb} \right] = \begin{cases} 1 & \text{if } a = b, \\ -\frac{1}{q-1} & \text{if } a \neq b. \end{cases}$$

2. Let $\psi_0, \psi_1, \dots, \psi_{q-1}$ denote the corners of the q -ary simplex in \mathbb{R}^{q-1} , translated so that the origin is its geometric center. Define the function $\psi: \mathbb{F}_q \rightarrow \mathbb{R}^{q-1}$ as $\psi(i) := \psi_i$. Again, the vectors satisfy

$$\langle \psi(a), \psi(b) \rangle = \begin{cases} 1 & \text{if } a = b, \\ -\frac{1}{q-1} & \text{if } a \neq b. \end{cases}$$

Remark 8.29. A \mathbb{F}_q -integral vector $v \in \mathbb{F}_q^N$ can be thought of as a \mathbb{F}_q -valued function over the measure space $([N], \mu)$ where μ is the uniform distribution over $[N]$.

Remark 8.30. The following notions are equivalent: Collection of \mathbb{F}_q -valued functions on some measure space $\mathcal{R} \iff$ Collection of jointly-distributed, \mathbb{F}_q -valued random variables \iff Distribution over \mathbb{F}_q -assignments.

For the case of \mathbb{F}_q -integral vector, the tensor product operation is to be defined carefully, in order to mimic the properties of the traditional tensor product. We will use the following definition for the tensor operation \otimes_q .

8. Limits of Semidefinite Relaxations

Definition 8.31. Given two \mathbb{F}_q -valued functions $u: \mathcal{R} \rightarrow \mathbb{F}_q$ and $u': \mathcal{R}' \rightarrow \mathbb{F}_q$, define the symmetrized tensor product $u \otimes_q u': (\mathcal{R} \times \mathbb{F}_q^*) \times (\mathcal{R}' \times \mathbb{F}_q^*) \rightarrow \mathbb{F}_q$ as

$$(u \otimes_q u')(r, t, r', t') \stackrel{\text{def}}{=} t \cdot u(r) + t' \cdot u'(r').$$

Lemma 8.32. For any \mathbb{F}_q -valued functions $u, v: \mathcal{R} \rightarrow \mathbb{F}_q$ and $u', v': \mathcal{R}' \rightarrow \mathbb{F}_q$,

$$\langle \psi(u \otimes_q u'), \psi(v \otimes_q v') \rangle = \langle \psi(u), \psi(v) \rangle \langle \psi(u'), \psi(v') \rangle.$$

Proof.

$$\begin{aligned} & \langle \psi(u \otimes_q u'), \psi(v \otimes_q v') \rangle \\ &= \langle \chi(u \otimes_q u'), \chi(v \otimes_q v') \rangle \quad (\text{using } \langle \psi_a, \psi_b \rangle = \langle \chi(a), \chi(b) \rangle) \\ &= \mathbb{E}_{(r,t)} \mathbb{E}_{(r',t')} \mathbb{E}_{\ell \in \mathbb{F}_q^*} \omega^{\ell t u(r) + \ell t' u'(r')} \cdot \omega^{-\ell t v(r) - \ell t' v'(r')} \quad (\text{by definitions of } \otimes_q \text{ and } \chi) \\ &= \mathbb{E}_{\ell \in \mathbb{F}_q^*} \left(\mathbb{E}_{(r,t)} \omega^{\ell t u(r) - \ell t v(r)} \right) \cdot \left(\mathbb{E}_{(r',t')} \omega^{\ell t' u'(r') - \ell t' v'(r')} \right) \\ &= \mathbb{E}_{\ell \in \mathbb{F}_q^*} \left(\mathbb{E}_r \langle \chi(\ell u(r)), \chi(\ell v(r)) \rangle \right) \cdot \left(\mathbb{E}_{r'} \langle \chi(\ell u'(r')), \chi(\ell v'(r')) \rangle \right) \\ &= \mathbb{E}_{\ell \in \mathbb{F}_q^*} \langle \chi(\ell u), \chi(\ell v) \rangle \langle \chi(\ell u'), \chi(\ell v') \rangle \\ &= \langle \chi(u), \chi(v) \rangle \langle \chi(u'), \chi(v') \rangle \quad (\text{using } \langle \chi(\ell a), \chi(\ell b) \rangle = \langle \chi(a), \chi(b) \rangle \text{ for } \ell \in \mathbb{F}_q^*) \\ &= \langle \psi(u), \psi(v) \rangle \langle \psi(u'), \psi(v') \rangle \quad (\text{using } \langle \psi_a, \psi_b \rangle = \langle \chi(a), \chi(b) \rangle) \end{aligned}$$

□

Remark 8.33. Unlike the ordinary tensor operation, the q -ary tensor operation we defined is not associative. Formally, we define the tensoring operation to be right-associative

$$u_1 \otimes_q u_2 \otimes_q \dots \otimes_q u_{k-1} \otimes_q u_k \stackrel{\text{def}}{=} u_1 \otimes_q \left(u_2 \otimes_q \left(\dots (u_{k-1} \otimes_q u_k) \dots \right) \right).$$

The lack of associativity will never be an issue in our constructions.

We need the following simple technical observation in one of our proofs.

Observation 8.34. Let $u, v: \mathcal{R} \rightarrow \mathbb{F}_q$ be two “symmetric” \mathbb{F}_q -integral vectors. that is, $\mathbb{P}_r\{u(r) - v(r) = a\} = \mathbb{P}_r\{u(r) - v(r) = b\}$ for all $a, b \in \mathbb{F}_q^*$. Then, for all $a, b \in \mathbb{F}_q$, we have $\mathbb{E}_r \langle \psi(a + u(r)), \psi(b + v(r)) \rangle = \langle a \otimes u, b \otimes v \rangle$.

8. Limits of Semidefinite Relaxations

We wish to point out that in our applications, the vectors u and v will be tensor powers. In this case, the symmetry condition is always satisfied.

Proof. Using the symmetry assumption, we see that

$$\begin{aligned} & \mathbb{P}_{r \sim \mathcal{R}, t, t' \in \mathbb{F}_q^*} \left\{ ta + t'u(r) = tb + t'v(r) \right\} \\ &= \mathbb{P}_{r \sim \mathcal{R}, t \in \mathbb{F}_q^*} \left\{ a - b = t \cdot (v(r) - u(r)) \right\} \\ &= \mathbb{P}_{r \sim \mathcal{R}} \left\{ a - b = v(r) - u(r) \right\} \end{aligned} \quad (8.6)$$

If we let ρ denote this probability, then we have $\langle a \otimes u, b \otimes v \rangle = \rho - (1 - \rho)/(q - 1)$ (using the left-hand side of Eq. (8.6) as well as $\mathbb{E}_r \langle \psi(a + u(r)), \psi(b + v(r)) \rangle = \rho - (1 - \rho)/(q - 1)$ (using the right-hand side of Eq. (8.6)). \square

The following procedure yields a way to generate \mathbb{F}_q -integral vectors from arbitrary vectors. The transformation is inspired by the rounding scheme for UNIQUE GAMES in Charikar et al. [CMM06a].

Observation 8.35. *Define the function $\zeta: \mathcal{G}^q \rightarrow \mathbb{F}_q$ on the Gaussian domain as follows:*

$$\zeta(x_1, \dots, x_q) = \operatorname{argmax}_{i \in [q]} x_i \quad (8.7)$$

Given a family of unit vectors $\{v_1, \dots, v_R\} \in \mathbb{R}^d$, define the set of \mathbb{F}_q -valued functions $v_1^, \dots, v_R^*: \mathcal{R} \rightarrow \mathbb{F}_q$ with $\mathcal{R} = (\mathcal{G}^d)^q$ —the Gaussian space of appropriate dimension—as follows:*

$$v_i^*(g_1, \dots, g_q) = \zeta(\langle v_i, g_1 \rangle, \dots, \langle v_i, g_q \rangle)$$

for $g_1, \dots, g_q \in (\mathcal{G}^d)^q$. The \mathbb{F}_q -valued functions $\{v_i^\}$ satisfy,*

1. $\langle u, v \rangle = 0 \implies \langle \psi(u^*), \psi(v^*) \rangle = 0$,
2. $\langle u, v \rangle = 1 - \varepsilon \implies \langle \psi(u^*), \psi(v^*) \rangle = 1 - f(\varepsilon, q) = 1 - O(\sqrt{\varepsilon \log q})$.

Proof. To see (1), observe that if $\langle u, v \rangle = 0$, then the sets of random variables $\{\langle u, g_1 \rangle, \dots, \langle u, g_q \rangle\}$ and $\{\langle v, g_1 \rangle, \dots, \langle v, g_q \rangle\}$ are completely independent of each other. Therefore,

$$\langle \psi(u^*), \psi(v^*) \rangle = \mathbb{E}_{r \in \mathcal{G}^{dq}} \left[\psi(u^*(r)) \right] \cdot \mathbb{E}_{r \in \mathcal{G}^{dq}} \left[\psi(v^*(r)) \right] = 0.$$

Assertion 2 follows from Lemma C.8 in [CMM06a]. \square

8.7. Local Distributions for UNIQUE GAMES

In this section, we will construct local distribution over labelings to a UNIQUE GAMES instance.

The following facts are direct consequences of the (symmetrized) ℓ_2^2 -triangle inequality.

Fact 8.36. *Let $a, b, c \in \cup \mathcal{B}$ with $|\langle a, b \rangle| = 1 - \eta_{ab}$ and $|\langle b, c \rangle| = 1 - \eta_{bc}$. Then, $|\langle a, c \rangle| \geq 1 - \eta_{ab} - \eta_{bc}$.*

Fact 8.37. *Let $A, B, C \in \mathcal{B}$ with $\rho(A, B) = 1 - \eta_{AB}$ and $\rho(B, C) = 1 - \eta_{BC}$. Then, $\rho(A, C) \geq 1 - \eta_{AB} - \eta_{BC}$.*

The construction in the proof of the next lemma is closely related to propagation-style UG algorithms [Tre05, AKK⁺08].

Definition 8.38. A set $\mathcal{S} \subseteq \mathcal{B}$ is *consistent* if

$$\forall A, B \in \mathcal{S}. \quad \rho(A, B) \geq 1 - 1/16.$$

Lemma 8.39. *If $\mathcal{S} \subseteq \mathcal{B}$ is consistent, there exists bijections $\{\pi_A: [R] \rightarrow A\}_{A \in \mathcal{S}}$ such that*

$$\forall A, B \in \mathcal{S}. \quad \pi_B = \pi_{B \leftarrow A} \circ \pi_A.$$

Proof. We can construct the bijections in a greedy fashion: Start with an arbitrary cloud $C \in \mathcal{S}$ and choose an arbitrary bijection $\pi_C: [R] \rightarrow C$. For all other clouds $B \in \mathcal{S}$, choose $\pi_B := \pi_{B \leftarrow C} \circ \pi_C$.

Let A, B be two arbitrary clouds in \mathcal{S} . Let $\sigma_{A \leftarrow B} := \pi_A \circ \pi_B^{-1}$. To prove the lemma, we have to verify that $\sigma_{A \leftarrow B} = \pi_{A \leftarrow B}$. By construction, $\sigma_{A \leftarrow B} = \pi_{A \leftarrow C} \circ \pi_{C \leftarrow B}$. Let $\eta = 1/16$. Since $\rho(A, C) \geq 1 - \eta$ and $\rho(B, C) \geq 1 - \eta$, we have $|\langle b, \sigma_{A \leftarrow B}(b) \rangle| \geq 1 - 2\eta$ for all $b \in B$ (using Fact 8.36). Since $(1 - 2\eta)^2 > 1 - 4\eta = 3/4$, Fact 8.18 (uniqueness of bijection) implies that $\sigma_{A \leftarrow B} = \pi_{A \leftarrow B}$. \square

Hence, for a consistent set of clouds \mathcal{S} , the distribution over local UNIQUE GAMES labelings $\mu_{\mathcal{S}}$ can be defined easily as follows:

Sample $\ell \in [R]$ uniformly at random, and for every cloud $A \in \mathcal{S}$,
assign $\pi_A(\ell)$ as label.

To construct a local distribution for a set \mathcal{S} which is not consistent, we partition the set \mathcal{S} into consistent clusters. To this end, we make the following definition:

8. Limits of Semidefinite Relaxations

Definition 8.40. A set $\mathcal{S} \subseteq \mathcal{B}$ is *consistent* with respect to a partition P of \mathcal{B} (denoted $\text{Cons}(\mathcal{S}, P)$) if

$$\forall \mathcal{C} \in P. \quad \forall A, B \in \mathcal{C} \cap \mathcal{S}. \quad \rho(A, B) \geq 1 - 1/16.$$

We use $\text{Incons}(\mathcal{S}, P)$ to denote the event that \mathcal{S} is not consistent with P . The following is a corollary of Lemma 8.39.

Corollary 8.41. Let P be a partition of \mathcal{B} and let $\mathcal{S} \subseteq \mathcal{B}$. If $\text{Cons}(\mathcal{S}, P)$, then there exists bijections $\{\pi_A: [R] \rightarrow A \mid A \in \mathcal{S}\}$ such that

$$\forall \mathcal{C} \in P. \quad \forall A, B \in \mathcal{C} \cap \mathcal{S}. \quad \pi_B = \pi_{B \leftarrow A} \circ \pi_A.$$

The following lemma relies on the fact that the correlations $\rho(A, B)$ behave up to a small errors like inner products of real vectors. In other words, there is a geometric representation of the correlations $\rho(A, B)$ that can be used for the decomposition. This insight has also been used in UG algorithms [AKK⁺08].

Lemma 8.42. For every $t \in \mathbb{N}$, there exists a distribution over partitions P of \mathcal{B} such that

– if $\rho(A, B) \geq 1 - \varepsilon$, then

$$\mathbb{P}\{P(A) = P(B)\} \geq 1 - O(t\sqrt{\varepsilon}).$$

– if $\rho(A, B) \leq 1 - 1/16$, then

$$\mathbb{P}\{P(A) = P(B)\} \leq (3/4)^t.$$

Proof. Let $s \in \mathbb{N}$ be even and large enough (we will determine the value of s later). For every set $B \in \mathcal{B}$, define a vector $\mathbf{v}_B \in \mathbb{R}^D$ with $D := d^s$ as

$$\mathbf{v}_B := \frac{1}{\sqrt{R}} \sum_{v \in B} v^{\otimes s}.$$

We consider the following distribution over partitions P of \mathcal{B} : Choose t random hyperplanes H_1, \dots, H_t through the origin in \mathbb{R}^D . Consider the partition of \mathbb{R}^D formed by these hyperplanes. Output the induced partition P of \mathcal{B} (two sets $A, B \in \mathcal{B}$ are in the same cluster of P if and only if \mathbf{v}_A and \mathbf{v}_B are not separated by any of the hyperplanes H_1, \dots, H_t).

8. Limits of Semidefinite Relaxations

Since s is even, Corollary 8.21 shows that for any two sets $A, B \in \mathcal{B}$,

$$\langle \mathbf{v}_A, \mathbf{v}_B \rangle = \rho(A, B)^s \pm 2 \cdot (3/4)^{-s/2}.$$

Furthermore, if $\rho(A, B) = 1 - \varepsilon$, then

$$\langle \mathbf{v}_A, \mathbf{v}_B \rangle \geq (1 - \varepsilon)^s \geq 1 - s\varepsilon.$$

Let $\eta = 1/16$. We choose s minimally such that $(1 - \eta)^s + 2 \cdot (3/4)^{-s/2} \leq 1/\sqrt{2}$. (So s is an absolute constant.) Then for any two sets $A, B \in \mathcal{B}$ with $\rho(A, B) \leq 1 - \eta$, their vectors have inner product $\langle \mathbf{v}_A, \mathbf{v}_B \rangle \leq 1/\sqrt{2}$. Thus, a random hyperplane through the origin separates \mathbf{v}_A and \mathbf{v}_B with probability at least $1/4$. Therefore,

$$\mathbb{P}\{P(A) = P(B)\} \leq (3/4)^t.$$

On the other hand, if $\rho(A, B) = 1 - \varepsilon$, then the vectors of A and B have inner product $\langle \mathbf{v}_A, \mathbf{v}_B \rangle \geq 1 - s\varepsilon$. Thus, a random hyperplane through the origins separates the vectors with probability at most $O(\sqrt{\varepsilon})$. Hence,

$$\mathbb{P}\{P(A) = P(B)\} \geq \left(1 - O(\sqrt{\varepsilon})\right)^t \geq 1 - O(t\sqrt{\varepsilon}).$$

□

Remark 8.43. Using a more sophisticated construction, we can improve the bound $1 - O(t\sqrt{\varepsilon})$ to $1 - O(\sqrt{t\varepsilon})$.

The previous lemma together with a simple union bound imply the next corollary.

Corollary 8.44. *The distribution over partitions from Lemma 8.42 satisfies the following property: For every set $\mathcal{S} \subseteq \mathcal{B}$,*

$$\mathbb{P}\{\text{Incons}(\mathcal{S}, P)\} \leq |\mathcal{S}|^2 \cdot (3/4)^t$$

Remark 8.45. Using a slightly more refined argument (triangle inequality), we could improve the bound $r^2 \cdot (3/4)^t$ to $r \cdot (3/4)^t$.

8.8. Construction of SDP Solutions for E2LIN(q)

In this section, we construct SDP vectors and local distributions for $\mathcal{B} \times \mathbb{F}_q^R$ that form the variables in the $\Phi_{\gamma, q}(\Phi)$ instance described in Section 8.5. The set $\mathcal{B} \times \mathbb{F}_q^R$ correspond to the set of vertices in the instance obtained by applying a q -ary long code based reduction on the UNIQUE GAMES instance encoded by \mathcal{B} . For a vertex $(B, x) \in \mathcal{B} \times \mathbb{F}_q^R$, we index the coordinates of x by the elements of B . Specifically, we have $x = (x_b)_{b \in B} \in \mathbb{F}_q^B$.

8. Limits of Semidefinite Relaxations

Geometric Partitioning. Apply Lemma 8.42 to the collection of sets of vectors \mathcal{B} . We obtain a distribution \mathcal{P} over partitions P of \mathcal{B} into T disjoint subsets $\{P_\alpha\}_{\alpha=1}^T$. For a subset $\mathcal{S} \subset \mathcal{B}$, let $\mathcal{S} = \{\mathcal{S}_\alpha\}_{\alpha=1}^T$ denote the partition induced on the set \mathcal{S} , that is, $\mathcal{S}_\alpha := P_\alpha \cap \mathcal{S}$. For a family $B \in \mathcal{B}$, let α_B denote the index of the set P_{α_B} in the partition P that contains B .

8.8.1. Vector Solution

For a vertex $(B, x) \in \mathcal{B} \times \mathbb{F}_q^R$, the corresponding SDP vectors are given by functions $\mathbf{V}_j^{B,x} : \mathcal{P} \times [T] \times \mathcal{R} \rightarrow \mathbb{R}^q$ defined as follows:

$$\mathbf{W}_j^{B,x}(r) = \frac{1}{\sqrt{R}} \sum_{b \in B} \psi(x_b - j + b^{\otimes t}(r)) \quad (8.8)$$

$$\mathbf{U}_j^{B,x}(P, \alpha, r) = P_\alpha(B) \cdot \mathbf{W}_j^{B,x}(r) \quad (8.9)$$

$$\mathbf{V}_j^{B,x} = \frac{1}{q} \mathbf{V}_0 + \frac{\sqrt{q-1}}{q} \mathbf{U}_j^{B,x} \quad (8.10)$$

Here \mathcal{R} is the measure space over which the tensored vectors $b^{\otimes t}$ are defined. The notation $P_\alpha(B)$ denotes the 0/1-indicator for the event $B \in P_\alpha$. Further, \mathbf{V}_0 is a unit vector orthogonal to all the vectors $\mathbf{U}_j^{B,x}$.

Let us evaluate the inner product between two vectors $\mathbf{V}_i^{A,x}$ and $\mathbf{V}_j^{B,y}$, (in this way, we also clarify the intended measure on the coordinate set)

$$\begin{aligned} \langle \mathbf{V}_i^{A,x}, \mathbf{V}_j^{B,y} \rangle &= \frac{1}{q^2} + \frac{q-1}{q^2} \langle \mathbf{U}_i^{A,x}, \mathbf{U}_j^{B,y} \rangle \\ &= \frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{E}_{P \sim \mathcal{P}} \sum_{\alpha=1}^T P_\alpha(A) P_\alpha(B) \langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle \\ &= \frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{P}_{P \sim \mathcal{P}} \{P(A) = P(B)\} \langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle \end{aligned} \quad (8.11)$$

Let us also compute the inner product of $\mathbf{W}_i^{A,x}$ and $\mathbf{W}_j^{B,y}$. Recall the notation $\langle u, v \rangle_\psi := \langle \psi(u), \psi(v) \rangle$.

$$\begin{aligned} \langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle &= \frac{1}{n} \sum_{a \in A, b \in B} \mathbb{E}_{r \sim \mathcal{R}} \langle x_a - i + a^{\otimes t}(r), y_b - j + b^{\otimes t}(r) \rangle_\psi \\ &= \frac{1}{n} \sum_{a \in A, b \in B} \langle (x_a - i) \otimes a^{\otimes t}, (y_b - j) \otimes b^{\otimes t} \rangle_\psi \quad (\text{by Observation 8.34}) \\ &= \frac{1}{n} \sum_{a \in A, b \in B} \langle \psi(x_a - i), \psi(y_b - j) \rangle \langle a, b \rangle_\psi^t \quad (\text{by Lemma 8.32}) \end{aligned} \quad (8.12)$$

8.8.2. Local Distributions

Fix a subset $\mathcal{S} \subset \mathcal{B}$ of size at most r . In this section, we will construct a local distribution over \mathbb{F}_q -assignments for the vertex set $S = \mathcal{S} \times \mathbb{F}_q^R$ (see Figure 8.2). Clearly, the same construction also yields a distribution for a general set of vertices $S' \subset \mathcal{B} \times \mathbb{F}_q^R$ of size at most r .

Remark 8.46. In the construction in Figure 8.2, the steps 6–7 are not strictly necessary, but they simplify some of the following calculations. Specifically, we could use the \mathbb{F}_q -assignment $\{F^{B,x}\}_{(B,x) \in \mathcal{S}}$ to define the local distribution for the vertex set S . The resulting collection of local distributions could be extended to an approximately feasible SDP solution (albeit using a slightly different vector solution).

We need the following two simple observations.

Observation 8.47. For all $a, b \in \mathbb{F}_q$, we have

$$\mathbb{P}_{\kappa \in \mathbb{F}_q} [a + \kappa = i \wedge b + \kappa = j] = \frac{1}{q^2} + \frac{q-1}{q^2} \langle \psi(a-i), \psi(b-j) \rangle.$$

Proof. If $a-i = b-j$ then both LHS and RHS are equal to $1/q$, otherwise both are equal to 0. \square

Observation 8.48. Fix $a, b \in \mathbb{F}_q$, over a random choice of $h_1, h_2 \in \mathbb{F}_q$,

$$\mathbb{E}_{h_1, h_2 \in \mathbb{F}_q} [\langle \psi(a+h_1), \psi(b+h_2) \rangle] = 0.$$

Proof. Follows easily from the fact that $\langle \psi(i), \psi(j) \rangle = 1$ if $i = j$ and $-1/q-1$ otherwise. \square

The next lemma shows that the second-order correlations of the distribution μ_S approximately match the inner products of the vector solution $\{\mathbf{V}_i^{A,x}\}$.

Lemma 8.49. For any two vertices $(A, x), (B, y) \in S$,

$$\mathbb{P}_{Z \sim \mu_S} [Z^{A,x} = i \wedge Z^{B,y} = j] = \langle \mathbf{V}_i^{A,x}, \mathbf{V}_j^{B,y} \rangle \pm 10|S|^2(3/4)^{t/2}.$$

Proof. Firstly, since $\mathbb{P}[\text{Cons}(S, P)] \geq 1 - |S|^2(3/4)^t$ (by Corollary 8.44),

$$\mathbb{P}_{\mu_S} [Z^{A,x} = i \wedge Z^{B,y} = j] = \mathbb{P}_{\mu_S} [Z^{A,x} = i \wedge Z^{B,y} = j \mid \text{Cons}(S, P)] \pm |S|^2(3/4)^t. \quad (8.13)$$

8. Limits of Semidefinite Relaxations

For $S = \mathcal{S} \times \mathbb{F}_q^R$, the local distribution μ_S over assignments \mathbb{F}_q^S is defined by the following sampling procedure:

Partitioning:

1. Sample a partition $P = \{P_\alpha\}_{\alpha=1}^T$ of \mathcal{B} from the distribution \mathcal{P} obtained by Lemma 8.42. Let α_A, α_B denote the indices of sets in the partition P that contain $A, B \in \mathcal{S}$ respectively.
2. If $\text{Incons}(\mathcal{S}, P)$ then output a uniform random \mathbb{F}_q -assignment to $S = \mathcal{S} \times \mathbb{F}_q^R$. Specifically, set

$$Z^{(B,x)} = \text{uniform random element from } \mathbb{F}_q \quad \forall B \in \mathcal{S}, x \in \mathbb{F}_q^R.$$

Choosing Consistent Representatives:

4. If $\text{Cons}(\mathcal{S}, P)$ then by Corollary 8.41, for every part $\mathcal{S}_\alpha = P_\alpha \cap \mathcal{S}$, there exists bijections $\Pi_{\mathcal{S}_\alpha} = \{\pi_B: [R] \rightarrow B \mid B \in \mathcal{S}_\alpha\}$ such that for every $A, B \in \mathcal{S}_\alpha$,

$$\pi_A = \pi_{A \leftarrow B} \circ \pi_B.$$

5. Sample $L = \{\ell_\alpha\}_{\alpha=1}^T$ by choosing each ℓ_α uniformly at random from $[R]$. For every cloud $B \in \mathcal{S}$, define $\ell_B = \ell_{\alpha_B}$. The choice of L determines a set of representatives for each $B \in \mathcal{S}$. Specifically, the representative of B is fixed to be $\pi_B(\ell_B)$.

Sampling Assignments:

5. Sample $r \in \mathcal{R}$ from the corresponding probability measure and assign

$$F^{B,x}(P, L, r) = x_{\pi_B(\ell_B)} + \pi_B(\ell_B)^{\otimes t}(r).$$

6. Sample $H = \{h_\alpha\}_{\alpha=1}^T$ by choosing each h_α uniformly at random from $[q]$. For every cloud $B \in \mathcal{B}$, define $h_B = h_{\alpha_B}$.

7. Sample κ uniformly at random from $[q]$.

8. For each $B \in \mathcal{S}_\alpha$ and $x \in \mathbb{F}_q^R$, set

$$Z^{B,x}(P, L, r, H, \kappa) = F^{B,x}(P, L, r) + h_B + \kappa.$$

9. Output the \mathbb{F}_q -assignment $\{Z^{B,x}\}_{(B,x) \in \mathcal{S}}$.

Figure 8.2.: Local distribution over \mathbb{F}_q -assignments

8. Limits of Semidefinite Relaxations

Using Observation 8.47, and the definition of $Z^{A,x}$ and $Z^{B,y}$ we can write

$$\begin{aligned} & \mathbb{P}_{\mu_S} \left[Z^{A,x} = i \wedge Z^{B,y} = j \mid \text{Cons}(\mathcal{S}, P) \right] \\ &= \frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{E}_{P,H,L,r} \left[\langle \psi(F^{A,x} + h_A - i), \psi(F^{B,y} + h_B - j) \rangle \mid \text{Cons}(\mathcal{S}, P) \right]. \end{aligned} \quad (8.14)$$

If A, B fall in the same set in the partition P (that is $\alpha_A = \alpha_B$), then we have $h_A = h_B$. If A, B fall in different sets (that is $\alpha_A \neq \alpha_B$), then h_A, h_B are independent random variables uniformly distributed over \mathbb{F}_q . Using Observation 8.48, we can write

$$\begin{aligned} & \mathbb{E}_{P,H,L,r} \left[\langle \psi(F^{A,x} + h_A - i), \psi(F^{B,y} + h_B - j) \rangle \mid \text{Cons}(\mathcal{S}, P) \right] \\ &= \mathbb{E}_{P,L,r} \left[\mathbf{1}(\alpha_A = \alpha_B) \langle \psi(F^{A,x} - i), \psi(F^{B,y} - j) \rangle \mid \text{Cons}(\mathcal{S}, P) \right]. \end{aligned} \quad (8.15)$$

Let P be a partition such that $\text{Cons}(\mathcal{S}, P)$ and $\alpha_A = \alpha_B = \alpha$. The bijections π_A, π_B (see step 4 Figure 8.2) satisfy $\pi_A = \pi_{A \leftarrow B} \circ \pi_B$. Note that therefore $a = \pi_{A \leftarrow B}(b)$ whenever $a = \pi_A(\ell)$ and $b = \pi_B(\ell)$ for some $\ell \in [R]$. Hence,

$$\begin{aligned} & \mathbb{E}_L \mathbb{E}_r \left[\langle \psi(F^{A,x}(P, L, r) - i), \psi(F^{B,y}(P, L, r) - j) \rangle \right] \\ &= \mathbb{E}_{\ell_\alpha} \mathbb{E}_r \left[\langle \psi(x_{\pi_A(\ell_\alpha)} - i + \pi_A(\ell_\alpha)^{\otimes t}(r), \psi(y_{\pi_B(\ell_\alpha)} - j + \pi_B(\ell_\alpha)^{\otimes t}(r)) \rangle \right] \\ &= \frac{1}{R} \sum_{\substack{a \in A, b \in B \\ a = \pi_{A \leftarrow B}(b)}} \mathbb{E}_r \langle \psi(x_a - i + a^{\otimes t}(r), \psi(y_b - j + b^{\otimes t}(r)) \rangle \quad (\text{using } \pi_A = \pi_{A \leftarrow B} \circ \pi_B) \\ &= \frac{1}{R} \sum_{\substack{a \in A, b \in B \\ a = \pi_{A \leftarrow B}(b)}} \langle \psi(x_a - i), \psi(y_b - j) \rangle \cdot \langle a, b \rangle_\psi^t \quad (\text{using Observation 8.34 and Lemma 8.32}) \\ &= \langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle \pm 2 \cdot (3/4)^{t/2} \quad (\text{using Eq. (8.12) and Lemma 8.20}). \end{aligned}$$

Combining the last equation with the previous equations (8.13)–(8.15), we

8. Limits of Semidefinite Relaxations

can finish the proof

$$\begin{aligned}
& \mathbb{P}_{\mu_S} \left[Z^{A,x} = i \wedge Z^{B,y} = j \right] \\
&= \frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{E}_P \left[\mathbf{1}(\alpha_A = \alpha_B) \mid \text{Cons}(\mathcal{S}, P) \right] \langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle \pm (|\mathcal{S}|^2 (3/4)^t + 2 \cdot (3/4)^{t/2}) \\
&= \frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{P} \left[P(A) = P(B) \right] \cdot \langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle \pm 10 |\mathcal{S}|^2 (3/4)^{t/2} \\
&\quad \text{(using } \mathbb{P} \{ \text{Cons}(\mathcal{S}, P) \} \geq 1 - |\mathcal{S}|^2 (3/4)^t \text{ and } |\langle \mathbf{W}_i^{A,x}, \mathbf{W}_j^{B,y} \rangle| \leq 1) \\
&= \langle \mathbf{V}_i^{A,x}, \mathbf{V}_j^{B,y} \rangle \pm 10 |\mathcal{S}|^2 (3/4)^{t/2} \quad \text{(using Eq. (8.11)).}
\end{aligned}$$

□

Lemma 8.50. *Let $\mathcal{S}' \subset \mathcal{S}$ be two subsets of \mathcal{B} and let $S' = \mathcal{S}' \times \mathbb{F}_q^R$ and $S = \mathcal{S} \times \mathbb{F}_q^R$. Then,*

$$\| \mu_{S'} - \text{margin}_{S'} S \|_1 \leq 2 |\mathcal{S}|^2 (3/4)^t.$$

Proof. For a partition $P \in \mathcal{P}$, let $\mu_{S|P}$ denote the distribution μ_S conditioned on the choice of partition P . Firstly, we will show the following claim:

Claim 8.51. *If $\text{Cons}(\mathcal{S}', P)$ and $\text{Cons}(\mathcal{S}, P)$, then $\mu_{S'|P} = \text{margin}_{S'} \mu_{S|P}$.*

Proof. Let $\{\mathcal{S}_\alpha\}$ and $\{\mathcal{S}'_\alpha\}$ denote the partitions induced by P on the sets \mathcal{S} and \mathcal{S}' respectively. Since $\mathcal{S}' \subseteq \mathcal{S}$, we have $\mathcal{S}'_\alpha \subseteq \mathcal{S}_\alpha$ for all $\alpha \in [T]$. By our assumption, each of the sets \mathcal{S}'_α are *consistent* in that $\rho(A, B) \geq 1 - 1/16$ for all $A, B \in \mathcal{S}'_\alpha$. Similarly, the sets \mathcal{S}_α are also *consistent*.

Let us consider the pair of sets $\mathcal{S}'_\alpha \subset \mathcal{S}_\alpha$ for some $\alpha \in [T]$. Intuitively, the vectors within these sets fall into R distinct clusters. Thus the distribution over the choice of consistent representatives are the same in $\mu_{S'|P}$ and $\text{margin}_{S'} \mu_{S|P}$. Formally, we have two sets of bijections $\Pi_{\mathcal{S}'_\alpha} = \{\pi'_A \mid A \in \mathcal{S}'_\alpha\}$ and $\Pi_{\mathcal{S}_\alpha} = \{\pi_A \mid A \in \mathcal{S}_\alpha\}$ satisfying the following property:

$$\pi_{A \rightarrow B} \circ \pi'_A(\ell) = \pi'_B(\ell) \quad \pi_{A \rightarrow B} \circ \pi_A(\ell) = \pi_B(\ell) \quad \forall A, B \in \mathcal{S}'_\alpha, \ell \in [R].$$

Fix a collection $A \in \mathcal{S}'_\alpha$. Let \sim denote that two sets of random variables are identically distributed.

$$\begin{aligned}
\{\pi'_B(\ell_\alpha) \mid B \in \mathcal{S}'_\alpha\} &\sim \{\pi_{A \rightarrow B} \circ \pi'_A(\ell_\alpha) \mid B \in \mathcal{S}'_\alpha\} \\
&\sim \{\pi_{A \rightarrow B}(a) \mid B \in \mathcal{S}'_\alpha, a \text{ is uniformly random in } A\} \\
&\sim \{\pi_{A \rightarrow B} \circ \pi_A(\ell_\alpha) \mid B \in \mathcal{S}'_\alpha\} \sim \{\pi_B(\ell_\alpha) \mid B \in \mathcal{S}'_\alpha\}.
\end{aligned}$$

8. Limits of Semidefinite Relaxations

The variables $L = \{\ell_\alpha\}$ are independent of each other. Therefore,

$$\{\pi'_B(\ell_B) \mid B \in \mathcal{S}'\} \sim \{\pi_B(\ell_B) \mid B \in \mathcal{S}'\}.$$

Notice that the choice of $r \in \mathcal{R}$, H and κ are independent of the set \mathcal{S} . Hence, the final assignments $\{Z^{B,x} \mid B \in \mathcal{S}', x \in \mathbb{F}_q^R\}$ are identically distributed in both cases. \square

Returning to the proof of Lemma 8.50, we can write

$$\begin{aligned} \|\mu_{\mathcal{S}'} - \text{margin}_{\mathcal{S}'} \mu_{\mathcal{S}}\|_1 &= \left\| \mathbb{E}_P \mu_{\mathcal{S}'|P} - \mathbb{E}_P \text{margin}_{\mathcal{S}'} \mu_{\mathcal{S}|P} \right\|_1 \\ &\leq \mathbb{E}_P \left[\|\mu_{\mathcal{S}'|P} - \text{margin}_{\mathcal{S}'} \mu_{\mathcal{S}|P}\|_1 \right] \quad (\text{using Jensen's inequality}) \\ &= \mathbb{P}[\text{Incons}(\mathcal{S}, P)] \cdot \mathbb{E}_P \left[\|\mu_{\mathcal{S}'|P} - \text{margin}_{\mathcal{S}'} \mu_{\mathcal{S}|P}\|_1 \mid \text{Incons}(\mathcal{S}, P) \right]. \end{aligned}$$

The first step uses that the operator $\text{margin}_{\mathcal{S}'}$ is linear. The final step in the above calculation makes use of Claim 8.51. The lemma follows by observing that $\mathbb{P}[\text{Incons}(\mathcal{S}, P)] \leq |\mathcal{S}|^2 (3/4)^t$ and $\|\mu_{\mathcal{S}'|P} - \text{margin}_{\mathcal{S}'} \mu_{\mathcal{S}|P}\|_1 \leq 2$. \square

The next corollary follows from the previous lemma (Lemma 8.50) and the triangle inequality.

Corollary 8.52. *Let $\mathcal{S}, \mathcal{S}'$ be two subsets of \mathcal{B} and let $\mathcal{S}' = \mathcal{S}' \times \mathbb{F}_q^R$ and $\mathcal{S} = \mathcal{S} \times \mathbb{F}_q^R$. Then,*

$$\|\text{margin}_{\mathcal{S} \cap \mathcal{S}'} \mu_{\mathcal{S}} - \text{margin}_{\mathcal{S} \cap \mathcal{S}'} \mu_{\mathcal{S}'}\|_1 \leq 4 \max(|\mathcal{S}|^2, |\mathcal{S}'|^2) (3/4)^t.$$

Proof. Suppose Φ is given by the vertex set V , the edge set $E \subseteq V \times V$, and the collection of permutations $\{\pi_e\}_{e \in E}$. Using Lemma 8.16, we obtain a weak \mathbb{F}_q -integral SDP solution $\mathcal{B} = \{B_u\}_{u \in V}$ of value $1 - O(\sqrt{\eta \log q})$ for Φ .

We construct a vector solution $\{\mathbf{V}_i^{B_u, x} \mid i \in \mathbb{F}_q, B_u \in \mathcal{B}, x \in \mathbb{F}_q^R\}$ and local distributions $\{\mu_S \mid S \subseteq \mathcal{B} \times \mathbb{F}_q^R\}$ as defined in Section (Section 8.8).

Note that since each set $B \in \mathcal{B}$ correspond to a vertices in $u \in V$, we can view these vectors and local distributions as an SDP solution for the E2Lin_q instance $\Phi_{\gamma, q}(\Phi)$. Specifically, we make the identifications $\mathbf{V}_i^{u, x} := \mathbf{V}_i^{B_u, x}$ and $\mu_S := \mu_{\{(B_u, x) \mid (u, x) \in S\}}$ for all $u \in V$, $x \in \mathbb{F}_q^R$, and sets $S \subseteq V \times \mathbb{F}_q^R$.

8. Limits of Semidefinite Relaxations

Lemma 8.49 and Corollary 8.52 show that this SDP solution is ε -infeasible for SA_r and LH_r , where $\varepsilon = O(r^2 \cdot (3/4)^{t/2})$. The value of the SDP solution for $\Phi_{\gamma,q}(\Phi)$ (see Fig. 8.1) is given by

$$\mathbb{E}_{v \in V} \mathbb{E}_{\substack{w, w' \in N(v) \\ \pi = \pi_{w,v}, \pi' = \pi_{w',v}}} \mathbb{E}_{\{x,y\}} \mathbb{E}_{c \in \mathbb{F}_q} \sum_{i=1}^q \langle \mathbf{V}_i^{w, (x \circ \pi + c \cdot \mathbb{1})}, \mathbf{V}_{i-c}^{w', y \circ \pi'} \rangle.$$

Using Eq. (8.11)–(8.12),

$$\begin{aligned} & \langle \mathbf{V}_i^{w, (x \circ \pi + c \cdot \mathbb{1})}, \mathbf{V}_{i-c}^{w', y \circ \pi'} \rangle \\ &= \frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{P}_{P \sim \mathcal{P}} [P(B_w) = P(B_{w'})] \cdot \frac{1}{n} \sum_{\ell, \ell' \in [R]} \langle \psi(x_{\pi(\ell)} + c - i), \psi(y_{\pi'(\ell')} - (i - c)) \rangle \langle b_{w,\ell}, b_{w',\ell'} \rangle_\psi^t. \end{aligned}$$

Note that $\langle \psi(x_{\pi(\ell)} + c - i), \psi(y_{\pi'(\ell')} - (i - c)) \rangle = \langle \psi(x_{\pi(\ell)}), \psi(y_{\pi'(\ell')}) \rangle$. Using Observation 8.14, we have $\pi_{(w,v)}(\ell) = \pi_{(w',v)}(\ell')$ if and only if $\ell = \pi_{B_w \leftarrow B_{w'}}(\ell')$. Hence, by Lemma 8.20,

$$\begin{aligned} & \frac{1}{n} \sum_{\ell, \ell' \in [R]} \langle \psi(x_{\pi(\ell)}), \psi(y_{\pi'(\ell')}) \rangle \langle b_{w,\ell}, b_{w',\ell'} \rangle_\psi^t \\ &= \frac{1}{n} \sum_{\ell} \langle \psi(x_{\pi(\ell)}), \psi(y_{\pi(\ell)}) \rangle \langle b_{w,\pi(\ell)}, b_{w',\pi(\ell)} \rangle_\psi^t \pm 2 \cdot r^2 (3/4)^{t/2} \\ &= \frac{1}{n} \sum_{\ell} \langle \psi(x_\ell), \psi(y_\ell) \rangle \rho(B_w, B_{w'})^t \pm O(\varepsilon). \end{aligned}$$

Note that the distribution of $\{x, y\}$ is independent of the vertices v, w, w' , and

$$\mathbb{E}_{\{x,y\}} \frac{1}{R} \sum_{\ell \in [R]} \langle \psi(x_\ell), \psi(y_\ell) \rangle = 1 - \gamma.$$

Therefore, if we let $\eta_{w,w'} = \rho(B_w, B_{w'})$, we can lower bound the value of the SDP solution as follows

$$\begin{aligned} & \mathbb{E}_{v \in V} \mathbb{E}_{\substack{w, w' \in N(v) \\ \pi = \pi_{w,v}, \pi' = \pi_{w',v}}} \mathbb{E}_{\{x,y\}} \mathbb{E}_{c \in \mathbb{F}_q} \sum_{i=1}^q \langle \mathbf{V}_i^{w, (x \circ \pi + c \cdot \mathbb{1})}, \mathbf{V}_{i-c}^{w', y \circ \pi'} \rangle \\ &= \mathbb{E}_{v \in V} \mathbb{E}_{w, w' \in N(v)} \left[\frac{1}{q^2} + \frac{q-1}{q^2} \mathbb{P}_{P \sim \mathcal{P}} [P(B_w) = P(B_{w'})] \cdot q \cdot \rho(B_w, B_{w'})^t (1 - \gamma) \right] \pm O(\varepsilon) \\ &\geq (1 - \gamma) \mathbb{E}_{v \in V} \mathbb{E}_{w, w' \in N(v)} \mathbb{P}_{P \sim \mathcal{P}} [P(B_w) = P(B_{w'})] \rho(B_w, B_{w'})^t \pm O(\varepsilon) \\ &\geq (1 - \gamma) \mathbb{E}_{v \in V} \mathbb{E}_{w, w' \in N(v)} (1 - O(t \sqrt{\eta_{w,w'}})) \pm O(\varepsilon) \quad (\text{using Lemma 8.42}) \end{aligned}$$

8. Limits of Semidefinite Relaxations

Using Jensen's inequality and the fact that $\mathbb{E}_{v,w,w'} \eta_{v,w,w'} = O(\sqrt{\eta \log q})$ (Lemma 8.16), we see that the value of our SDP solution is at least $1 - \gamma - O(\varepsilon + t\eta^{1/4})$ (recall that we assume q to be constant).

On smoothing the SDP solution using Theorem 8.10, we lose $O(r^2\varepsilon) = O(r^4(3/4)^t)$ in the SDP value. Thus we can set $t = o(\eta^{-1/4})$ and $r = (3/4)^{t/10}$ in order to get a feasible SDP solution for LH_r with value $1 - \gamma - o_{\eta,\delta}(1)$.

On smoothing the SDP solution using Theorem 8.11, we lose $O(q^r\varepsilon) = O(q^r(3/4)^t)$ in the SDP value. Thus we can set, $t = o(\eta^{-1/4})$ and $r = t/\log^2 q$, we would get a feasible SDP solution for SA_r with value $1 - \gamma - o_{\eta,\delta}(1)$. \square

Proof of Theorems 8.1–8.2.. Using Theorem 8.23 with the Khot–Vishnoi integrality gap instance (Lemma 8.15), we have $N = 2^{2^{\log(1/\delta)/\eta}}$ and thus $r = 2^{O((\log \log N)^{1/4})}$. Similarly for SA_r , we get $r = O((\log \log N)^{1/4})$.

8.9. Smoothing

Let Σ be a finite alphabet of size q . Let $\{\chi_1, \dots, \chi_q\}$ be an orthonormal basis for the vector space $\{f: \Sigma \rightarrow \mathbb{R}\}$ such that $\chi_1(a) = 1$ for all $a \in \Sigma$. (Here, orthonormal means $\mathbb{E}_{a \in \Sigma} \chi_i(a)\chi_j(a) = \delta_{ij}$ for all $i, j \in [q]$.) For $R \in \mathbb{N}$, let $\{\chi_\sigma \mid \sigma \in [q]^R\}$ be the orthonormal basis of the vector space $\{f: \Sigma^R \rightarrow \mathbb{R}\}$ defined by

$$\chi_\sigma(x) \stackrel{\text{def}}{=} \chi_{\sigma_1}(x_1) \cdots \chi_{\sigma_R}(x_R), \quad (8.16)$$

where $\sigma = (\sigma_1, \dots, \sigma_R) \in [q]^R$ and $x = (x_1, \dots, x_R) \in \Sigma^R$.

For a function $f: \Sigma^R \rightarrow \mathbb{R}$, we denote

$$\hat{f}(\sigma) \stackrel{\text{def}}{=} \sum_{x \in \Sigma^R} f(x) \chi_\sigma(x). \quad (8.17)$$

Using the fact $\mathbb{E}_{\sigma \in [q]^R} \chi_\sigma(x)\chi_\sigma(y) = \delta_{xy}$ for all $x, y \in \Sigma^R$, we see that

$$f = \mathbb{E}_{\sigma \in [q]^R} \hat{f}(\sigma) \chi_\sigma.$$

We define the following norm for functions $\hat{f}: [q]^R \rightarrow \mathbb{R}$,

$$\|\hat{f}\|_1 \stackrel{\text{def}}{=} \sum_{\sigma \in [q]^R} |\hat{f}(\sigma)|.$$

8. Limits of Semidefinite Relaxations

We say $f: \Sigma^R \rightarrow \mathbb{R}$ is a *distribution* if $f(x) \geq 0$ for all $x \in \Sigma^R$ and $\sum_{x \in \Sigma^R} f(x) = 1$. We define

$$K \stackrel{\text{def}}{=} \max_{\sigma \in [q]^R, x \in \Sigma^R} |\chi_\sigma(x)|.$$

In the next lemma, we give a proof of the following intuitive fact: If a function $g: \Sigma^R \rightarrow \mathbb{R}$ satisfies the normalization constraint $\sum_{x \in \Sigma^R} g(x) = 1$ and it is close to a distribution in the sense that there exists a distribution f such that $\|\hat{f} - \hat{g}\|$ is small, then g can be made to a distribution by “smoothing” it. Here, smoothing means to move slightly towards the uniform distribution (where every assignment has probability q^{-R}).

Lemma 8.53. *Let $f, g: \Sigma^R \rightarrow \mathbb{R}$ be two functions with $\hat{f}(\mathbb{1}) = \hat{g}(\mathbb{1}) = 1$. Suppose f is a distribution. Then, the following function is also a distribution*

$$(1 - \varepsilon)g + \varepsilon q^{-R} \quad \text{where } \varepsilon = \|\hat{f} - \hat{g}\|_1 \cdot K.$$

Proof. It is clear that the function $h = (1 - \varepsilon)g + \varepsilon q^{-R}$ satisfies the constraint $\hat{h}(\mathbb{1}) = 1$. For every $x \in \Sigma^R$, we have

$$\begin{aligned} h(x) &= (1 - \varepsilon)g(x) + \varepsilon q^{-R} \\ &\geq (1 - \varepsilon)(g(x) - f(x)) + \varepsilon q^{-R} \quad (\text{using } f(x) \geq 0) \\ &= \varepsilon q^{-R} + (1 - \varepsilon) \mathbb{E}_{\sigma \in [q]^R} (\hat{g}(\sigma) - \hat{f}(\sigma)) \chi_\sigma(x) \\ &\geq \varepsilon q^{-R} - (1 - \varepsilon) \mathbb{E}_{\sigma \in [q]^R} |\hat{g}(\sigma) - \hat{f}(\sigma)| \cdot K \\ &= \varepsilon q^{-R} - (1 - \varepsilon)K \|\hat{f} - \hat{g}\|_1 \cdot q^{-R} \\ &\geq 0. \quad (\text{by our choice of } \varepsilon) \end{aligned}$$

□

Let V be a set. For a function $f: \Sigma^V \rightarrow \mathbb{R}$ and a subset $S \subseteq V$, we define the function $\text{margin}_S f: \Sigma^S \rightarrow \mathbb{R}$ as

$$\text{margin}_S f(x) \stackrel{\text{def}}{=} \sum_{y \in \Sigma^{V \setminus S}} f(x, y).$$

Note that if f is a distribution over Σ -assignments to V then $\text{margin}_S f$ is its marginal distribution over Σ -assignments to T .

8. Limits of Semidefinite Relaxations

Lemma 8.54. For every $f: \Sigma^V \rightarrow \mathbb{R}$ and $S \subseteq V$,

$$\text{margin}_S f = \mathbb{E}_{\sigma \in [q]^S} \hat{f}(\sigma, \mathbf{1}) \chi_\sigma.$$

Here, $\sigma, \mathbf{1}$ denotes the Σ -assignment to V that agrees with σ on S and assigns 1 to all variables in $V \setminus S$.

Proof.

$$\begin{aligned} \text{margin}_S f(x) &= \sum_{y \in \Sigma^{V \setminus S}} f(x, y) \\ &= \sum_{y \in \Sigma^{V \setminus S}} \mathbb{E}_{\sigma \in [q]^V} \hat{f}(\sigma) \chi_\sigma(x, y) \\ &= \sum_{y \in \Sigma^{V \setminus S}} \mathbb{E}_{\sigma \in [q]^S} \mathbb{E}_{\sigma' \in [q]^{V \setminus S}} \hat{f}(\sigma) \chi_\sigma(x) \chi_{\sigma'}(y) \\ &= \mathbb{E}_{\sigma \in [q]^S} \mathbb{E}_{\sigma' \in [q]^{V \setminus S}} \hat{f}(\sigma, \sigma') \chi_\sigma(x) \cdot \sum_{y \in \Sigma^{V \setminus S}} \chi_{\sigma'}(y) \\ &= \mathbb{E}_{\sigma \in [q]^S} \hat{f}(\sigma, \mathbf{1}) \chi_\sigma(x). \quad (\text{using } \sum_{y \in [q]^{V \setminus S}} \chi_{\sigma'}(y) = 0 \text{ for } \sigma' \neq \mathbf{1}.) \end{aligned}$$

□

The margin operator has the following useful property (which is clear from its definition).

Lemma 8.55. For every function $f: \Sigma^V \rightarrow \mathbb{R}$ and any sets $T \subseteq S \subseteq V$,

$$\text{margin}_T \text{margin}_S f = \text{margin}_T f.$$

Lemma 8.56. Let V be a set and let $\{\mu_S: \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ be a collection of distributions. Suppose that for all sets $A, B \subseteq V$ with $|A|, |B| \leq R$,

$$\|\text{margin}_{A \cap B} \mu_A - \text{margin}_{A \cap B} \mu_B\|_1 \leq \eta.$$

Then, there exists a collection of distributions $\{\mu'_S: \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ such that

– for all $A, B \subseteq V$ with $|A|, |B| \leq R$,

$$\text{margin}_{A \cap B} \mu'_A = \text{margin}_{A \cap B} \mu'_B.$$

8. Limits of Semidefinite Relaxations

– for all $S \subseteq V$ with $|S| \leq R$,

$$\|\mu'_S - \mu_S\|_1 \leq O(\eta q^R K^2),$$

The previous lemma is not enough to establish the robustness of our SDP relaxations. The issue is that we not only require that the distributions are consistent among themselves but they should also be consistent with the SDP vectors.

The following lemma allows us to deal with this issue.

Lemma 8.57. *Let V be a set and let $\{\mu_S : \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ be a collection of distributions. Suppose that*

– for all sets $A, B \subseteq V$ with $|A|, |B| \leq R$,

$$\|\text{margin}_{A \cap B} \mu_A - \text{margin}_{A \cap B} \mu_B\|_1 \leq \eta.$$

– for all sets $A, B \subseteq V$ with $|A|, |B| \leq 2$,

$$\text{margin}_{A \cap B} \mu_A = \text{margin}_{A \cap B} \mu_B.$$

Then, for $\varepsilon \geq q^R K^2 \eta$, there exists a collection of distributions $\{\mu'_S : \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ such that

– for all $A, B \subseteq V$ with $|A|, |B| \leq R$,

$$\text{margin}_{A \cap B} \mu'_A = \text{margin}_{A \cap B} \mu'_B. \tag{8.18}$$

– for all $S \subseteq V$ with $|S| \leq R$,

$$\|\mu'_S - \mu_S\|_1 \leq O(K^2 \eta q^R), \tag{8.19}$$

– for all $S \subseteq V$ with $|S| \leq 2$,

$$\mu'_S = (1 - \varepsilon)\mu_S + \varepsilon \cdot q^{-|S|}. \tag{8.20}$$

Proof. For $\sigma \in [q]^V$, let $\text{supp}(\sigma)$ denote the set of coordinates of σ not equal to 1, and let $|\sigma|$ denote the number of such coordinates,

$$\text{supp}(\sigma) \stackrel{\text{def}}{=} \{i \in V \mid \sigma_i \neq 1\} \quad \text{and} \quad |\sigma| \stackrel{\text{def}}{=} |\text{supp}(\sigma)|.$$

8. Limits of Semidefinite Relaxations

For every $\sigma \in [q]^V$ with $|\sigma| \leq R$, we define

$$\hat{f}(\sigma) := \mathbb{E}_{x \sim \mu_S} \chi_\sigma(x) \quad \text{where } S = \text{supp}(\sigma).$$

For every σ with $|\sigma| > R$, we set $\hat{f}(\sigma) := 0$. We define μ'_S in terms of $f = \mathbb{E}_\sigma \hat{f}(\sigma) \chi_\sigma$,

$$\mu'_S := \text{margin}_S(1 - \varepsilon)f + \varepsilon q^{-|V|}.$$

By Lemma 8.55, this choice of μ'_S satisfies condition (8.18).

First, let us argue that the functions μ'_S are distributions. Let $S \subseteq V$ with $|S| \leq R$. For $\sigma \in [q]^S$ with $T := \text{supp}(\sigma) \subseteq S$, we have

$$\begin{aligned} |\hat{f}(\sigma, \mathbf{1}) - \mathbb{E}_{x \sim \mu_S} \chi_\sigma(x)| &= \left| \mathbb{E}_{x \sim \mu_T} \chi_\sigma(x) - \mathbb{E}_{x \sim \mu_S} \chi_\sigma(x) \right| \\ &\leq \|\mu_T - \text{margin}_T \mu_S\|_1 \cdot \max |\chi_\sigma| \\ &\leq \eta \cdot K. \end{aligned} \tag{8.21}$$

Let f_S denote the function $\text{margin}_S f$. By Lemma 8.54, $\hat{f}_S(\sigma) = \hat{f}(\sigma, \mathbf{1})$ for all $\sigma \in [q]^S$. Hence, $\|\hat{f} - \hat{\mu}_S\|_1 \leq q^R \cdot K \eta$. It follows that for $\varepsilon \geq q^R K^2 \eta$, the function $\mu'_S = (1 - \varepsilon)f_S + \varepsilon q^{-|S|}$ is a distribution (using Lemma 8.53).

Next, let us verify that (8.19) holds. We have

$$\begin{aligned} \|\mu'_S - \mu_S\|_1 &\leq O(\varepsilon) + \|\text{margin}_S f - \mu_S\|_1 \\ &\stackrel{\text{La. 8.54}}{=} O(\varepsilon) + \left\| \mathbb{E}_{\sigma \in [k]^S} \left(\hat{f}(\sigma, \mathbf{1}) - \mathbb{E}_{x \sim \mu_S} \chi_\sigma(x) \right) \chi_\sigma \right\|_1 \\ &\stackrel{(8.21)}{\leq} O(\eta K^2 \cdot k^R) \quad (\text{using } |\hat{f}(\sigma, \mathbf{1}) - \hat{\mu}_S(\sigma)| \leq \eta K \text{ and } |\chi_\sigma(x)| \leq K). \end{aligned}$$

Finally, we show that the new distributions satisfy (8.20). Let $S \subseteq V$ be a set of size at most 2. It follows from the consistency assumption that for all $\sigma \in [k]^S$, we have $\hat{f}(\sigma, \mathbf{1}) = \hat{\mu}_S(\sigma)$. Hence, $f_S = \mu_S$, which implies (8.20). \square

Lemma 8.58. *Let V be a set and let $\{\mu_S : \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ be a collection of distributions. Suppose that*

– for all sets $A, B \subseteq V$ with $|A|, |B| \leq R$,

$$\|\text{margin}_{A \cap B} \mu_A - \text{margin}_{A \cap B} \mu_B\|_1 \leq \eta.$$

8. Limits of Semidefinite Relaxations

– for all sets $A, B \subseteq V$ with $|A|, |B| \leq 2$,

$$\text{margin}_{A \cap B} \mu_A = \text{margin}_{A \cap B} \mu_B.$$

Then, for $\varepsilon \geq kR^2K^2\eta$, there exists a collection of distributions $\{\mu'_S : \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ such that

– for all $A, B \subseteq V$ with $|A|, |B| \leq R$ with $|A \cap B| \leq 2$,

$$\text{margin}_{A \cap B} \mu'_A = \text{margin}_{A \cap B} \mu'_B. \quad (8.22)$$

– for all $S \subseteq V$ with $|S| \leq R$,

$$\|\mu'_S - \mu_S\|_1 \leq O(K^2\eta kR^2), \quad (8.23)$$

– for all $S \subseteq V$ with $|S| \leq 2$,

$$\mu'_S = (1 - \varepsilon)\mu_S + \varepsilon \cdot k^{-|S|}. \quad (8.24)$$

Proof. The proof is along the lines of the proof of the previous lemma.

Define $\hat{f} : [k]^R \rightarrow \mathbb{R}$ as before. We define new functions $\{\mu^*_S : \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ such that

$$\hat{\mu}^*_S(\sigma) = \begin{cases} \hat{\mu}_S(\sigma) & \text{if } \text{supp}(\sigma) > 2, \\ \hat{f}(\sigma, \mathbf{1}) & \text{if } 1 \leq \text{supp}(\sigma) \leq 2, \\ 1 & \text{otherwise.} \end{cases}$$

Since $|\hat{f}(\sigma, \mathbf{1}) - \hat{\mu}_S(\sigma)| \leq K\eta$ (see proof of previous lemma), we can upper bound $\|\hat{\mu}^*_S - \hat{\mu}_S\|_1 \leq kR^2 \cdot K\eta$ (there are not more than kR^2 different $\sigma \in [k]^S$ with $\hat{f}(\sigma, \mathbf{1}) \neq \hat{\mu}_S(\sigma)$). By Lemma 8.53, for $\varepsilon \geq kR^2K^2\eta$, the functions $\{\mu'_S : \Sigma^S \rightarrow \mathbb{R} \mid S \subseteq V, |S| \leq R\}$ defined by $\mu'_S := (1 - \varepsilon)\mu^*_S + \varepsilon k^{-|S|}$ are the desired distributions. We can check that the assertions of the lemma are satisfied in the same way as for the proof of the previous lemma. \square

Proofs of Theorem 8.10 and Theorem 8.11 (Sketch). We apply Lemma 8.58 or Lemma 8.57 to the local distributions $\{\mu_S\}$ of the ε -infeasible LH_R or SA_R solution, respectively. We get a new set of local distributions $\{\mu'_S\}$ that have the desired consistency properties. It remains to change the vectors so that their

8. Limits of Semidefinite Relaxations

inner product match the corresponding probabilities in the local distributions. Suppose $\{v_{i,a}\}$ is the original vector assignment. Let $\{u_{i,a}\}$ be the vector assignment that corresponds to the uniform distribution over all possible assignments to the variables (this vector assignment is the geometric center of the set of all vector assignments). Then, we define the new vector assignment $\{v'_{i,a}\}$ as

$$v'_{i,a} = \sqrt{1-\delta} \cdot v_{i,a} \oplus \sqrt{\delta} u_{i,a},$$

where δ is the smoothing parameter in Lemma 8.58 or Lemma 8.57. It is easy to verify that $\{v'_{i,a}\}$ together with $\{\mu'_S\}$ form a feasible LH_R or SA_R solution. \square

8.10. Notes

The material presented in this chapter is based on the paper “Integrality Gaps for Strong SDP Relaxations of Unique Games” [RS09a], joint with Prasad Raghavendra. A preliminary version appeared at FOCS 2009.

Related Work

In a breakthrough result, Arora et al. [ARV04] used a strong semidefinite program with triangle inequalities to obtain $O(\sqrt{\log n})$ approximation for the SPARSEST CUT problem. Inspired by this work, stronger semidefinite programs have been utilized to obtain better approximation algorithms for certain graph coloring problems [Ch107, ACC06, CS08]. We wish to point out that the work of Chlamtac and Singh [CS08] uses the SA_r hierarchy to obtain approximation algorithms for the hypergraph coloring problem.

In this light, hierarchies of stronger SDP relaxations such as Lovász–Schröder [LS91], Lasserre [Las01], and Sherali–Adams hierarchies [SA90] (See [Lau03] for a comparison) have emerged as possible avenues to obtain better approximation ratios.

Considerable progress has been made in understanding the limits of linear programming hierarchies. Building on a sequence of works [ABL02, ABLT06, Tou05, Tou06], Schoenebeck et al. [STT07a] obtained a $2 - \epsilon$ -factor integrality gap for $\Omega(n)$ rounds of Lovász–Schröder LS hierarchy. More recently, Charikar et al. [CMM09] constructed integrality gaps for $\Omega(n^\delta)$ rounds of Sherali–Adams hierarchy for several problems like MAX CUT, VERTEX COVER, SPARSEST CUT and MAXIMUM ACYCLIC SUBGRAPH. Furthermore, the same work also exhibits $\Omega(n^\delta)$ -round Sherali–Adams integrality gap for UNIQUE GAMES, in

8. Limits of Semidefinite Relaxations

turn obtaining a corresponding gap for every problem to which UNIQUE GAMES is reduced to.

Lower bound results of this nature are fewer in the case of semidefinite programs. A $\Omega(n)$ LS_+ round lower bound for proving unsatisfiability of random 3-SAT formulae was obtained in [BOGH⁺06, AAT05]. In turn, this leads to $\Omega(n)$ -round LS_+ integrality gaps for problems like SET COVER, HYPERGRAPH VERTEX COVER where a matching NP-hardness result is known. Similarly, the $\frac{7}{6}$ -integrality gap for $\Omega(n)$ rounds of LS_+ in [STT07b] falls in a regime where a matching NP-hardness result has been shown to hold. A significant exception is the result of Georgiou et al. [GMPT07] that exhibited a $2 - \varepsilon$ -integrality gap for $\Omega\left(\sqrt{\frac{\log n}{\log \log n}}\right)$ rounds of LS_+ hierarchy. More recently, building on the beautiful work of [Sch08] on Lasserre integrality gaps for Random 3-SAT, Tulsiani [Tul09] obtained a $\Omega(n)$ -round Lasserre integrality gap matching the corresponding UG-hardness for k -CSP [ST06].

9. Open Problems

Number of Large Eigenvalues versus Small-Set Expansion. Consider the following hypothesis.

Hypothesis 9.1. *There exists an absolute constant $\varepsilon_0 > 0$ such that for every constant $\delta > 0$, there exists a function $f: \mathbb{N} \rightarrow \mathbb{N}$ with $f(n) \leq n^{o(1)}$ such that every graph with n vertices and at least $f(n)$ eigenvalues larger than $1 - \varepsilon_0$ contains a vertex set with volume δ and expansion at most $1 - \varepsilon_0$.*

The results in [Section 5.2](#) (Subspace Enumeration), e.g. [Theorem 5.7](#), demonstrate that this hypothesis implies that the Small-Set Expansion Hypothesis is false (assuming that NP does not have $\exp(n^{o(1)})$ -time algorithms).

The graphs that come closest to refuting [Hypothesis 9.1](#) are Boolean noise graphs (which are important tools for UG-hardness reductions). There exists (Boolean noise) graphs with n vertices and $\text{poly log } n$ eigenvalues larger than $1 - 1/\log \log n$, but no set with volume $O(1/\log n)$ and expansion less than $1/2$. These graphs demonstrate that even if [Hypothesis 9.1](#) were true, the function $f(n)$ in the hypothesis must be larger than $\text{poly log } n$.

We conclude that a resolution of [Hypothesis 9.1](#) would either lead to an algorithm refuting the Small-Set Expansion Hypothesis or it would lead to a family of graphs with “better properties” (in the sense above) than Boolean noise graphs. Such a family of graphs could potentially lead to more efficient hardness reductions or improved integrality gap constructions.

Subexponential Algorithm for Sparsest Cut. An interesting question raised by the subexponential algorithms for SMALL-SET EXPANSION and UNIQUE GAMES (see [Chapter 5](#)) is whether similar subexponential algorithms exist for other problems that are SSE-hard or UG-hard. Concretely, we can ask if for every constant $\varepsilon > 0$, there exists an algorithm for SPARSEST CUT with constant approximation ratio $C = C(\varepsilon)$ and running time $\exp(n^\varepsilon)$. A positive resolution of the following conjecture would imply such an algorithm. (We omit the proof at this point. It follows from extensions of the results presented in [Section 5.2](#), techniques in [[ARV09](#)], and an alternative characterization of the threshold rank of a graph.)

9. Open Problems

Conjecture 9.2. *For every $\varepsilon > 0$, there exists positive constants $\eta = \eta(\varepsilon)$ and $\delta = \delta(\varepsilon)$ such that the following holds: For every collection of unit vectors $v_1, \dots, v_n \in \mathbb{R}^n$ with $\mathbb{E}_{i,j \in [n]} |\langle v_i, v_j \rangle| \leq n^{-\varepsilon}$, there exists two sets $S, T \subseteq \{1, \dots, n\}$ with $|S|, |T| \geq \delta n$ and $\|v_i - v_j\|^2 \geq \eta$ for all $i \in S$ and $j \in T$.*

In words, the conjecture says that if the *global correlation* $\mathbb{E}_{i,j} |\langle v_i, v_j \rangle|$ of a collection of unit vectors $v_1, \dots, v_n \in \mathbb{R}^n$ is polynomially small ($\leq n^{-\varepsilon}$), then an $\Omega_\varepsilon(1)$ fraction of the vectors is $\Omega_\varepsilon(1)$ -separated from another $\Omega_\varepsilon(1)$ fraction of the vectors. (An alternative formulation is that the graph formed by joining any two vectors with distance $\eta = \eta(\varepsilon) > 0$ is not a good vertex expander.)

The configuration of vectors that comes closest to refuting [Conjecture 9.2](#) is the $\Omega(\sqrt{d})$ -fold tensored Boolean hypercube $\{\pm 1/\sqrt{d}\}^d$. Here, one can show that any two subsets containing an $\Omega(1)$ fraction of the vectors are not $\Omega(1)$ separated (using the vertex expansion of the usual graph defined on Boolean hypercube). However, the global correlation of this configuration is at least $2^{-\tilde{O}(\sqrt{\log n})}$, which is not polynomially small.

Resolving [Conjecture 9.2](#) would either lead to an interesting algorithm for SPARSEST CUT or to a family of geometric graphs with “nicer properties” (in the sense above) than the Boolean hypercube.

Bibliography

- [AAT05] Michael Alekhnovich, Sanjeev Arora, and Iannis Tourlakis, *Towards strong nonapproximability results in the Lovász–Schrijver hierarchy*, STOC, 2005, pp. 294–303. [157](#)
- [ABL02] Sanjeev Arora, Béla Bollobás, and László Lovász, *Proving integrality gaps without knowing the linear program*, FOCS, 2002, pp. 313–322. [7](#), [119](#), [156](#)
- [ABLT06] Sanjeev Arora, Béla Bollobás, László Lovász, and Iannis Tourlakis, *Proving integrality gaps without knowing the linear program*, Theory of Computing 2 (2006), no. 1, 19–51. [156](#)
- [ABS10] Sanjeev Arora, Boaz Barak, and David Steurer, *Subexponential algorithms for unique games and related problems*, Manuscript, 2010. [13](#), [78](#)
- [ACC06] Sanjeev Arora, Eden Chlamtac, and Moses Charikar, *New approximation guarantee for chromatic number*, STOC, 2006, pp. 215–224. [156](#)
- [AIMS10] Sanjeev Arora, Russell Impagliazzo, William Matthews, and David Steurer, *Improved algorithms for unique games via divide and conquer*, Tech. Report ECCC TR10-041, ECCC, 2010. [46](#), [80](#)
- [AK09] Noga Alon and Bo’az Klartag, *Economical toric spines via Cheeger’s inequality*, J. Topol. Anal. **1** (2009), no. 2, 101–111. MR MR2541756 [63](#)
- [AKK⁺08] Sanjeev Arora, Subhash Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K. Vishnoi, *Unique games on expanding constraint graphs are easy*, STOC, 2008, pp. 21–28. [13](#), [45](#), [133](#), [140](#), [141](#), [164](#)

Bibliography

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy, *Proof verification and the hardness of approximation problems*, J. ACM **45** (1998), no. 3, 501–555. [3](#)
- [Alo86] Noga Alon, *Eigenvalues and expanders*, Combinatorica **6** (1986), no. 2, 83–96. [9](#)
- [AM85] Noga Alon and V. D. Milman, λ_1 , *isoperimetric inequalities for graphs, and superconcentrators*, J. Comb. Theory, Ser. B **38** (1985), no. 1, 73–88. [9](#)
- [AMS07] Christoph Ambühl, Monaldo Mastrolilli, and Ola Svensson, *Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling*, FOCS, 2007, pp. 329–337. [9](#), [106](#)
- [ARV04] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani, *Expander flows, geometric embeddings and graph partitioning*, STOC, 2004, pp. 222–231. [9](#), [156](#)
- [ARV09] ———, *Expander flows, geometric embeddings and graph partitioning*, J. ACM **56** (2009), no. 2. [82](#), [158](#)
- [AS98] Sanjeev Arora and Shmuel Safra, *Probabilistic checking of proofs: A new characterization of NP*, J. ACM **45** (1998), no. 1, 70–122. [3](#)
- [Aus07] Per Austrin, *Towards sharp inapproximability for any 2-CSP*, FOCS, 2007, pp. 307–317. [120](#), [121](#)
- [BHH⁺08] Boaz Barak, Moritz Hardt, Ishay Haviv, Anup Rao, Oded Regev, and David Steurer, *Rounding parallel repetitions of unique games*, FOCS, 2008, pp. 374–383. [13](#), [62](#)
- [BOGH⁺06] Joshua Buresh-Oppenheimer, Nicola Galesi, Shlomo Hoory, Avner Magen, and Toniann Pitassi, *Rank bounds and integrality gaps for cutting planes procedures*, Theory of Computing **2** (2006), no. 1, 65–90. [157](#)
- [Bor85] Christer Borell, *Geometric bounds on the Ornstein-Uhlenbeck velocity process*, Z. Wahrsch. Verw. Gebiete **70** (1985), no. 1, 1–13. MR 795785 (87k:60103) [31](#)

Bibliography

- [Bro97] Andrei Broder, *On the resemblance and containment of documents*, Compression and Complexity of Sequences, 1997, pp. 21–29. [63](#)
- [Che70] Jeff Cheeger, *A lower bound for the smallest eigenvalue of the Laplacian*, Problems in analysis (Papers dedicated to Salomon Bochner, 1969), Princeton Univ. Press, Princeton, N. J., 1970, pp. 195–199. MR MR0402831 (53 #6645) [8](#), [9](#), [20](#)
- [Chl07] Eden Chlamtac, *Approximation algorithms using hierarchies of semidefinite programming relaxations*, FOCS, 2007, pp. 691–701. [156](#)
- [CKK⁺06] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar, *On the hardness of approximating multicut and sparsest-cut*, Computational Complexity **15** (2006), no. 2, 94–114. [121](#)
- [CMM06a] Moses Charikar, Konstantin Makarychev, and Yury Makarychev, *Near-optimal algorithms for unique games*, STOC, 2006, pp. 205–214. [13](#), [48](#), [139](#)
- [CMM06b] Eden Chlamtac, Konstantin Makarychev, and Yury Makarychev, *How to play unique games using embeddings*, FOCS, 2006, pp. 687–696. [31](#), [63](#)
- [CMM09] Moses Charikar, Konstantin Makarychev, and Yury Makarychev, *Integrality gaps for Sherali–Adams relaxations*, STOC, 2009, pp. 283–292. [119](#), [120](#), [122](#), [156](#)
- [CS08] Eden Chlamtac and Gyanit Singh, *Improved approximation guarantees through higher levels of sdp hierarchies*, APPROX-RANDOM, 2008, pp. 49–62. [156](#)
- [DI98] Tassos Dimitriou and Russell Impagliazzo, *Go with the winners for graph bisection*, SODA, 1998, pp. 510–520. [21](#)
- [DKSV06] Nikhil R. Devanur, Subhash Khot, Rishi Saket, and Nisheeth K. Vishnoi, *Integrality gaps for sparsest cut and minimum linear arrangement problems*, STOC, 2006, pp. 537–546. [121](#)
- [Dod84] Jozef Dodziuk, *Difference equations, isoperimetric inequality and transience of certain random walks*, Trans. Amer. Math. Soc. **284** (1984), no. 2, 787–794. MR 743744 (85m:58185) [9](#)

Bibliography

- [Fei02] Uriel Feige, *Relations between average case complexity and approximation complexity*, IEEE Conference on Computational Complexity, 2002, p. 5. [35](#)
- [FKO07] Uriel Feige, Guy Kindler, and Ryan O’Donnell, *Understanding parallel repetition requires understanding foams*, IEEE Conference on Computational Complexity, 2007, pp. 179–192. [6](#)
- [FL92] Uriel Feige and László Lovász, *Two-prover one-round proof systems: Their power and their problems (extended abstract)*, STOC, 1992, pp. 733–744. [26](#), [48](#)
- [GMPT07] Konstantinos Georgiou, Avner Magen, Toniann Pitassi, and Iannis Tourlakis, *Integrality gaps of $2 - o(1)$ for vertex cover sdps in the Lovász–Schrijver hierarchy*, FOCS, 2007, pp. 702–712. [157](#)
- [GMR08] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra, *Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph*, FOCS, 2008, pp. 573–582. [120](#), [121](#), [122](#)
- [GMT06] Sharad Goel, Ravi Montenegro, and Prasad Tetali, *Mixing time bounds via the spectral profile*, Electron. J. Probab. **11** (2006), no. 1, 1–26 (electronic). MR MR2199053 (2007e:60075) [21](#)
- [GW95] Michel X. Goemans and David P. Williamson, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, J. ACM **42** (1995), no. 6, 1115–1145. [2](#), [136](#)
- [Hås01] Johan Håstad, *Some optimal inapproximability results*, J. ACM **48** (2001), no. 4, 798–859. [3](#)
- [Hol09] Thomas Holenstein, *Parallel repetition: Simplification and the no-signaling case*, Theory of Computing **5** (2009), no. 1, 141–172. [62](#), [63](#)
- [IPZ01] R. Impagliazzo, R. Paturi, and F. Zane, *Which problems have strongly exponential complexity?*, Journal of Computer and System Sciences **63** (2001), no. 4, 512–530. [1](#), [4](#), [64](#), [82](#)

Bibliography

- [Kar72] Richard M. Karp, *Reducibility among combinatorial problems*, Complexity of Computer Computations, 1972, pp. 85–103. [1](#)
- [Kho02] Subhash Khot, *On the power of unique 2-prover 1-round games*, STOC, 2002, pp. 767–775. [4](#), [27](#)
- [Kho10] ———, *On the unique games conjecture (invited survey)*, IEEE Conference on Computational Complexity, 2010, pp. 99–121. [5](#)
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell, *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?*, SIAM J. Comput. **37** (2007), no. 1, 319–357. [27](#), [31](#), [120](#), [121](#), [122](#), [123](#), [124](#), [134](#), [135](#)
- [Kol10] Alexandra Kolla, *Spectral algorithms for unique games*, IEEE Conference on Computational Complexity, 2010, pp. 122–130. [46](#), [78](#), [79](#)
- [KORW08] Guy Kindler, Ryan O’Donnell, Anup Rao, and Avi Wigderson, *Rounding schemes and cubical tilings with sphere-like surface area*, FOCS, 2008. [62](#), [63](#)
- [KRT08] Julia Kempe, Oded Regev, and Ben Toner, *Unique games with entangled provers are easy*, FOCS, 2008, pp. 457–466. [26](#), [63](#)
- [KT02] Jon M. Kleinberg and Éva Tardos, *Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields*, J. ACM **49** (2002), no. 5, 616–639. [63](#)
- [KT07] Alexandra Kolla and Madhur Tulsiani, *Playing random and expanding unique games*, Unpublished manuscript available from the authors’ webpages, to appear in journal version of [AKK⁺08], 2007. [46](#), [78](#), [79](#)
- [KV05] Subhash Khot and Nisheeth K. Vishnoi, *The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1* , FOCS, 2005, pp. 53–62. [45](#), [46](#), [79](#), [119](#), [120](#), [121](#), [122](#), [123](#), [124](#), [126](#), [131](#), [133](#), [134](#)
- [Las01] Jean B. Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM Journal on Optimization **11** (2001), no. 3, 796–817. [35](#), [156](#)

Bibliography

- [Lau03] Monique Laurent, *A comparison of the sherali-adams, lovász-schrijver, and lasserre relaxations for 0-1 programming*, Math. Oper. Res. **28** (2003), no. 3, 470–496. [156](#)
- [LR99] Frank Thomson Leighton and Satish Rao, *Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms*, J. ACM **46** (1999), no. 6, 787–832. [9](#), [82](#)
- [LS91] L. Lovász and A. Schrijver, *Cones of matrices and set-functions and 0-1 optimization*, SIAM J. Optim. **1** (1991), no. 2, 166–190. [156](#)
- [MM09] Konstantin Makarychev and Yury Makarychev, *How to play unique games on expanders*, CoRR **abs/0903.0367** (2009). [45](#), [46](#)
- [MNRS08] Rajsekar Manokaran, Joseph Naor, Prasad Raghavendra, and Roy Schwartz, *SDP gaps and UGC-hardness for multiway cut, 0-extension, and metric labeling*, STOC, 2008, pp. 11–20. [120](#), [121](#), [122](#)
- [MOO05] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz, *Noise stability of functions with low influences invariance and optimality*, FOCS, 2005, pp. 21–30. [31](#)
- [MR08] Dana Moshkovitz and Ran Raz, *Two query pcp with sub-constant error*, FOCS, 2008, pp. 314–323. [3](#), [6](#), [64](#)
- [MS07] Rajat Mittal and Mario Szegedy, *Product rules in semidefinite programming*, FCT, 2007, pp. 435–445. [26](#)
- [Rag08] Prasad Raghavendra, *Optimal algorithms and inapproximability results for every CSP?*, STOC, 2008, pp. 245–254. [5](#), [28](#), [120](#), [122](#)
- [Rao08] Anup Rao, *Parallel repetition in projection games and a concentration bound*, STOC, 2008, pp. 1–10. [82](#)
- [Raz98] Ran Raz, *A parallel repetition theorem*, SIAM J. Comput. **27** (1998), no. 3, 763–803. [3](#), [47](#)
- [Raz08] ———, *A counterexample to strong parallel repetition*, FOCS, 2008, pp. 369–373. [26](#), [62](#), [63](#)

Bibliography

- [RS09a] Prasad Raghavendra and David Steurer, *Integrality gaps for strong SDP relaxations of unique games*, FOCS, 2009, pp. 575–585. [15](#), [156](#)
- [RS09b] ———, *Towards computing the Grothendieck constant*, SODA, 2009, pp. 525–534. [120](#), [121](#)
- [RS10] ———, *Graph expansion and the unique games conjecture*, STOC, 2010, To appear. [14](#), [46](#), [105](#)
- [RST10a] Prasad Raghavendra, David Steurer, and Prasad Tetali, *Approximations for the isoperimetric and spectral profile of graphs and related parameters*, STOC, 2010, To appear. [21](#), [83](#)
- [RST10b] Prasad Raghavendra, David Steurer, and Madhur Tulsiani, *Reductions between expansion problems*, Manuscript, 2010. [14](#), [117](#)
- [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 J. Discrete Math. **3** (1990), no. 3, 411–430. MR MR1061981 (91k:90116) [120](#), [156](#)
- [Sch08] Grant Schoenebeck, *Linear level Lasserre lower bounds for certain k -CSPs*, FOCS, 2008, pp. 593–602. [35](#), [157](#)
- [ST06] Alex Samorodnitsky and Luca Trevisan, *Gowers uniformity, influence of variables, and PCPs*, STOC, 2006, pp. 11–20. [157](#)
- [Ste10a] David Steurer, *Fast SDP algorithms for constraint satisfaction problems*, SODA, 2010, To appear. [28](#)
- [Ste10b] ———, *Improved rounding for parallel repeated unique games*, APPROX-RANDOM, 2010, pp. 724–737. [13](#), [62](#)
- [STT07a] Grant Schoenebeck, Luca Trevisan, and Madhur Tulsiani, *A linear round lower bound for lovasz-schrijver sdp relaxations of vertex cover*, IEEE Conference on Computational Complexity, 2007, pp. 205–216. [122](#), [156](#)
- [STT07b] ———, *Tight integrality gaps for Lovász–Schrijver LP relaxations of vertex cover and max cut*, STOC, 2007, pp. 302–310. [120](#), [157](#)

Bibliography

- [Tou05] Iannis Tourlakis, *Towards optimal integrality gaps for hypergraph vertex cover in the lovász-schrijver hierarchy*, APPROX-RANDOM, 2005, pp. 233–244. [156](#)
- [Tou06] ———, *New lower bounds for vertex cover in the Lovász–Schrijver hierarchy*, IEEE Conference on Computational Complexity, 2006, pp. 170–182. [156](#)
- [Tre05] Luca Trevisan, *Approximation algorithms for unique games*, FOCS, 2005, pp. 197–205. [45](#), [80](#), [140](#)
- [TSSW00] Luca Trevisan, Gregory B. Sorkin, Madhu Sudan, and David P. Williamson, *Gadgets, approximation, and linear programming*, SIAM J. Comput. **29** (2000), no. 6, 2074–2097 (electronic). [3](#)
- [Tul09] Madhur Tulsiani, *CSP gaps and reductions in the Lasserre hierarchy*, STOC, 2009, pp. 303–312. [120](#), [122](#), [157](#)