Solving Unique Games over Globally Hypercontractive Graphs

Mitali Bafna
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

Dor Minzer
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

We study the complexity of affine Unique-Games (UG) over *globally hypercontractive graphs*, which are graphs that are not small set expanders but admit a useful and succinct characterization of all small sets that violate the small-set expansion property. This class of graphs includes the Johnson and Grassmann graphs, which have played a pivotal role in recent PCP constructions for UG, and their generalizations via high-dimensional expanders.

We show new rounding techniques for higher degree sum-of-squares (SoS) relaxations for worst-case optimization. In particular, our algorithm shows how to round "low-entropy" pseudodistributions, broadly extending the algorithmic framework of [5]. At a high level, [5] showed how to round pseudodistributions for problems where there is a "unique" good solution. We extend their framework by exhibiting a rounding for problems where there might be "few good solutions".

Our result suggests that UG is easy on globally hypercontractive graphs, and therefore highlights the importance of graphs that lack such a characterization in the context of PCP reductions for UG.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis; Theory of computation → Complexity theory and logic

Keywords and phrases unique games, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.CCC.2024.3

Related Version *Full Version*: `https://arxiv.org/abs/2304.07284` [10]

Funding *Dor Minzer*: Supported by a Sloan Research Fellowship, NSF CCF award 2227876 and NSF CAREER award 2239160.

1 Introduction

The main goal of this paper is to design efficient algorithms that solve instances of the Unique-Games problem whose underlying graph is globally hypercontractive graphs, an extension of the class of small set expanders. The motivation for our investigation is three-fold.

Candidate hard instances for Unique-Games

Recent progress towards the UGC [31, 21, 22, 32] showed that it is NP-hard to distinguish $1/2$-satisfiable instances of UG from ε-satisfiable instances. These works crucially relied on the use of globally hypercontractive graphs. Our algorithms allow us to examine the hard instances arising from their reduction. We try to identify the source of hardness and thus suggest a natural class of graphs that might be hard for UG (Section 1.2.1).

New rounding techniques for Higher Degree SoS

In doing so we build new rounding techniques for higher degree SoS. The study of algorithms for UG has led to the development of general algorithmic techniques such as sophisticated graph partitioning tools [1] and new rounding techniques for SoS [16, 13]. These techniques have in turn led to breakthroughs in robust statistics [14, 37, 12] and other average case problems [17]. The setting of SoS for worst-case optimization is much less understood though. We give new techniques that might be useful for other problems in the worst case setting. We elaborate on our rounding techniques in Section 1.2.2.

The emergence of Unique-Games instances in other contexts

Unique-Games instances on structured graphs naturally appear in other contexts in theoretical computer science, and the tools developed by trying to design algorithms are often helpful. In the context of the current paper, the type of Unique-Games instance we study turn out to be crucial in the field of high-dimensional expanders. Indeed, in [9, 18] the authors investigate a conjecture due to Dinur and Kaufman [20], which asks whether one can construct sparse, low soundness, 2-query direct product testers using high dimensional expanders. The results of [9, 18] assert that a sufficient condition for a high-dimensional expander to admit such direct product testers is the existence of certain local algorithm to approximate affine instance of Unique-Games defined on graphs associated with the complex; the authors refer to this property as UG coboundary expansion. In a follow-up work in progress, the authors and Lifshitz [8] have constructed complexes that are UG coboundary expanders, and some of the ideas developed herein are crucial. See Section 1.2.3 for more details.

1.1 Unique-Games

The Unique Games Conjecture (UGC in short) is a central open problems in Complexity Theory [27]. In short, the UGC says that distinguishing between almost satisfiable (value $\geq 1-\varepsilon$) and highly unsatisfiable (value $\leq \varepsilon$) instances of a certain 2-variable constraint satisfaction problem (CSP) called *Unique Games* is NP-hard. The primary reason for the interest in UGC is that, if true, it implies a large number of hardness of approximation results that are often times tight [33, 29, 4, 34, 40] (see [28, 43]). One of the most striking consequences of UGC is that it implies that a class of semi-definite programs (SDP), namely the basic SDP or degree 2 Sum-of-Squares, achieves the best possible approximation ratio (among all efficient algorithms) for all CSPs [40].

▶ **Definition 1.** *A instance of Unique-Games Ψ consists of a graph $G=(V,E)$, a finite alphabet Σ and a collection of constraints, $\Phi=\{\Phi_e\}_{e\in E}$, one for each edge in G. For all $e\in E$, the constraint Φ_e takes the form $\Phi_e=\{(\sigma,\phi_e(\sigma)) \mid \sigma\in\Sigma\}$, where $\phi_e\colon\Sigma\to\Sigma$ is a 1-to-1 map.*

The goal in the Unique-Games problem is to find an assignment $A\colon V\to\Sigma$ that satisfies the maximum number of constraints possible, that is, satisfies that $(A(u),A(v))\in\Phi_e$ for the largest number of edges $e=(u,v)\in E$ as possible. We define the value of the instance Ψ by:

$$\mathsf{val}(\Psi)=\max_{A\colon V\to\Sigma}\frac{\#\{e \mid A \text{ satisfies } e\}}{|E|}.$$

With this in mind, the Unique-Games Conjecture is the following statement:

▶ **Conjecture 2.** *For all $\varepsilon,\delta>0$ there is $k\in\mathbb{N}$ such that given a Unique-Games instance Ψ with alphabet size at most k, it is NP-hard to distinguish between:*

YES case: $\mathsf{val}(\Psi)\geq 1-\varepsilon$.

NO case: $\mathsf{val}(\Psi)\leq\delta$.

It turns out that the topology of the underlying graph G plays a crucial role in the complexity of the UG instance defined over it. In particular, it turns out that UG over expander graphs is easy:

▶ **Definition 3.** *Given a regular graph $G = (V, E)$ and a set of vertices $S \subseteq V$, the edge expansion of S is defined by:*

$$\Phi(S) = \Pr_{u\in S, v\in\Gamma(u)}[v \notin S].$$

The results of [3, 38, 2] assert that UG instances with completeness close to 1 over expanders are easy. A graph G is a called a (γ, ξ)-small set-expander (SSE) if for every $S \subseteq V$ of size at most $\xi|V|$ it holds that $\Phi(S) \geq \gamma$. In [5], it is shown that UG is easy over "certifiable" small-set expanders, that in fact captures all currently known small-set expanders. Thus to find hard instances of UG, one must look beyond graphs that are expanders and small set expanders.

1.1.1 NP-hardness Reduction for 2-2 Games and Global Hypercontractivity

Indeed, recent progress towards UGC [31, 21, 22, 32] has utilized graphs which are not small-set expanders. In these works it is proved that 2-to-1-Games are NP-hard (which is a very similar problem to UG, except that each one of the maps ϕ_e defining the constraints is a 2-to-1 map). This implies that for all $\varepsilon > 0$, given a UG instance Ψ over sufficiently large alphabet, it is NP-hard to distinguish between the case that $\mathsf{val}(\Psi) \geq 1/2$ and the case that $\mathsf{val}(\Psi) \leq \varepsilon$. To prove these results, these works use graphs that are not small set expanders in two different ways:

1. Smooth Parallel Repetition: A key step in the reduction of [31, 21, 22, 32] is an application of the Parallel Repetition Theorem [42] to get a hardness result for a sufficiently smooth outer PCP construction. Roughly speaking, this step in the process may be associated with the Johnson graph with a *large intersection* parameter. That is, with the graph $J(n, \ell, t)$ in which the vertices are $\binom{[n]}{\ell}$, and two vertices A and B are adjacent if $|A \cap B| = \ell - t$, and we think of t as much smaller than ℓ (say, $t = \sqrt{\ell}$).
2. Composition with the Grassmann encoding: The Grassmann encoding is an encoding of linear functions based on the Grassmann graph $\mathsf{Grass}(n, \ell)$ over $\mathbb{F}_2$. The Grassmann graph over $\mathbb{F}_2$ is the graph whose vertices are all ℓ-dimensional subspaces of $\mathbb{F}_2^n$, denoted by $\begin{bmatrix} n \\ \ell \end{bmatrix}$, and two vertices L and L' are adjacent if $\mathsf{dim}(L \cap L') = \ell - 1$.

Both of the graphs above, namely the Johnson graph with large intersection sizes, as well as the Grassmann graph, are not small set expanders. However very importantly, the class of small sets in the Grassmann graph with bad expansion has a succinct and intuitive characterization, and the proof of the 2-to-1 Games Theorem heavily relies on this characterization.

Though the term is not formally defined, we refer to graphs such as the Grassmann graph above as globally hypercontractive graphs. By that, we mean that there is a collection of "obviously-non-expanding local sets", such that any small set that doesn't expand well must have a large intersection with one of the sets from the collection (see the full version of the paper for a semi-formal definition). Aside from the Grassmann graph, this class of graphs includes Johnson graphs with small intersection sizes [30], certain Cayley graphs over the symmetric group [23], p-biased cubes for $p = o(1)$, other product domains [26] as well as high dimensional expanders [24, 7].

Thus a natural approach to isolating hard instances of UG is to study the complexity of UG over graphs that are globally hypercontractive, in particular the Johnson and Grassmann graphs. The study of this question was initiated in [5] for the class of *Affine Unique-Games* (Definition 4) over Johnson graphs. Unfortunately their algorithm gave parameters that were insufficient to shed light on the source of hardness in reduction above and therefore to derive any of the consequences of our results (see full version).

1.2 Our Results

Our results improve upon [5] in two main aspects: we are able to deal with instances with arbitrarily small (but constant) completeness, and most importantly, their algorithm gets a soundness guarantee that degrades with other parameters of the graph (which in all PCP constructions grow with the alphabet size), whereas our doesn't. To describe our results we start with the definition of Affine Unique Games.

▶ **Definition 4.** *An instance of Affine-UG is an instance of Unique-Games in which the alphabet is the ring of integers modulo q, $\mathbb{Z}_q$, and all of the constraint maps ϕ_e are affine shifts, that is, ϕ_e of the form $\phi_e(\sigma) = \sigma + b_e$ for some $b_e \in \mathbb{Z}_q$.*

An equivalent but slightly different way to view the Affine-UG problem is as a system of linear equations (X, E) over $\mathbb{Z}_q$. Each equation in E is of the form $x_i - x_j = b$ where $x_i, x_j \in X$ are variables and $b \in \mathbb{Z}_q$ is some constant. Despite looking very restrictive, it is known [29] that the UGC is true if and only if it holds for the class of Affine UG and furthermore this class captures many interesting optimization problems such as Max-Cut and graph coloring, thus we shall focus our attention on Affine UG henceforth. [1]

Our main result asserts that there is a polynomial time algorithm for solving Affine UG over globally hypercontractive graphs. As the term globally hypercontractive graph is not formally defined, below are some concrete instances of graphs on which this applies. In the full version we give a semi-formal definition of globally hypercontractive graphs and also show how our algorithm and analysis can be abstracted to solve UG on such graphs, as long as one is provided with an SoS certificate of global hypercontractivity.

We first consider the Johnson graph with small intersection sizes, which we henceforth refer to as the noisy-Johnson graph. This is the regime in which a characterization theorem for non-expanding sets holds. Formally the "α-noisy" Johnson graph is $J(n, \ell, t)$ in the case that $t = \alpha\ell$, for $\alpha \in (0, 1)$ bounded away from 0 and thought of as a fixed constant independent of ℓ. The first result for Johnson graphs addresses the case that the completeness of the instance is close to 1, in which case our algorithm matches the guarantee of the algorithm of [5] for certifiable small-set expanders, and in particular the α-noisy hypercube graph:

▶ **Theorem 5.** *There is $\varepsilon_0 > 0$ such that for all $\alpha \in (0, 1)$ the following holds for all $0 < \varepsilon \le \varepsilon_0$. There exists an algorithm whose running time is $n^{\mathrm{poly}(\ell, |\Sigma|, 1/\varepsilon)}$ which, on input Ψ which is an affine UG instance over $J(n, \ell, \alpha\ell)$ promised to be at least $(1-\varepsilon)$-satisfiable, finds an assignment that satisfies at least $2^{-O\left(\frac{\sqrt{\varepsilon}}{\alpha}\right)}$-fraction of the constraints in Ψ.*

The second result addresses the case of UG instances with arbitrarily small (but bounded away from 0) completeness, in which case our algorithm satisfies a constant fraction of the constraints:

▶ **Theorem 6.** *For all $\alpha \in (0, 1)$ and $c > 0$, there is $\delta > 0$ such that the following holds. There exists an algorithm whose running time is n^D with $D = \ell^{\mathrm{poly}(|\Sigma|\ell^{1/c})}$[2] which on input Ψ, an affine UG instance over $J(n, \ell, \alpha\ell)$ promised to be at least c-satisfiable, finds an assignment that satisfies at least δ-fraction of the constraints in Ψ.*

---

[1] We remark that the reduction of [29] does not preserve the topology of the graph. We are therefore not able to translate our results directly to the class of general UG, and believe this is an interesting direction for further study.

[2] We note that we have not optimized for D and the $\exp(\ell)$-dependence arises due to the degree of the SoS proofs. We used a blackbox statement to convert some of the proofs to SoS proofs, and we conjecture that one can in fact improve the SoS degree to $O(\ell)$ when done carefully.

We remark that the soundness guarantee in the theorems above does not depend on ℓ (when $\alpha = \Omega(1)$), which in most PCP constructions grows with the alphabet size of the instance. But note that this guarantee degrades when $\alpha = o(1)$ and becomes useless if α depends on the alphabet size. In Section 1.2.1 below we discuss why this is interesting – in fact $o(1)$-Noisy Johnson graphs are a natural candidate for hard instances of UG.

We can get similar results given any of the globally hypercontractive graphs mentioned earlier. Below we give a corollary for the Grassmann graph. We show that there is a polynomial time algorithm solving affine UG over the Grassmann graph, even on instances with small completeness:

▶ **Theorem 7.** *For all $c > 0$ there exists $\delta > 0$ such that the following holds. There exists an algorithm whose running time is n^D with $D = \ell^{\mathrm{poly}(|\Sigma|\ell^{1/c})}$ which on input Ψ, an affine UG instance over $\mathsf{Grass}(n, \ell)$ promised to be at least c-satisfiable, finds an assignment that satisfies at least δ-fraction of the constraints in Ψ.*

Note that since the spectral gap of the Grassmann graph is $1/2$, UG algorithms over expanders already imply Theorem 7 for $c \gg 1/2$. Thus, the main contribution of Theorem 7 is the algorithm on Grassmann graphs that works for arbitrarily small completeness.

Below we state our result for random walks on high dimensional expanders (HDX), a large class of graphs that generalize the Johnson graphs but do not necessarily possess its strong symmetries. These include graphs stemming from cut-offs of [36]'s construction of Ramanujan complexes, or [25]'s construction of coset complex expanders. These graphs exhibit the nice high-dimensional expansion properties (e.g. global hypercontractivity) of the Johnson graphs yet are substantially different in other aspects, such as being of bounded degree.

▶ **Theorem 8.** *For all $\alpha \in (0, 1)$ and $c > 0$, there exists $\delta > 0$ such that the following holds. Let X be any d-dimensional two-sided γ-local-spectral expander with $\gamma \ll o_\ell(1)$ and $d > \ell$. There exists an algorithm whose running time is n^D with $D = \ell^{\mathrm{poly}(|\Sigma|\ell^{1/c})}$ which on input Ψ, an affine UG instance over the canonical walk M on $X(\ell)$ of depth α, promised to be at least c-satisfiable, finds an assignment that satisfies at least δ-fraction of the constraints in Ψ.*

Since we have not defined any of the HDX terminology, let us note that this is indeed a generalization of Theorem 6. The Johnson graph corresponds to the complete complex X (which is the simplest instantiation of a two-sided local spectral expander), and the α-noisy Johnson graph $J(n, \ell, \alpha\ell)$ corresponds to a "canonical" random-walk on $X(\ell)$ that goes down $\alpha\ell$-levels and comes back up randomly to $X(\ell)$ while ensuring that it changes exactly $\alpha\ell$ elements in a vertex. In fact, in the above theorem we can allow M to be any complete random walk on $X(\ell)$ and our soundness guarantee will only depend on c and certain parameters of M that are inherently independent of ℓ [3].

1.2.1 Candidate Hard Instances for Unique Games

Our results suggest that the hardness in the instances of UG obtained via the reduction of [31, 21, 22, 32] does not come from the Grassmann graph (which is globally hypercontractive), but rather from the smooth parallel repetition step. Recall that this step uses a Johnson

[3] Concretely it depends on the stripped threshold rank of M above a certain threshold as defined in [6]. For example, when M is the canonical random walk with depth α on $X(\ell)$, and the completeness is $c = 1 - \varepsilon$, this quantity is $r(M) = O(\sqrt{\varepsilon}/\alpha)$ and our soundness guarantee is $\exp(-r)$, matching that of Theorem 5.

graph with a large intersection parameter ($J(n, \ell, \alpha\ell)$ with $\alpha \approx 0$), that is not globally-hypercontractive. Therefore combining the knowledge from the reduction and our algorithm we get that the α-noisy-Johnson graphs should be hard for UG when $\alpha = o(1)$ and become easy when α is bounded away from 0. This also explains why our soundness guarantee decays with α. Indeed, we would be able to make a stronger assertion provided that our results held for general UG (as opposed to only affine UG) or if the reduction above produced instances of Affine UG. Though we believe an algorithm for general UG should exist along the lines of our algorithm, we do not know how to prove so and leave this is an interesting direction to investigate.

Albeit, ignoring the subtlety between general and affine UG, this means that the $o(1)$-noisy Johnson graphs and shallow random walks on HDX provide a natural candidate for constructing SoS lower bounds for UG.

1.2.2 New Rounding Scheme for Higher Degree SoS

Our algorithms are obtained via a novel rounding scheme and analysis for the standard higher degree Sum-of-Squares SDP relaxation for Unique Games. Raghavendra's [40] groundbreaking result showing the optimality of the basic SDP for all CSPs under the UGC, led to efforts to refute the UGC using higher degree SoS relaxations [35, 39]. The study of SoS algorithms has since produced numerous algorithmic advances across many fronts: high-dimensional robust statistics [14, 37, 11, 12], quantum computation [15] and algorithms for semi-random models [17], to name a few. Most of these works use the sum-of-squares method for *average-case* problems though and unfortunately there remains a dearth of techniques for analysing higher degree SoS relaxations for *worst-case* optimization problems. The handful of techniques known for worst-case rounding are the *global correlation rounding* technique from [16, 41] and its generalization via reweightings in [15].

There is an intuitive reason for why this is the case: all aforementioned algorithms for average-case problems rely on a strong "uniqueness" property for the solution space. That is, given an average-case optimization problem, the key observation in the analysis is that the solution to the problem is *unique* upto small perturbations. These algorithms then proceed by converting a proof of uniqueness into an SoS algorithm for finding such a solution, via the proofs-to-algorithms framework for designing Sum-of-Squares algorithms [13].

Such strong uniqueness properties are too much to expect for worst-case problems. Recently [5] showed how to round UG instances on certifiable SSEs. The key property of such instances was a certain "weak uniqueness" of the solution space: any two solutions to the UG instance on an SSE are weakly correlated to each other, i.e. they "agree" on 1% of the vertices. [5] then exploited this observation to give a novel analysis of a higher degree SoS rounding.

For many worst-case problems though the solution space might not be so structured and in fact could allow for many distinct solutions. It turns out that this is precisely the case for UG on globally hypercontractive graphs. Our main technical contribution is to strengthen and broadly extend the [5] framework. At a high-level we show that in our case, the solution space is supported over "few good solutions". That is, there is a small list of solutions such that every good solution is 1% correlated with one of these. We give a new rounding for higher degree SoS that exploits this "weak few good solutions" property. This turns out to be significantly more challenging than the case where we have "weak uniqueness". We expect that with this strengthening, the framework of weak uniqueness to algorithms should be broadly applicable for other worst-case optimization problems. In Section 2 we provide a detailed overview of our techniques, starting out with the framework of [5].

1.2.3 The Emergence of Unique-Games Instances in Other Contexts

Affine instances of Unique-Games naturally appear in the context of high-dimensional expanders. For instance, given a graph $G = (V, E)$ and a labeling $\Pi\colon E \to \mathbb{F}_2$, one may think of (G, Π) as an instance of Unique-Games, wherein the goal is to find a labeling $A\colon V \to \mathbb{F}_2$ such that $A(u) - A(v) = \Pi(u, v)$ for as many edges $(u, v) \in E$ as possible. Note that if (u, v, w) is a triangle in G and $\Pi(u, v) + \Pi(v, w) + \Pi(w, u) \neq 0$, then no assignment can simultanousely satisfy all of the edges $(u, v), (v, w)$ and (w, u). We call such triangles inconsistent triangles. The coboundary constant of G (with coefficients in $\mathbb{F}_2$) is defined as the ratio

$$\max_{\Pi} \frac{1 - \mathrm{val}(G, \Pi)}{\text{fraction of inconsistent triangles in } G}.$$

The coboundary expansion of a graph (and its higher degree analogs for simplicial complexes) are important notions of topological expansion. These notions are inherently different from the more traditional spectral-type expansion notions studied for graphs (and simplicial complexes), and therefore they provide us additional understanding of graphs/ complexes. For instance, recently the works [9, 18] proved that spectral expansion of simplicial complexes is insufficient if one wishes to construct low soundness direct product testers. Instead, one needs spectral expansion as well as coboundary expansion with respect to some non-Abelian groups. The connection between Unique-Games and expansion is useful in studying these new notions of coboundary expansion, and in a recent work we use it to construct such coboundary expanders [8, 19].

1.3 Open Problems

We end this introductory section by stating a few open directions that are of interest for future research. Perhaps the most pertinent question that arises out of our work is whether one can build better integrality gaps for UG:

▶ **Problem 1.** *Can we get higher degree SoS lower bounds for UG using non-globally hypercontractive graphs such as the Johnson graph in the $o(1)$-noise regime?*

The second problem asks whether our results continue to hold for non-affine unique games:

▶ **Problem 2.** *For globally hypercontractive graphs G such as the Johnson graph (with small intersection size) and the Grassmann graph, is there a polynomial time algorithm that given a UG instance Ψ over G with $\mathsf{val}(\Psi) \geq 1 - \varepsilon$ (where $\varepsilon > 0$ is thought of as small), finds an assignment satisfying at least δ fraction of the constraints in Ψ? How about the case that $\mathsf{val}(\Psi) \geq c$, where c is bounded away from 1?*

The third problem asks whether there are other combinatorial optimization problems for which our techniques may yield improved algorithms. Informally, we show how to round SoS relaxations for problems that admit a few good solutions. We believe that this technique should be useful outside the context of UG – given any problem for which one can prove (in SoS) that there are only a "few good solutions", one can apply similar rounding techniques to obtain one such solution.

▶ **Problem 3.** *Can one use the low-entropy rounding framework to get improved run-time for other combinatorial optimization problems, such as coloring 3-colorable graphs using as few colors as possible or improved subexponential time algorithms for Max-Cut?*

2 Overview of Our Techniques

We now elaborate on our techniques starting with the framework of [5]. They proposed a new technique for rounding relaxations of UG that have "low-entropy" measured via a function called the shift-partition size. Given two fixed assignments for the instance, their shift-partition size is roughly defined as the fraction of variables on which these assignments agree (upto symmetry). Taking the equivalent view of the SDP solution as a distribution $\mathcal{D}$ over non-integral solutions, called a pseudodistribution, the expected shift-partition size of two random assignments drawn from $\mathcal{D}$ is then roughly equal to an average of local collision probabilities under $\mathcal{D}$ and thus a proxy for the *entropy* of $\mathcal{D}$. Their analysis proceeds by showing: (1) when the expected shift-partition size (equivalently collision probability) is large, one can round to a high-valued solution, and moreover (2) when the graph is a certifiable small-set expander, the pseudodistribution always has large shift-partition size! They were not able to extend this idea to get high-valued solutions for the broader class of globally hypercontractive graphs though, since in this case the pseudodistribution might be supported over multiple assignments and therefore does not have high collision probability. It turns out though that even in this harder case, the pseudodistribution $\mathcal{D} \times \mathcal{D}$ has large expected shift-partition size after *conditioning* on an event E. But they could not exploit this property since after conditioning the shift-partition could be large for trivial reasons[4] and therefore is no longer a good proxy for the collision probability/entropy of the distribution $\mathcal{D}$.

To get around this barrier, we show that after a suitable preprocessing step on the pseudodistribution, one can in fact condition on any event E (with not too small probability) while preserving most of the desired local independence properties of the distribution. Thus, even after conditioning on E, the expected shift-partition size of $\mathcal{D}\times\mathcal{D} \mid E$ being large signifies that the pseudodistribution $\mathcal{D}$ has high collision probability. One can then use a simple rounding procedure to obtain a high-valued UG solution. Conditioning pseudodistributions is one of the few ways we know of harnessing the power of higher-degree pseudodistributions, hence we believe that the idea of gaining structural control over the distribution after conditioning may be applicable in the analysis of other SoS algorithms too. To explain further details, we start by describing the approach of [5].

2.1 The Approach of [5]: Rounding analysis via the Shift Partition

Fix an Affine Unique-Games instance $\Psi = (G = (V, E), \mathbb{F}_q, \Phi)$. In the SoS relaxation of the Unique-Games problem we have a collection of variables $X_{v,\sigma}$, one for pair of vertex $v \in V$ and label to it $\sigma \in \Sigma$. The output of the program is a pseudoexpectation operator $\widetilde{\mathbb{E}}$, which assigns to each monomial involving at most d of the variables a real-number, under which:

1. The value is high:

$$\widetilde{\mathbb{E}}\left[\sum_{(u,v)\in E}\sum_{\sigma\in\Sigma} X_{v,\sigma}X_{u,\phi_{u,v}(\sigma)}\right] \geq c \cdot |E|.$$

2. $\widetilde{\mathbb{E}}$ is a linear, positive semi-definite operator (when viewed as a matrix over $\mathbb{R}^{M\times M}$ where M is the set of monomials of degree at most $d/2$) satisfying various Booleanity constraints on $X_{u,\sigma}$.
3. Scaling: $\widetilde{\mathbb{E}}[1] = 1$.

[4] In the worst case, the event E could collapse the product distribution over two random assignments to set the second random assignment to be always equal to the first one. In this case a pair of assignments drawn from $\mathcal{D} \times \mathcal{D} \mid E$ being equal does not say anything about the collision probability of $\mathcal{D}$.

Morally, the pseudoexpectation $\widetilde{\mathbb{E}}$ should be thought of in the following way: there is an unknown distribution $\mathcal{D}$ over assignments $A_1, \ldots, A_m$ that each have value at least c. For the assignment A_i we think of Boolean valued assignment to the variables $X_{u,\sigma}$ that assigns to a variable 1 if and only if $A_i(u) = \sigma$, and associate with it the expectation operator $\mathbb{E}_i$ which maps monomials to Boolean values in the natural way according to A_i. The operator $\widetilde{\mathbb{E}}$ then is the average of the operators $\mathbb{E}_i$ according to $i \sim \mathcal{D}$. [5]

Shift-partition

Given $\widetilde{\mathbb{E}}$, one can construct a different pseudoexpectation operator that allows access to moments of two assignments $X = A_i, X' = A_j$ where $i, j \sim \mathcal{D}$ are chosen independently. In expectation, we get that at least c^2 fraction of the edges get satisfies by both X and X'; the algorithm attempts to satisfy these edges. Towards this end, given two fixed assignments X and X' we define the shift-partition of the vertices of V: $V = \cup_{s\in\mathbb{F}_q} F_s$ where for each $s \in \mathbb{F}_q$ we define

$$F_s(X, X') = \{v \in V \mid X(v) - X'(v) = s\}.$$

The shift-partition size is then defined as:

$$\widetilde{\mathbb{E}}_{X,X'\sim\mathcal{D}}\left[\sum_{s\in\Sigma}\left(\frac{|F_s(X,X')|}{|V(G)|}\right)^2\right].$$

After rearranging, we get that when X and X' are independent, this expression is an average of some local collision probabilities (precisely $\mathbb{E}_{u,v}[CP(X_u - X_v)]$), and hence the shift-partition size being large in expectation turns out to be useful for rounding.

On the other hand, observe that if an edge $(u, v) \in E$ is satisfied by both X and X', then $X(u) - X(v) = X'(u) - X'(v)$ and rearranging we conclude that u and v are in the same part F_s of the shift partition. We therefore conclude that in expectation over $X, X' \sim \mathcal{D}$ at least c^2 fraction of the edges of G stay inside the same part of the shift partition, implying that the expansion of the shift-partition is small.

Small-set expanders

If the graph G is a small-set expander, then the above implies that at least one of the sets F_s/the shift-partition size is large and the following rounding procedure works in such cases:

1. Sample a vertex $v \in V$ and choose $A(v) = \sigma$ according to the distribution $p(\sigma) = \widetilde{\mathbb{E}}[X_{v,\sigma}]$.
2. For any $u \in V$, sample $A(u)$ according to the distribution $p(a) = \frac{\widetilde{\mathbb{E}}[X_{u,a}X_{v,\sigma}]}{\widetilde{\mathbb{E}}[X_{v,\sigma}]}$.

To get an understanding to why this rounding scheme works, think of X as fixed and X' as random. Thus, the fact that part s of the shift partition is large implies that $X' = X + s$ on a constant fraction of the vertices. Therefore, once we sampled the assignment to v in the first part of the algorithm, the value of s is determined. In the second step we are sampling the assignment to other nodes conditioned on the value of v. However, there is one value for u which is much more likely than others – namely $X(u) + s$, and so we can expect that $X'(u) = X(u) + s$ for a constant fraction of the vertices u. In particular, for any edge (u, w) inside F_s that is satisfied by X, we will have that the assignments sampled for u and w are

[5] Formally speaking, when given $\widetilde{\mathbb{E}}$ we are not guaranteed that there exists an actual distribution $\mathcal{D}$ over good assignments as above, however this intuition will be good enough for the sake of this informal presentation.

$X(u)+s$ and $X(w)+s$ respectively with constant probability, in which case we manage to satisfy (u,w). To analyse this rounding strategy formally, [5] crucially use the independence of X and X'.

In essence, the above asserts that the shift-partition being large implies that the solution space of X must have high collision probability, which can then be used for rounding. By that, we mean that our distribution essentially consists of only one assignment (upto shift-symmetry) and its perturbations.

Non small-set expanders

Consider a graph which is not a small set expander, say that G is the Johnson graph $J(n,\ell,t=\ell/2)$. In that case the above reasoning no longer works as F_s may indeed be all small sets. However, as explained earlier, using global hypercontractivity we can infer that one of the sets F_s must posses a certain structure – it must have large density inside one of the canonical non-expanding sets. In the case of the Johnson graph specifically, these canonical sets take the following form:

$$H_R = \left\{ A \in \binom{n}{\ell} \;\middle|\; A \supseteq R \right\},$$

for $R \subset [n]$ and $|R| = r = O(1)$. In fact, global hypercontractivity gives the following stronger structural property: the set $H = \bigcup_{R\in\mathcal{R}} H_R$ where $\mathcal{R}$ consists of all R's inside which some part F_s is dense, has a constant measure. Doing simple accounting, it follows that $|\mathcal{R}| \geq \Omega(n^r/\ell^r)$ and as there are at most $\binom{n}{r}$ different canonical sets it follows that $|\mathcal{R}|$ contains an $\Omega(1/\ell^r)$ fraction of these sets.

For each choice of X and X' though we may have a different collection of dense subcubes $\mathcal{R}$. But since $\mathcal{R}$ contains an $\Omega(1/\ell^r)$ fraction of all the subcubes, we get that there must be at least one subcube H_R that is dense with probability $\Omega(1/\ell^r)$ over $X, X' \sim \mathcal{D}$. Let H_R be such a subcube and $E_R(X,X')$ be the event that H_R is dense. Ideally, at this point one would like to condition on E_R so that one of the parts inside the shift partition F_s becomes large inside H_R, and then hope that as was the case for small-set expanders, we can satisfy many of the edges inside $F_s \cap H_R$.

Unfortunately, this hope does not materialize – after conditioning on E_R even though the shift-partition is large, the rounding strategy above may break. Indeed, for the rounding procedure we wanted the values of $X(u)$ and $X'(u)$ for $X, X' \sim \mathcal{D}$ to be independent for every vertex u. However, after conditioning the joint distribution $\mathcal{D} \times \mathcal{D} \mid E_R$ over (X,X') might have correlations between X and X'. In particular this distribution could even be supported on pairs (X,X') that are always equal to each other, in which case the shift-partition is large because of trivial reasons and therefore its large size doesn't imply anything about the collision probability/entropy of $\mathcal{D}$.

Hence in [5] the authors don't manage to do this conditioning, and instead settle for satisfying an $\Omega\left(\frac{1}{\ell^{2r}}\right)$-fraction of the constraints on H_R. After that they iterate this algorithm many times to satisfy an $\Omega\left(\frac{1}{\ell^{2r}}\right)$-fraction of the constraints of the whole graph.

2.2 Our Approach: Conditioning on the Event E via (Eliminating) Global Correlations

Our main contribution to the above framework is to show that by adding an additional preprocessing step, we can ensure that even after conditioning on the event E_R above, the assignments X and X' will remain highly independent. In particular, the fact that some part in the shift partition becomes large must happen – just like in the case of small-set expanders – due to the fact that our distribution has high collision probability.

As the event $E = E_R(X, X')$ has probability at least $\Omega\left(\frac{1}{\ell^r}\right)$, if we are sufficiently high up in the SoS hierarchy ($\Theta(\ell^r)$ levels will do, for an overall running time of $n^{\Theta(\ell^r)}$), we do have access to the conditional pseudoexpectation

$$\widetilde{\mathbb{E}}[Y \mid E] = \frac{\widetilde{\mathbb{E}}[Y 1_E]}{\widetilde{\mathbb{E}}[1_E]}.$$

This means that we can sample labels of vertices conditioned on the event E. To make this useful though, we must change the rounding procedure. To get some intuition consider the extreme case in which after conditioning on $E(X, X')$ there are huge correlations between X and X' that remain in our distribution.

Namely, suppose that after conditioning on E it holds that $X(u) = X'(u)$ for almost all vertices u. In that case, if we sampled X, X' from $\mathcal{D}\times\mathcal{D}$ (not conditioned on E), we would get that with probability at least $\Pr[E] \geq \Omega\left(\frac{1}{\ell^r}\right)$ the event E holds, in which case X and X' agree on almost all vertices. This means that if $\mathcal{D}$ was an actual distribution the assignments have a large global correlation: fix $X' = X_0$ for X_0 that satisfies $\Pr_{\mathcal{D}}[E(X, X_0) = 1] \geq \Pr_{\mathcal{D}\times\mathcal{D}}[E]$. Once E holds, we have that $X(u) - X(v) = X_0(u) - X_0(v)$ for almost all pairs of vertices, hence the values of the assignment X to the vertices u and v is correlated across $\mathcal{D}$. Therefore, a natural idea is to avoid this issue by transforming $\mathcal{D}$ to another distribution lacking global correlations, in the sense that the assignments to a typical pair of vertices u and v are almost independent.

For this purpose we use an idea from [41], which adapted to our setting says that for any $\tau > 0$ there is $d = d(\tau, |\Sigma|)$ such that conditioning $\widetilde{\mathbb{E}}$ on the values of d randomly chosen vertices ensures that the global correlation is at most τ. That is, the values of $X(u)$ and $X(v)$ for two typical vertices u and v are at most τ-correlated, and the same holds for X'. In the full version of the paper we then show that if we start with such a pseudodistribution that lacks global correlations, then one can condition on the event E and retain near independence between the assignments X and X', at least on most vertices. To be more precise, we show that for $Y_{u,v} = (X(u), X(v))$ and $Y'_{u,v} = (X'(u), X'(v))$, the statistical distance between $Y_{u,v}, Y'_{u,v} \mid E$ and $Y_{u,v}, Y'_{u,v}$ is small for almost all pairs of vertices u, v.[6]

Using this idea we are able to get an $\Omega(1)$-valued solution on some basic set H_R. To summarize, we first preprocess the pseudodistribution to eliminate global correlations. We can then find an event $E(X, X')$, corresponding to the fact that some part F_s in the shift partition has becomes dense in some basic set H_R. Furthermore, conditioning on E most pairs $(X(u), X(v)), (X'(u), X'(v))$ remain almost-independent. Then running a simple rounding procedure on H_R (as in [5]), we are able to satisfy a good fraction of the edges inside H_R. H_R might be a $o(1)$ fraction of the graph though, therefore like [5] we repeat this procedure multiple times to get an $\Omega(1)$-valued solution for the whole graph. This gives an efficient algorithm for affine UG over the Johnson graphs as in Theorem 5.

To prove Theorem 6 (namely, the regime where c is not close to 1) more work is needed. Indeed, in the case that c is close to 1 we are able to conclude that essentially all edges stay within some part F_s of the shift partition. Thus, as long as our sets H_R cover a constant fraction of the edges that stay within some F_s, they are automatically guaranteed to cover a constant fraction of the edges that are satisfied by both X and X', and these are the edges our rounding procedure manages to satisfy. If c is just bounded away from 0 we can no longer make such an argument, and it is no longer even clear that the sets H_R cover some edges that we have a hope of satisfying.

---

[6] To make our rounding succeed we need to use a more complicated version of $Y_{u,v}$ (see the full version of the paper).

2.3 Getting Small Completeness: Capturing all of the Non-expanding Edges

To design our algorithm for the case when the completeness c is just guaranteed to be bounded away from 0 we must first argue that in the shift partition, we are able to capture almost all of the edges that stay within a part F_s using the basic sets H_R (so as to ensure we are including the edges that X and X' both satisfy).

Towards this end we require a more refined corollary of global hypercontractivity, asserting that if we have a small set of vertices F in the Johnson graph that has edge expansion at most $1-\eta$, then we can find a collection $\mathcal{R}$ of basic sets such that:

1. **Bounded and dense:** each $R \in \mathcal{R}$ has size $|R| = O(1)$ and F is dense inside each H_R. That is, $\delta(F \cap H_R) \geq \Omega_\eta(\delta(H_R))$ for each $R \in \mathcal{R}$.
2. **Maximally dense:** For all $R \in \mathcal{R}$ and all $R' \subsetneq R$, F is not very dense in $H_{R'}$.
3. **Capture almost all non-expanding edges:** Almost all the edges that stay inside F also stay inside H_R for some $R \in \mathcal{R}$.

Indeed, we show that a global hypercontractive inequality such as the one in [30] can be used to prove such a result (in a black-box manner).

Using this result, we are able to argue that that the edges that stay inside the subcubes H_R for $R \in \mathcal{R}$ cover most of the edges that stay within the same part in the shift partition. There are several subtleties here that one has to deal with, for example, "regularity issues" such as, how many different R's cover a given edge. The goal of the second item above is to handle such concerns, and it roughly says that no vertex nor edge gets over-counted by a lot. After that, we are able to condition on an event E, where as before E indicates that some part F_s becomes dense inside some basic set H_R, so that the resulting distribution has a large shift-partition inside H_R. At this point, we are (morally) back to the problem of rounding the SoS solution on a set with a large shift-partition, except that now our solution has value $c' > 0$ (as opposed to close to 1). We remark that again, we use the "elimination of global correlations" idea presented earlier to retain near independence after conditioning. With more care, we use a similar analysis to the one presented for completeness close to 1 to finish the proof when c is arbitrarily small.

References

1 Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. *J. ACM*, 62(5):42:1–42:25, 2015. doi:10.1145/2775105.

2 Sanjeev Arora, Russell Impagliazzo, William Matthews, and David Steurer. Improved algorithms for unique games via divide and conquer. *Electron. Colloquium Comput. Complex.*, 17:41, 2010. URL: http://eccc.hpi-web.de/report/2010/041, arXiv:TR10-041.

3 Sanjeev Arora, Subhash Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 21–28, 2008. doi:10.1145/1374376.1374380.

4 Per Austrin. Balanced max 2-sat might not be the hardest. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 189–197. ACM, 2007. doi:10.1145/1250790.1250818.

5 Mitali Bafna, Boaz Barak, Pravesh K. Kothari, Tselil Schramm, and David Steurer. Playing unique games on certified small-set expanders. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 1629–1642. ACM, 2021.

6 Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett. High dimensional expanders: Eigenstripping, pseudorandomness, and unique games. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022*, pages 1069–1128, 2022.

7 Mitali Bafna, Max Hopkins, Tali Kaufman, and Shachar Lovett. Hypercontractivity on high dimensional expanders. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 185–194, 2022.

8 Mitali Bafna, Noam Lifshitz, and Dor Minzer. Constant degree direct product testers with small soundness. *arXiv preprint*, 2024. `arXiv:2402.00850`.

9 Mitali Bafna and Dor Minzer. Characterizing direct product testing via coboundary expansion. *Electron. Colloquium Comput. Complex.*, TR23-120, 2023. `arXiv:TR23-120`.

10 Mitali Bafna and Dor Minzer. Solving unique games over globally hypercontractive graphs. *CoRR*, abs/2304.07284, 2023. `doi:10.48550/arXiv.2304.07284`.

11 Ainesh Bakshi, Ilias Diakonikolas, Samuel B. Hopkins, Daniel Kane, Sushrut Karmalkar, and Pravesh K. Kothari. Outlier-robust clustering of gaussians and other non-spherical mixtures. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 149–159. IEEE, 2020.

12 Ainesh Bakshi, Ilias Diakonikolas, He Jia, Daniel M. Kane, Pravesh K. Kothari, and Santosh S. Vempala. Robustly learning mixtures of k arbitrary gaussians. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 – 24, 2022*, pages 1234–1247. ACM, 2022.

13 Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 – 22, 2012*, pages 307–326, 2012. `doi:10.1145/2213977.2214006`.

14 Boaz Barak, Jonathan A. Kelner, and David Steurer. Rounding sum-of-squares relaxations. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 – June 03, 2014*, pages 31–40, 2014. `doi:10.1145/2591796.2591886`.

15 Boaz Barak, Pravesh K. Kothari, and David Steurer. Quantum entanglement, sum of squares, and the log rank conjecture. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 975–988, 2017. `doi:10.1145/3055399.3055488`.

16 Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 472–481, 2011. `doi:10.1109/FOCS.2011.95`.

17 Rares-Darius Buhai, Pravesh K. Kothari, and David Steurer. Algorithms approaching the threshold for semi-random planted clique. *CoRR*, abs/2212.05619, 2022. `doi:10.48550/arXiv.2212.05619`.

18 Yotam Dikstein and Irit Dinur. Agreement theorems for high dimensional expanders in the small soundness regime: the role of covers. *Electron. Colloquium Comput. Complex.*, TR23-119, 2023. `arXiv:TR23-119`.

19 Yotam Dikstein, Irit Dinur, and Alexander Lubotzky. Low acceptance agreement tests via bounded-degree symplectic hdxs, 2024. `arXiv:2402.01078`.

20 Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 974–985. IEEE Computer Society, 2017. `doi:10.1109/FOCS.2017.94`.

21 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 940–951. ACM, 2018. `doi:10.1145/3188745.3188806`.

22 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 376–389, 2018.

23 Yuval Filmus, Guy Kindler, Noam Lifshitz, and Dor Minzer. Hypercontractivity on the symmetric group. *arXiv preprint*, 2020. arXiv:2009.05503.

24 Tom Gur, Noam Lifshitz, and Siqi Liu. Hypercontractivity on high dimensional expanders. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 176–184, 2022.

25 Tali Kaufman and Izhar Oppenheim. Construction of new local spectral high dimensional expanders. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 773–786, 2018.

26 Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer. Global hypercontractivity and its applications. *arXiv preprint*, 2021. arXiv:2103.04604.

27 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 767–775, 2002. doi:10.1145/509907.510017.

28 Subhash Khot. On the unique games conjecture (invited survey). In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity, CCC 2010, Cambridge, Massachusetts, USA, June 9-12, 2010*, pages 99–121. IEEE Computer Society, 2010. doi:10.1109/CCC.2010.19.

29 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(1):319–357, 2007. doi:10.1137/S0097539705447372.

30 Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra. Small set expansion in the johnson graph. *Electron. Colloquium Comput. Complex.*, TR18-078, 2018. arXiv:TR18-078.

31 Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 576–589, 2017. doi:10.1145/3055399.3055432.

32 Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. *Annals of Mathematics*, 198(1):1–92, 2023.

33 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-epsilon. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008. doi:10.1016/j.jcss.2007.06.019.

34 Guy Kindler, Assaf Naor, and Gideon Schechtman. The UGC hardness threshold of the Lp grothendieck problem. *Math. Oper. Res.*, 35(2):267–283, 2010. doi:10.1287/moor.1090.0425.

35 Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2000/01. doi:10.1137/S1052623400366802.

36 Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Explicit constructions of ramanujan complexes of type ad. *European Journal of Combinatorics*, 26(6):965–993, 2005.

37 Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 438–446. IEEE, 2016.

38 Konstantin Makarychev and Yury Makarychev. How to play unique games on expanders. In Klaus Jansen and Roberto Solis-Oba, editors, *Approximation and Online Algorithms – 8th International Workshop, WAOA 2010, Liverpool, UK, September 9-10, 2010. Revised Papers*, volume 6534 of *Lecture Notes in Computer Science*, pages 190–200. Springer, 2010. doi:10.1007/978-3-642-18318-8_17.

39 Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization.* PhD thesis, California Institute of Technology, 2000.

40 Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 245–254, 2008. doi:10.1145/1374376.1374414.

41 Prasad Raghavendra and Ning Tan. Approximating csps with global cardinality constraints using sdp hierarchies. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 373–387. SIAM, 2012.

42 Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. `doi:10.1137/S0097539795280895`.

43 Luca Trevisan. On khot's unique games conjecture. *Bulletin (New Series) of the American Mathematical Society*, 49(1), 2012.