Unique Games on Expanding Constraint Graphs are Easy

[Extended Abstract]

Sanjeev Arora*
Princeton University
arora@cs.princeton.edu

Subhash A. Khot†
Courant Institute
khot@cs.nyu.edu

Alexandra Kolla ‡
UC Berkeley
akolla@cs.berkeley.edu

David Steurer*
Princeton University
dsteurer@cs.princeton.edu

Madhur Tulsiani§
UC Berkeley
madhurt@cs.berkeley.edu

Nisheeth K. Vishnoi¶
IBM India Research Lab
nisheeth.vishnoi@gmail.com

ABSTRACT

We present an efficient algorithm to find a good solution to the Unique Games problem when the constraint graph is an expander.

We introduce a new analysis of the standard SDP in this case that involves correlations among distant vertices. It also leads to a parallel repetition theorem for unique games when the graph is an expander.

Categories and Subject Descriptors

F.2.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity—*Nonnumerical Algorithms and Problems*

General Terms

Algorithms, Theory

Keywords

Semidefinite Programming, Approximation Algorithms, Expander Graphs

1. INTRODUCTION

*Supported by NSF grants MSPA-MCS 0528414, and ITR 0205594.

†Supported by NSF CAREER award CCF-0643626, Microsoft New Faculty Fellowship and Sloan Foundation Fellowship.

‡Supported by NSF grant ITR CCR-0121555.

§Supported by NSF grant CCF 0515231.

¶Part of this work was done while the author was being financially supported by Umesh Vazirani and being hosted by Sanjeev Arora in Princeton University.

STOC'08, May 17–20, 2008, Victoria, British Columbia, Canada.

UNIQUE GAMES is a constraint satisfaction problem where one is given a constraint graph $G = (V, E)$, a label set $[k]$ and for each edge $e = (u, v)$, a bijective mapping $\pi_{uv} : [k] \mapsto [k]$. The goal is to assign to each vertex in G a label from $[k]$ so as to maximize the fraction of the constraints that are "satisfied," where an edge $e = (u, v)$ is said to be *satisfied* by an assignment if u is assigned a label i and v is assigned a label j such that $\pi_{uv}(i) = j$. The value of a labeling $\Lambda\colon V \to [k]$ is the fraction of the constraints satisfied by it and is denoted by $\mathsf{val}(\Lambda)$. For a UNIQUE GAMES instance $\mathcal{U}$, we denote by $\mathsf{opt}(\mathcal{U})$ the maximum value of $\mathsf{val}(\Lambda)$ over all labelings. This optimization problem was first considered by Cai, Condon, and Lipton [3]. The Unique Games Conjecture (UGC) of Khot [12] asserts that for such a constraint satisfaction problem, for arbitrarily small constants $\eta, \zeta > 0$, it is NP-hard to decide whether there is a labeling that satisfies $1 - \eta$ fraction of the constraints or, for every labeling, the fraction of the constraints satisfied is at most ζ as long as the size of the label set, k, is allowed to grow as a function of η and ζ.

Since its origin, the UGC has been successfully used to prove (often optimal) hardness of approximation results for several important NP-hard problems such as MIN-2SAT-DELETION [12], VERTEX COVER [14], MAXIMUM CUT [13], GRAPH COLORING [8], and non-uniform SPARSEST CUT [5, 15]. However, one fundamental problem that has resisted attempts to prove inapproximability results, even assuming UGC, is the (uniform) SPARSEST CUT problem. This problem has a $O(\sqrt{\log n})$ approximation algorithm by Arora, Rao, and Vazirani [2], but no hardness result beyond NP-hardness is known (recently, in [1], a PTAS is ruled out under a complexity assumption stronger than $\mathsf{P} \neq \mathsf{NP}$). In fact, it seems unlikely that there is a reduction from UNIQUE GAMES to SPARSEST CUT, unless one assumes that the starting UNIQUE GAMES instance has some expansion property. This is because if the UNIQUE GAMES instance itself has a sparse cut, then the instance of SPARSEST CUT produced by such a reduction also has a sparse cut (this is certainly the case for known reductions, i.e. [5, 15]), irrespective of whether the UNIQUE GAMES instance is a YES or a NO instance. This motivates the following question: is UNIQUE GAMES problem hard even with the promise that the constraint graph is an expander? A priori, this could be true even with a very strong notion of expansion (as some of the authors of this paper speculated), leading to a superconstant hardness result for SPARSEST CUT and related problems like MINIMUM LINEAR ARRANGEMENT.

In this paper, we show that the UNIQUE GAMES problem is actually easy when the constraint graph is even a relatively weak expander. One notion of expansion that we consider in this paper is when the second smallest eigenvalue of the normalized Laplacian of a graph G, denoted by $\lambda := \lambda_2(G)$, is bounded away from 0. We note that the size of balanced cuts (relative to the total number of edges) in a graph is also a useful notion of expansion and the results in this paper can be extended to work in that setting.

Our main result.

We show the following theorem in Section 2:

THEOREM 1.1. *There is a polynomial time algorithm for* UNIQUE GAMES *that, given $\eta > 0$, distinguishes between the following two cases:*

- YES *case: There is a labeling which satisfies at least $1 - \eta$ fraction of the constraints.*
- NO *case: Every labeling satisfies less than $1 - O(\frac{\eta}{\lambda}\log(\frac{\lambda}{\eta}))$ fraction of the constraints.*

A consequence of the result is that when the UNIQUE GAMES instance is $(1-\eta)$-satisfiable and $\lambda \gg \eta$, the algorithm finds a labeling to the UNIQUE GAMES instance that satisfies 99% of the constraints. An important feature of the algorithm is that its performance does not depend on the number of labels k.

Comparison to previous work.

Most of the algorithms for UNIQUE GAMES (which can be viewed as attempts to disprove the UGC) are based on the SDP relaxation proposed by Feige and Lovász [9]. Their paper showed that if the UNIQUE GAMES instance is unsatisfiable, then the value of the SDP relaxation is bounded away from 1, though they did not give quantitative bounds. Khot [12] gave a SDP-rounding algorithm to find a labeling that satisfies $1 - O(k^2\eta^{1/5}\log(1/\eta))$ fraction of the constraints when there exists a labeling that satisfies $1 - \eta$ fraction of the constraints. The SDP's analysis was then revisited by many papers. On an $(1-\eta)$-satisfiable instance, these papers obtain a labeling that satisfies at least $1 - f(\eta, n, k)$ fraction of the constraints where $f(\eta, n, k)$ is $\sqrt[3]{\eta \log n}$ in Trevisan [22], $\sqrt{\eta \log k}$ in Charikar, Makarychev, and Makarychev [4], $\eta\sqrt{\log n \log k}$ in Chlamtac, Makarychev, and Makarychev [6], and $\eta \log n$ via an LP based approach in Gupta and Talwar [10]. Trevisan [22] also gave a combinatorial algorithm that works well on expanders. On an $(1 - \eta)$-satisfiable instance, he showed how to obtain a labeling satisfying $1 - \eta \log n \log \frac{1}{\lambda}$ fraction of the constraints. All these results require η to go to 0 as either n or k go to infinity in order to maintain their applicability[1]. Our main result is the first of its kind where under an additional promise of a natural graph property, namely expansion, the performance of the algorithm is independent of k and n. Furthermore, our analysis steps away from the edge-by-edge analysis of previous papers in favor of a more global analysis of correlations, which may be useful for other problems. We also provide an integrality gap for this SDP to show that, quantitatively, our main result is tight up to log factors.

We note that if we impose a certain structure on our constraints, namely if they are of the form ΓMAX2LIN, our results continue to hold when λ is replaced by stronger relaxations for the expansion of G, similar in spirit to the relaxations obtained by SDP hierarchies [18, 16, 17]. In particular, we show that λ can be replaced by the value of such a relaxation for expansion of G after a constant number of rounds.

Application to parallel repetition.

Since our main result shows an upper bound on the integrality gap for the standard SDP, the analysis of Feige and Lovász [9] allows us to prove (see Section 3) a *parallel repetition theorem* for unique games with expansion. We show that the r-round parallel repetition value of a UNIQUE GAMES instance with value at most $1-\varepsilon$ is at most $(1 - \Omega(\varepsilon \cdot \lambda/\log \frac{1}{\varepsilon}))^r$. In addition to providing an alternate proof, when $\lambda \gg \varepsilon^2 \log(1/\varepsilon)$, this is better than the general bound for nonunique games, where the best bound is $(1-\Omega(\varepsilon^3/\log k))^r$ by Holenstein [11], improving upon Raz's Theorem [19]. We note that recently, Safra and Schwartz [21] also showed a parallel repetition theorem for games with expansion, and their result works even for general games. Also, Rao [20] has proved a better parallel repetition theorem for, so called, projection games, which are more general than unique games. His result does not assume any expansion of the game graph.

Randomly generated games.

For many constraint satisfaction problems such as 3SAT, solving randomly generated instances is of great interest. For instance, proving unsatisfiability of formulae on n variables and with dn randomly chosen clauses seems very difficult for $d \ll \sqrt{n}$. Our results suggest that it will be hard to define a model of probabilistic generation for unique games that will result in very difficult instances, since the natural models all lead to instances with high expansion.

2. MAIN RESULT

Let $\mathcal{U} = (G(V,E), [k], \{\pi_{uv}\}_{(u,v)\in E})$ be a UNIQUE GAMES instance. We use the standard SDP relaxation for the problem, which involves finding a *vector assignment* for each vertex.

For every $u \in V$, we associate a set of k orthogonal vectors $\{\mathbf{u}_1, \ldots, \mathbf{u}_k\}$. The intention is that if $i_0 \in [k]$ is a label for vertex $u \in V$, then $\mathbf{u}_{i_0} = \sqrt{k}\mathbf{1}$, and $\mathbf{u}_i = \mathbf{0}$ for all $i \neq i_0$. Here, $\mathbf{1}$ is some fixed unit vector and $\mathbf{0}$ is the zero-vector. Of course, in a general solution to the SDP this may no longer be true and $\{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_k\}$ is just any set of orthogonal vectors.

Our proof will use the fact that the objective function (1) can be rewritten as

$$1 - \tfrac{1}{2}\mathbf{E}_{e=(u,v)\in E}\mathbf{E}_{i\in[k]}\left\|\mathbf{u}_i - \mathbf{v}_{\pi_{uv}(i)}\right\|^2 \qquad (5)$$

2.1 Overview

Let $\mathcal{U} = (G(V,E), [k], \{\pi_{uv}\}_{(u,v)\in E})$ be a UNIQUE GAMES instance, and let $\{\mathbf{u}_i\}_{u\in V, i\in[k]}$ be an optimal SDP solution. Assume wlog that its value is $1-\eta$, since otherwise we know already that the instance is a NO instance. How do we extract a labeling from the vector solution?

[1]On the other hand, the UGC allows k to grow arbitrarily as a function of η, and therefore, all known algorithms fall short of disproving UGC.

$$\text{Maximize} \quad \mathbf{E}_{e=(u,v)\in E}\mathbf{E}_{i\in[k]} \left\langle \mathbf{u}_i, \mathbf{v}_{\pi_{uv}(i)} \right\rangle \tag{1}$$

Subject to

$$\forall\, u \in V \quad \mathbf{E}_{i\in[k]} \, \|\mathbf{u}_i\|^2 = 1 \tag{2}$$

$$\forall\, u \in V \;\; \forall\, i \neq j \qquad \langle \mathbf{u}_i, \mathbf{u}_j \rangle = 0 \tag{3}$$

$$\forall\, u, v \in V \;\; \forall\, i, j \qquad \langle \mathbf{u}_i, \mathbf{v}_j \rangle \geq 0 \tag{4}$$

Figure 1: SDP **for** UNIQUE GAMES

Constraint (2) suggests an obvious way to view the vectors corresponding to vertex u as a distribution on labels, namely, one that assigns probability label i to u with probability $\frac{1}{k}\|\mathbf{u}_i\|^2$. The most naive idea for a rounding algorithm would be to use this distribution to pick a label for each vertex, where the choice for different vertices is made independently. Of course, this doesn't work since all labels could have equal probability under this distribution and thus the chance that the labels i, j picked for vertices u, v in an edge e satisfy $\pi_e(i) = j$ is only $1/k$.

More sophisticated roundings use the fact that if the SDP value is $1 - \eta$ for some small η, then the vector assignments to the vertices of an average edge $e = (u, v)$ are highly correlated, in the sense that for "many" i, $\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_{\pi(i)} \rangle > 1 - \Omega(\eta)$ where $\overline{\mathbf{u}}_i$ denotes the unit vector in the direction of $\mathbf{u}_i$. This suggests many rounding possibilities as explored in previous papers [12, 22, 4], but counterexamples [15] show that this edge-by-edge analysis can only go so far: high correlation for edges does not by itself imply that a good global assignment exists.

The main idea in our work is to try to understand and exploit correlations in the vector assignments for vertices that are not necessarily adjacent. If u, v are not adjacent vertices we can try to identify the correlation between their vector assignments by noting that since the $\mathbf{v}_j$'s are mutually orthogonal, for each $\mathbf{u}_i$ there is at most one $\mathbf{v}_j$ such that $\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j \rangle > 1/\sqrt{2}$. Thus we can set up a maximal partial matching among their labels where the matching contains label pairs (i, j) such that $\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j \rangle > 1/\sqrt{2}$. The vector assignments to the two vertices can be thought of as highly correlated if the sum of squared ℓ_2 norm of all the $\mathbf{u}_i$'s (resp, all $\mathbf{v}_j$'s) involved in this matching is close to k. (This is a rough idea; see precise definition later.)

Our main contribution is to show that if the constraint graph is an expander then high correlation over edges necessarily implies high expected correlation between a randomly-chosen pair of vertices (which may be quite distant in the constraint graph). We also show that this allows us to construct a good global assignment. This is formalized below.

2.2 Rounding procedure and correctness proof

Now we describe our randomized rounding procedure $\mathcal{R}$, which outputs a labeling $\Lambda_{\text{alg}} \colon V \to [k]$. This uses a more precise version of the greedy matching outlined in the above overview. For a pair u, v of vertices (possibly nonadjacent), let $\sigma_{uv} \colon [k] \to [k]$ be a bijective mapping that maximizes $\mathbf{E}_{i\in[k]} \langle \mathbf{u}_i, \mathbf{v}_{\sigma_{uv}(i)} \rangle$; note that it can be efficiently found using max-weight bipartite matching. The procedure is as follows:

1. Pick a random vertex u.
2. Pick a label i for u from the distribution, where every label $i' \in [k]$ has probability $\frac{1}{k}\|\mathbf{u}_{i'}\|^2$.
3. Define $\Lambda_{\text{alg}}(v) := \sigma_{uv}(i)$ for every vertex $v \in V$.

(Of course, the rounding can be trivially derandomized since there are only nk choices for u, i.)

To analyse this procedure we define the *distance* $\rho(u, v)$ of a pair u, v of vertices as

$$\rho(u,v) := \frac{1}{2}\mathbf{E}_{i\in[k]} \left\| \mathbf{u}_i - \mathbf{v}_{\sigma_{uv}(i)} \right\|^2 \tag{6}$$

$$= 1 - \mathbf{E}_{i\in[k]} \left\langle \mathbf{u}_i, \mathbf{v}_{\sigma_{uv}(i)} \right\rangle \qquad \text{(using (2))}.$$

We think of two vertices u and v as highly correlated if $\rho(u, v)$ is small, i.e., $\mathbf{E}_{i\in[k]} \langle \mathbf{u}_i, \mathbf{v}_{\sigma_{uv}(i)} \rangle \approx 1$.

The following easy lemma shows that if the average vertex pair in G is highly correlated, then the above rounding procedure is likely to produce a good a labeling. Here we assume that G is a regular graph. Using standard arguments, all results can be generalized to the case of non-regular graphs. A proof of the lemma can be found in Section 2.2.

LEMMA 2.1 (GLOBAL CORR. $\implies$ HIGH VALUE). *The expected fraction of constraints satisfied by the labeling Λ_{alg} computed by the rounding procedure is*

$$\mathbf{E}_{\Lambda_{\text{alg}} \leftarrow \mathcal{R}} [\,\mathsf{val}(\Lambda_{\text{alg}})\,] \geq 1 - 3\eta - 6\mathbf{E}_{u,v\in V} [\,\rho(u,v)\,].$$

It is easy to see that if the SDP value is $1 - \eta$ then the average correlation on edges is high. For an edge $e = (u, v)$ in G, let $\eta_e := \frac{1}{2}\mathbf{E}_{i\in[k]} \left\| \mathbf{u}_i - \mathbf{v}_{\pi_{uv}(i)} \right\|^2$. Note, $\mathbf{E}_e[\eta_e] = \eta$. Then we have

$$\rho(u,v) = \frac{1}{2}\mathbf{E}_{i\in[k]} \left\| \mathbf{u}_i - \mathbf{v}_{\sigma_{uv}(i)} \right\|^2 = 1 - \mathbf{E}_{i\in[k]} \left\langle \mathbf{u}_i, \mathbf{v}_{\sigma_{uv}(i)} \right\rangle$$

$$\leq 1 - \mathbf{E}_{i\in[k]} \left\langle \mathbf{u}_i, \mathbf{v}_{\pi_{uv}(i)} \right\rangle = \eta_e$$

$$\text{(since } \sigma_{uv} \text{ is a max-weight matching).}$$

As mentioned in the overview, we show that high correlation on edges implies (when the constraint graph is an expander) high correlation on the average pair of vertices. The main technical contribution in this proof is a way to view a vector solution to the above SDP as a vector solution for SPARSEST CUT. This involves mapping any sequence of k vectors to a *single* vector in a nicely continuous way, which allows us to show that the distances $\rho(u, v)$ essentially behave like squared Euclidean distances. We defer the proof of the next lemma to Section 2.3.

LEMMA 2.2 (LOW DISTORTION EMBEDDING OF ρ). *For every positive even integer t and every SDP solution $\{\mathbf{u}_i\}_{u\in V, i\in[k]}$, there exists a set of vectors $\{\mathbf{V}_u\}_{u\in V}$ such that for every pair u, v of vertices*

$$\tfrac{1}{2t} \|\mathbf{V}_u - \mathbf{V}_v\|^2 \leq \rho(u,v) \leq \|\mathbf{V}_u - \mathbf{V}_v\|^2 + O(2^{-t/2}).$$

COROLLARY 2.3 (LOCAL CORR. $\implies$ GLOBAL CORR.).

$$\mathbf{E}_{u,v\in V} [\,\rho(u,v)\,] \leq 2t\eta/\lambda + O(2^{-t/2}).$$

PROOF. We use the following characterization of λ for regular graphs G

$$\lambda = \min \frac{\mathbf{E}_{(u,v)\in E} \|\mathbf{z}_u - \mathbf{z}_v\|^2}{\mathbf{E}_{u,v\in V} \|\mathbf{z}_u - \mathbf{z}_v\|^2}, \tag{7}$$

where the minimum is over all sets of vectors $\{\mathbf{z}_u\}_{u\in V}$. This characterization also shows that λ scaled by

$n^2/|E|$ is a relaxation for the SPARSEST CUT problem $\min_{\emptyset\neq S\subset V} |E(S,\overline{S})|/|S||\overline{S}|$ of G . Now using the previous Lemma we have

$$\begin{aligned}\mathbf{E}_{u,v\in V}[\,\rho(u,v)\,] &\le \mathbf{E}_{u,v\in V}\,\|\mathbf{V}_u-\mathbf{V}_v\|^2 + O(2^{-t/2})\\ &\le \tfrac{1}{\lambda}\mathbf{E}_{(u,v)\in E}\,\|\mathbf{V}_u-\mathbf{V}_v\|^2 + O(2^{-t/2})\\ &\le \tfrac{2t}{\lambda}\mathbf{E}_{(u,v)\in E}[\,\rho(u,v)\,] + O(2^{-t/2}).\end{aligned}$$

□

By combining the Corollary 2.3 and Lemma 2.1, we can show the following theorem.

THEOREM 2.4 (IMPLIES THEOREM 1.1). *There is a polynomial time algorithm that computes a labeling Λ with*

$$\mathsf{val}(\Lambda) \ge 1 - O\left(\tfrac{\eta}{\lambda}\log\left(\tfrac{\lambda}{\eta}\right)\right)$$

if the optimal value of the SDP *in Figure 1 for $\mathcal{U}$ is $1-\eta$.*

PROOF. By Corollary 2.3 and Lemma 2.1, the labeling $\Lambda_{\rm alg}$ satisfies a $1-O(t\eta/\lambda+2^{-t/2})$ fraction of the constraints of $\mathcal{U}$. If we choose t to be an integer close to $2\log(\lambda/\eta)$, it follows that $\mathsf{opt}(\mathcal{U}) \ge 1 - O(\frac{\eta}{\lambda}\log(\frac{\lambda}{\eta}))$. Since the rounding procedure $\mathcal{R}$ can easily be derandomized, a labeling Λ with $\mathsf{val}(\Lambda) \ge 1 - O(\frac{\eta}{\lambda}\log(\frac{\lambda}{\eta}))$ can be computed in polynomial time. □

We can show that the integrality gap (in terms of expansion) implied above is tight up to a logarithmic factor. The next theorem can be derived using the techniques in [15, 7]. The proof is deferred to Appendix A.3.

THEOREM 2.5. *For every $\eta > 0$ small enough and for every n large enough, there is a* UNIQUE GAMES *instance $\mathcal{U}_\eta$ on $\Theta(\log(n))$ labels and a constraint graph with $\lambda = \Omega(\eta)$, such that (1) $\mathsf{opt}(\mathcal{U}_\eta) \le 1/\log^\eta n$, and (2) there is an* SDP *solution for $\mathcal{U}_\eta$ of value at least $1-O(\eta)$.*

The next theorem shows that, assuming UGC, the approximation guarantee of Theorem 2.4 cannot be improved by more than a constant factor. The proof is deferred to Appendix A.3.

THEOREM 2.6. *Assuming* UGC, *for every $\eta,\delta > 0$, there exists $k = k(\eta,\delta)$ such that for a* UNIQUE GAMES *instance $\mathcal{U} = (G(V,E),[k],\{\pi_{uv}\}_{(u,v)\in E})$ it is* NP*-hard to distinguish between*

- *YES Case:* $\mathsf{opt}(\mathcal{U}) \ge 1-\eta$,
- *NO Case:* $\mathsf{opt}(\mathcal{U}) \le \delta$ *and* $\lambda > \Omega(\eta)$.

2.3 Proof of Lemma 2.1

We consider the labeling $\Lambda_{\rm alg}$ computed by the randomized rounding procedure $\mathcal{R}$. Recall that $\Lambda_{\rm alg}(v) = \sigma_{uv}(i)$ where the vertex u is chosen uniformly at random and the label i is chosen with probability proportional to $\|\mathbf{u}_i\|^2$. For notational ease we assume that σ_{uu} is the identity permutation and σ_{uv} is the inverse permutation of σ_{vu}. The following claim gives an estimate on the probability that the constraint between an edge $e=\{v,w\}$ is satisfied by $\Lambda_{\rm alg}$. Here we condition on the choice of u.

CLAIM 2.7. *For every vertex u and every edge $e=(v,w)$, $\mathbf{Pr}_{\Lambda_{\rm alg}}[\Lambda_{\rm alg}(w) \neq \pi_{v,w}(\Lambda_{\rm alg}(v)) \mid u] \le 3\cdot(\rho(u,v)+\eta_e+\rho(w,u))$.*

PROOF. We may assume that both σ_{uv} and σ_{uw} are the identity permutation. Let $\pi = \pi_{vw}$. First note $\mathbf{Pr}_{\Lambda_{\rm alg}}[\Lambda_{\rm alg}(w)\neq\pi(\Lambda_{\rm alg}(v)) \mid u] = \mathbf{E}_{i\in[k]}\left[\|\mathbf{u}_i\|^2\chi_{i\neq\pi(i)}\right]$, where $\chi_{\mathcal{E}}$ denotes the indicator random variable for an event $\mathcal{E}$. By orthogonality of the vectors $\{\mathbf{u}_i\}_{i\in[k]}$, it follows that

$$\begin{aligned}\mathbf{E}_{i\in[k]}\left[\|\mathbf{u}_i\|^2\chi_{i\neq\pi(i)}\right] &\le \tfrac{1}{2}\mathbf{E}_{i\in[k]}\left[\left(\|\mathbf{u}_i\|^2+\left\|\mathbf{u}_{\pi(i)}\right\|^2\right)\chi_{i\neq\pi(i)}\right]\\ &= \tfrac{1}{2}\mathbf{E}_{i\in[k]}\left\|\mathbf{u}_i-\mathbf{u}_{\pi(i)}\right\|^2.\end{aligned}$$

By triangle inequality, $\left\|\mathbf{u}_i-\mathbf{u}_{\pi(i)}\right\| \le \|\mathbf{u}_i-\mathbf{v}_i\| + \left\|\mathbf{v}_i-\mathbf{w}_{\pi(i)}\right\| + \left\|\mathbf{w}_{\pi(i)}-\mathbf{u}_{\pi(i)}\right\|$. Now we square both sides of the inequality and take expectations, $\mathbf{E}_{i\in[k]}\left\|\mathbf{u}_i-\mathbf{u}_{\pi(i)}\right\|^2 \le 3\mathbf{E}_{i\in[k]}\|\mathbf{u}_i-\mathbf{v}_i\|^2 + 3\mathbf{E}_{i\in[k]}\left\|\mathbf{v}_i-\mathbf{w}_{\pi(i)}\right\|^2 + 3\mathbf{E}_{i\in[k]}\left\|\mathbf{w}_{\pi(i)}-\mathbf{u}_{\pi(i)}\right\|^2 = 6\rho(u,v)+6\eta_e+6\rho(w,u)$. □

PROOF OF LEMMA 2.1. From Claim 2.7 it follows

$$\mathbf{E}_{\Lambda_{\rm alg}}[\,\mathsf{val}(\Lambda_{\rm alg})\,] \ge 1-3\mathbf{E}_{u\in V}\mathbf{E}_{e=(vw)\in E}\left[\rho(u,v)+\eta_e+\rho(w,u)\right].$$

Since G is a regular graph, both (u,v) and (w,u) are uniformly distributed over all pairs of vertices. Hence $\mathbf{E}_{\Lambda_{\rm alg}}[\,\mathsf{val}(\Lambda_{\rm alg})\,] \ge 1-3\eta-6\mathbf{E}_{u,v\in V}[\,\rho(u,v)\,]$. □

2.4 Proof of Lemma 2.2; the tensoring trick

Let t be an integer greater than or equal to 4, and $\{\mathbf{u}_i\}_{u\in V,i\in[k]}$ be an SDP solution for $\mathcal{U}$. Define $\overline{\mathbf{u}}_i = \frac{1}{\|\mathbf{u}_i\|}\mathbf{u}_i$ and $\mathbf{V}_u = \frac{1}{\sqrt{k}}\sum_{i\in[k]}\|\mathbf{u}_i\|\,\overline{\mathbf{u}}_i^{\otimes t}$, where $\otimes t$ denotes t-wise tensoring. Notice that the vectors $\mathbf{V}_u$ are unit vectors. Consider a pair u,v of vertices in G. The following claim implies the lower bound on $\rho(u,v)$ in Lemma 2.2.

CLAIM 2.8. $\|\mathbf{V}_u-\mathbf{V}_v\|^2 \le t\cdot\mathbf{E}_{i\in[k]}\left\|\mathbf{u}_i-\mathbf{v}_{\sigma_{uv}(i)}\right\|^2$

PROOF. Since $\mathbf{V}_u$ is a unit vector for each u, it suffices to prove $\langle\mathbf{V}_u,\mathbf{V}_v\rangle \ge 1-t\rho(u,v)$. Let $\sigma=\sigma_{uv}$. By Cauchy-Schwarz,

$$\tfrac{1}{k}\textstyle\sum_i\|\mathbf{u}_i\|\|\mathbf{v}_{\sigma(i)}\| \le \tfrac{1}{k}\big(\sum_i\|\mathbf{u}_i\|^2\big)^{1/2}\big(\sum_i\|\mathbf{v}_{\sigma(i)}\|^2\big)^{1/2} \le 1.$$

Thus there is some $\alpha \ge 1$ such that the following random variable X is well-defined: it takes value $\langle\overline{\mathbf{u}}_i,\overline{\mathbf{v}}_{\sigma(i)}\rangle$ with probability $\alpha\cdot\frac{1}{k}\|\mathbf{u}_i\|\|\mathbf{v}_{\sigma(i)}\|$. By Jensen's Inequality, $(\mathbf{E}[X])^t \le \mathbf{E}[X^t]$. Hence,

$$\begin{aligned}1-\rho(u,v)t \le (1-\rho(u,v))^t &= \left(\mathbf{E}_{i\in[k]}\left[\|\mathbf{u}_i\|\|\mathbf{v}_{\sigma(i)}\|\,\langle\overline{\mathbf{u}}_i,\overline{\mathbf{v}}_{\sigma(i)}\rangle\right]\right)^t\\ &= (\mathbf{E}[X/\alpha])^t \le (\mathbf{E}[X])^t/\alpha\\ &\le \mathbf{E}[X^t/\alpha] = \langle\mathbf{V}_u,\mathbf{V}_v\rangle.\end{aligned}$$

This proves the claim. □

Matching between two label sets.

In order to finish the proof of Lemma 2.2, it remains to prove the upper bound on $\rho(u,v)$ in terms of the distance $\|\mathbf{V}_u-\mathbf{V}_v\|^2$. For this part of the proof, it is essential that the vectors $\mathbf{V}_u$ are composed of (high) tensor powers of the vectors $\mathbf{u}_i$. For a pair u,v of vertices, consider the following set of label pairs

$$M = \{(i,j)\in[k]\times[k] \mid \langle\overline{\mathbf{u}}_i,\overline{\mathbf{v}}_j\rangle^2 > 1/2\}.$$

Since $\{\overline{\mathbf{u}}_i\}_{i\in[k]}$ and $\{\overline{\mathbf{v}}_j\}_{j\in[k]}$ are sets of ortho-normal vectors, M as bipartite graph between the labels for u and the labels for v is a (partial) matching, that is, every label for u

has at most one neighbor among the labels for v. Let σ be an arbitrary permutation of $[k]$ that agrees with the M on the matched labels, i.e., for all $(i,j) \in M$, we have $\sigma(i) = j$. The following claim shows the upper bound on $\rho(u,v)$ of Lemma 2.2.

Claim 2.9.

$$\frac{1}{2}\mathbf{E}_{i\in[k]} \left\|\mathbf{u}_i - \mathbf{v}_{\sigma(i)}\right\|^2 \le \frac{1}{2}\left\|\mathbf{V}_u - \mathbf{V}_v\right\|^2 + O(2^{-t/2}).$$

Proof. Let $\delta = \|\mathbf{V}_u - \mathbf{V}_v\|^2$. Note that

$$\tfrac{1}{k}\sum_{i,j} \|\mathbf{u}_i\|\|\mathbf{v}_j\| \left\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j\right\rangle^t = 1 - \delta/2. \tag{8}$$

We may assume that σ is the identity permutation. Then, $\rho(u,v)$ is at most

$$\begin{aligned}
\tfrac{1}{2}\mathbf{E}_{i\in[k]} \|\mathbf{u}_i - \mathbf{v}_i\|^2 &= 1 - \mathbf{E}_{i\in[k]} \langle \mathbf{u}_i, \mathbf{v}_i\rangle \\
&\le 1 - \tfrac{1}{k}\sum_{i\in[k]} \|\mathbf{u}_i\|\,\|\mathbf{v}_i\| \left\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_i\right\rangle^t \\
&\qquad\qquad (\text{using } \langle \mathbf{u}_i, \mathbf{v}_i\rangle \ge 0) \\
&= \delta/2 + \tfrac{1}{k}\sum_{i\ne j} \|\mathbf{u}_i\|\|\mathbf{v}_j\| \left\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j\right\rangle^t \\
&\qquad\qquad (\text{by } (8)) \\
&= \delta/2 + \langle \mathbf{p}, A\,\mathbf{q}\rangle\,,
\end{aligned}$$

where $p_i = \frac{1}{\sqrt{k}}\|\mathbf{u}_i\|$, $q_j = \frac{1}{\sqrt{k}}\|\mathbf{v}_j\|$, $A_{ii} = 0$, and for $i \ne j$, $A_{ij} = \langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j\rangle^t$. Since both $\mathbf{p}$ and $\mathbf{q}$ are unit vectors, $\langle \mathbf{p}, A\,\mathbf{q}\rangle$ is bounded by the largest singular value of A. As the matrix A has only non-negative entries, its largest singular value is bounded by the maximum sum of a row or a column. By symmetry, we may assume that the first row has the largest sum among all rows and columns. We rearrange the columns in such a way that $A_{11} \ge A_{12} \ge \ldots \ge A_{1k}$. Since $\overline{\mathbf{u}}_1$ is a unit vector and $\{\overline{\mathbf{v}}_j\}_{j\in[k]}$ is a set of orthonormal vectors, we have $\sum_j \langle \overline{\mathbf{u}}_1, \overline{\mathbf{v}}_j\rangle^2 \le 1$. Hence, $\langle \overline{\mathbf{u}}_1, \overline{\mathbf{v}}_j\rangle^2 \le 1/j$ and therefore $A_{1j} \le (1/j)^{t/2}$. On the other, every entry of A is at most $2^{-t/2}$, since all pairs (i,j) with $\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j\rangle^2 > 1/2$ participate in the matching M, and hence, $A_{ij} = 0$ for all i,j with $\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j\rangle^2 > 1/2$. It follows that the sum of the first row can be upper bounded by

$$\sum_{j\in[k]} A_{1j} \le A_{11} + \sum_{j\ge 2}^{\infty} (\tfrac{1}{j})^{t/2} \le 2^{-t/2} + \sum_{j\ge 2}^{\infty} (\tfrac{1}{j})^{t/2} = O(2^{-t/2}).$$

We conclude that the largest singular value of A is at most $O(2^{-t/2})$, and thus $\rho(u,v)$ can be upper bounded by $\delta/2 + O(2^{-t/2}) = \frac{1}{2}\|\mathbf{V}_u - \mathbf{V}_v\| + O(2^{-t/2})$, as claimed. □

3. STRONGER RELAXATIONS

In this section, we consider stronger SDP relaxations for Unique Games and for Sparsest Cut. A systematic way to obtain stronger relaxations is provided by SDP hierarchies. We choose to state our results in terms of Lasserre's SDP hierarchy [16, 17]. The results in this section apply only to a special case of Unique Games, called ΓMAX2LIN. We say a Unique Games instance $\mathcal{U} = (G(V,E), [k], \{\pi_{uv}\}_{(u,v)\in E})$ has ΓMAX2LIN form, if the label set $[k]$ can be identified with the group $\mathbb{Z}_k$ in such a way that every constraint permutation π_{uv} satisfies $\pi_{uv}(i+s) = \pi_{uv}(i) + s \in \mathbb{Z}_k$ for all $s, i \in \mathbb{Z}_k$. In other words, π_{uv} encodes a constraint of the form $x_u - x_v = c_{uv} \in \mathbb{Z}_k$. The ΓMAX2LIN property implies that we can find an optimal SDP solution $\{\mathbf{u}_i\}_{i\in[k]}$ for $\mathcal{U}$ that is *shift-invariant*, i.e., for all $s \in \mathbb{Z}_k$ we have $\langle \mathbf{u}_{i+s}, \mathbf{v}_{j+s}\rangle = \langle \mathbf{u}_i, \mathbf{v}_j\rangle$. In particular, every vector $\mathbf{u}_i$ has unit norm.

Alternative Embedding for ΓMAX2LIN.

The following lemma can be seen as alternative to Lemma 2.2. We emphasize that the lemma only holds for ΓMAX2LIN instances and shift-invariant SDP solutions.

Lemma 3.1. *Let Λ_{opt} be a labeling for $\mathcal{U}$ with $\mathsf{val}(\Lambda_{\text{opt}}) = 1 - \varepsilon$. Then the set of vectors $\{\mathbf{V}_u\}_{u\in V}$ with $\mathbf{V}_u = \mathbf{u}_{\Lambda_{\text{opt}}(u)}$ has the following two properties:*

1. *$\rho(u,v) \le \frac{1}{2}\|\mathbf{V}_u - \mathbf{V}_v\|^2$ for every pair u, v of vertices*
2. *$\frac{1}{2}\mathbf{E}_{(u,v)\in E}\|\mathbf{V}_u - \mathbf{V}_v\|^2 \le \eta + 2\varepsilon$*

Together with Lemma 2.1, the above lemma implies that the randomized rounding procedure $\mathcal{R}$ computes a labeling that satisfies at least a $1 - O(\varepsilon/\lambda)$ fraction of the constraints of $\mathcal{U}$, whenever $\mathsf{opt}(\mathcal{U}) \ge 1 - \varepsilon$. In this sense, the above lemma allows to prove the main result of this paper for the special case of ΓMAX2LIN.

Proof. Item 1 holds, since, by shift invariance,

$$\begin{aligned}
\rho(u,v) = \tfrac{1}{2}\mathbf{E}_{i\in[k]}\left\|\mathbf{u}_i - \mathbf{v}_{\sigma_{uv}(i)}\right\|^2 &= \tfrac{1}{2}\left\|\mathbf{u}_{\Lambda_{\text{opt}}(u)} - \mathbf{v}_{\sigma_{uv}(\Lambda_{\text{opt}}(u))}\right\|^2 \\
&\le \tfrac{1}{2}\left\|\mathbf{u}_{\Lambda_{\text{opt}}(u)} - \mathbf{v}_{\Lambda_{\text{opt}}(v)}\right\|^2.
\end{aligned}$$

Here we could assume, again by shift invariance, that $\left\|\mathbf{u}_i - \mathbf{v}_{\sigma_{uv}(i)}\right\|^2 = \min_j \|\mathbf{u}_i - \mathbf{v}_j\|^2$ for all i.

It remains to verify Item 2. By shift invariance,

$$\eta_{uv} = \tfrac{1}{2}\mathbf{E}_{i\in[k]}\left\|\mathbf{u}_i - \mathbf{v}_{\pi_{uv}(i)}\right\|^2 = \tfrac{1}{2}\left\|\mathbf{u}_{\Lambda_{\text{opt}}(u)} - \mathbf{v}_{\pi_{uv}(\Lambda_{\text{opt}}(u))}\right\|^2.$$

Hence, if Λ_{opt} satisfies the constraint on an edge $(u,v) \in E$, then $\frac{1}{2}\|\mathbf{V}_u - \mathbf{V}_v\|^2 = \eta_{uv}$. On the other hand, $\frac{1}{2}\|\mathbf{V}_u - \mathbf{V}_v\|^2 \le 2$ because every vector $\mathbf{V}_u$ has unit norm. Finally, since a $1-\varepsilon$ fraction of the edges is satisfied by Λ_{opt},

$$\mathbf{E}_{(u,v)\in E}\frac{1}{2}\|\mathbf{V}_u - \mathbf{V}_v\|^2 \le (1-\varepsilon)\cdot \mathbf{E}_{(u,v)\in E}[\,\eta_{uv}\,] + \varepsilon \cdot 2.$$

□

Stronger Relaxations for Sparsest Cut.

Let r be a positive integer. Denote by $\mathcal{I}$ the set of all subsets of V that have cardinality at most r. For every subset $I \in \mathcal{I}$, we have a variable $\mathbf{x}_I$. We consider a strengthening of the spectral relaxation for Sparsest Cut (Figure 2).

$$\text{Minimize} \quad \frac{\mathbf{E}_{(u,v)\in E}\|\mathbf{z}_u - \mathbf{z}_v\|^2}{\mathbf{E}_{u,v\in V}\|\mathbf{z}_u - \mathbf{z}_v\|^2} \tag{9}$$

Subject to

$$\forall\, I, J \in \mathcal{I},\ \forall\, I', J' \in \mathcal{I} \qquad \langle \mathbf{x}_I, \mathbf{x}_J\rangle = \langle \mathbf{x}_{I'}, \mathbf{x}_{J'}\rangle \quad \text{if } I \cup J = I' \cup J' \tag{10}$$

$$\forall\, u \in V,\ \forall\, i \in [k] \qquad \mathbf{x}_{\{u\}} = \mathbf{z}_u \tag{11}$$

$$\|\mathbf{x}_\emptyset\|^2 = 1 \tag{12}$$

Figure 2: Stronger relaxation for Sparsest Cut.

The variables $\mathbf{x}_I$ are intended to have values $\mathbf{0}$ or $\mathbf{1}$, where $\mathbf{1}$ is some fixed unit vector. If the intended cut is $(S, \overline{S})$, we would assign $\mathbf{1}$ to all variables $\mathbf{x}_{\{u\}} = \mathbf{z}_u$ with $u \in S$. The variables $\mathbf{x}_I$ are relaxations of boolean variables x_I. The intended value of x_I is the product of the variables x_t, $t \in I$.

Let $z_r(G)$ denote the optimal value of the SDP in Figure 2. We have

$$\lambda \leq z_1(G) \leq \ldots \leq z_n(G) = \tfrac{n^2}{|E|} \min_{\emptyset \neq S \subset V} \frac{|E(S,\overline{S})|}{|S||\overline{S}|}.$$

It can also be seen that the relaxation $z_3(G)$ is at least as strong as the relaxation for Sparsest Cut considered in [2]. The relaxations $z_r(G)$ are inspired by Lasserre's SDP hierarchy [16, 17].

The proof of the following theorem is similar to the proof of Theorem 2.4. The main difference is that we use Lemma 3.1, instead of Lemma 2.2, in order to show that local correlation implies global correlation. By strengthening the SDP for Unique Games, the vectors $\mathbf{V}_u$ obtained from Lemma 3.1 can be extended to a solution for the stronger SDP for Sparsest Cut in Figure 2. This allows us to replace the parameter λ by the parameter $z_r(G)$ in the below theorem.

Theorem 3.2. *There is an algorithm that computes in time $(kn)^{O(r)}$ a labeling Λ with*

$$\mathsf{val}(\Lambda) \geq 1 - O(\varepsilon / z_r(G))$$

if $\mathsf{opt}(\mathcal{U}) \geq 1 - \varepsilon$ and $\mathcal{U}$ has ΓMAX2LIN form.

The proof of the above theorem is deferred to Appendix A.1.

4. PARALLEL REPETITION FOR EXPANDING UNIQUE GAMES

In this section, we consider *bipartite* unique games, i.e., Unique Games instances $\mathcal{U} = (G(V, W, E), [k], \{\pi_{vw}\}_{(v,w)\in E})$ such that $G(V, W, E)$ is a bipartite graph with bipartition (V, W). A bipartite unique game can be seen as a 2-prover, 1-round proof system [9]. The two parts V, W correspond to the two provers. The edge set E corresponds to the set of questions asked by the verifier to the two provers and π_{vw} is the accepting predicate for the question corresponding to the edge (v, w).

In this section, we give an upper bound on the *amortized value* $\overline{w}(\mathcal{U}) = \sup_r \mathsf{opt}(\mathcal{U}^{\otimes r})^{1/r}$ of bipartite unique game $\mathcal{U}$ in terms of the expansion of its constraint graph. Here $\mathcal{U}^{\otimes r}$ denotes the game obtained by playing the game $\mathcal{U}$ for r rounds in parallel. We follow an approach proposed by Feige and Lovász [9]. Their approach is based on the SDP in Figure 3, which is a relaxation for the value of a bipartite unique game. Let $\overline{\sigma}(\mathcal{U})$ denote the value of this SDP relaxation. The following theorem is a consequence of the fact $\overline{\sigma}(\mathcal{U}^{\otimes r}) = \overline{\sigma}(\mathcal{U})^r$.

Theorem 4.1 ([9]). *For every bipartite unique game $\mathcal{U}$, $\overline{w}(\mathcal{U}) \leq \overline{\sigma}(\mathcal{U})$.*

We observe that the SDP in Figure 1 cannot be much stronger than the relaxation $\overline{\sigma}(\mathcal{U})$. The proof mostly uses standard arguments. We defer it to Appendix A.2.

Lemma 4.2. *If $\overline{\sigma}(\mathcal{U}) = 1 - \eta$ then the value of the SDP in Figure 1 is at least $1 - 2\eta$.*

$$\text{Maximize} \quad \mathbf{E}_{(v,w)\in E} \mathbf{E}_{i\in[k]} \left\langle \mathbf{v}_i, \mathbf{w}_{\pi_{vw}(i)} \right\rangle \qquad (13)$$

Subject to

$$\forall\, v \in V,\ w \in W,\ i, j \in [k] \qquad \langle \mathbf{v}_i, \mathbf{w}_j \rangle \geq 0 \qquad (14)$$

$$\forall\, v \in V,\ v' \in V \qquad \textstyle\sum_{i,i'} |\langle \mathbf{v}_i, \mathbf{v}'_{i'} \rangle| \leq k \qquad (15)$$

$$\forall\, w \in W,\ w' \in W \qquad \textstyle\sum_{j,j'} |\langle \mathbf{w}_j, \mathbf{w}'_{j'} \rangle| \leq k \qquad (16)$$

Figure 3: Feige-Lovasz SDP for Unique Games

Theorem 4.3. *If $\mathcal{U}$ is 2-prover 1-round unique game on alphabet $[k]$ with value at most $1-\varepsilon$, then the value $\mathcal{U}$ played in parallel for r rounds is at most $(1 - \Omega(\varepsilon \cdot \lambda / \log \frac{1}{\varepsilon}))^r$, where G is the graph corresponding to the questions to the two provers. In particular, the amortized value $\overline{w}(\mathcal{U})$ is at most $1 - \Omega(\varepsilon \cdot \lambda / \log \frac{1}{\varepsilon})$.*

Proof. Following the approach in [9], it is sufficient to show $\overline{\sigma}(\mathcal{U}) \leq 1 - \Omega(\varepsilon\lambda / \log \frac{1}{\varepsilon})$. Suppose that $\overline{\sigma}(\mathcal{U}) = 1 - \eta$. Then by Lemma 4.2, the value of the SDP in Figure 1 is at least $1 - 2\eta$. By Theorem 2.4, it follows that $\mathsf{opt}(\mathcal{U}) \geq 1 - O(\eta \log \frac{\lambda}{\eta} / \lambda)$. On the other hand, $\mathsf{opt}(\mathcal{U}) \leq 1 - \varepsilon$. Hence, $\varepsilon = O(\eta \log \frac{\lambda}{\eta} / \lambda)$ and therefore $\eta = \Omega(\lambda\varepsilon \log \frac{1}{\varepsilon})$, as claimed. □

5. ON REDUCTIONS TO SPARSEST CUT

Let us return to the main motivation of this paper, namely, the possibility of a reduction from Unique Games to Sparsest Cut. As remarked in the Introduction, for such a reduction to work, we need to assume, at the very least, that the Unique Games problem is hard even when the constraint graph has *some* expansion. The question is how much (and exactly what kind of) expansion is required to prove inapproximability of Sparsest Cut and whether one can expect Unique Games to be hard with the required expansion, in light of Theorem 1.1.

Let us focus on the size of balanced cuts (say $(\frac{1}{10}, \frac{9}{10})$-balanced cuts) as a measure of expansion. As the algorithm in Theorem 1.1 shows, when $\lambda \gg \eta$, the problem is easy. In terms of balanced cuts, the algorithm shows that the problem is easy when every balanced cut in G has size $\gg \sqrt{\eta}$. Thus if we were to hypothesize that Unique Games are hard with expansion, we could only hypothesize hardness on instances where every balanced cut is of size at least η^t for some $t > \frac{1}{2}$. Would such a hypothesis be enough to prove inapproximability of Sparsest Cut? We do not know the answer, though we can prove inapproximability of Sparsest Cut under a somewhat stronger hypothesis (we omit the proof; the reduction is same as in [15]):

Hypothesis 5.1. *For some fixed t, $\frac{1}{2} < t < 1$, for every $\eta, \delta > 0$, there exists $k = k(\eta, \delta)$ such that for a Unique Games instance $\mathcal{U}$ with k labels, it is NP-hard to distinguish between:*

- *YES Case: $\mathsf{opt}(\mathcal{U}) \geq 1 - \eta$.*
- *NO Case: $\mathsf{opt}(\mathcal{U}) \leq \delta$ and for any partition of V into three sets A, B, C, such that $|C| \leq \frac{|V|}{1000}$ and $|A|, |B| \geq \frac{|V|}{10}$, there are at least η^t fraction of the edges between the sets A and B.*

Theorem 5.2. *If Hypothesis 5.1 is true, then for any $\eta > 0$, it is NP-hard to distinguish whether a given graph*

H has a balanced cut of size at most η, or every balanced cut is of size at least η^t. In particular, the BALANCED SEPARATOR *problem, and therefore the* SPARSEST CUT *problem, is hard to approximate within any constant factor.*

Acknowledgments

We like to thank James Lee and Ryan O'Donnell for valuable discussions.

6. REFERENCES

[1] Christoph Ambühl, Monaldo Mastrolilli, and Ola Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *Annual Symposium on Foundations of Computer Science*, number 48, 2007.

[2] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 36, pages 222–231, 2004.

[3] Jin-yi Cai, Anne Condon, and Richard J. Lipton. Playing games of incomplete information. In *the 7th Annual Symposium on Theoretical Aspects of Computer Science*, pages 58–69, 1990.

[4] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for unique games. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 38, pages 205–214, 2006.

[5] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. In *Computational Complexity*, 15(2):94–114, 2006.

[6] Eden Chlamtac, Konstantin Makarychev, Yury Makarychev. How to play unique games using embeddings. In *Annual Symposium on Foundations of Computer Science*, number 47, pages 687–696, 2006.

[7] Nikhil Devanur, Subhash Khot, Rishi Saket, and Nisheeth K. Vishnoi. Integrality gaps for sparsest cut and minimum linear arrangement problems. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 38, pages 537–546, 2006.

[8] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 38, pages 344–353, 2006.

[9] Uriel Feige and László Lovász. Two-prover one-round proof systems, their power and their problems. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 24, pages 733–744, 1992.

[10] Anupam Gupta and Kunal Talwar. Approximating unique games. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 99–106, 2006.

[11] Thomas Holenstein. Parallel repetition: simplifications and the no-signaling case. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 39, pages 411–419, 2007.

[12] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the ACM Symposium on the Theory of Computing*, number 34, pages 767–775, 2002.

[13] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable CSPs? In *Annual Symposium on Foundations of Computer Science*, number 45, pages 146–154, 2004.

[14] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\varepsilon$. In *Proceedings of the Annual IEEE Conference on Computational Complexity*, number 18, pages 379–386, 2003.

[15] Subhash Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *Annual Symposium on Foundations of Computer Science*, number 46, pages 53–62, 2005.

[16] Jean Bernard Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11:796–817, 2001.

[17] Monique Laurent. A comparison of the Sherali-Adams, Lovász-Schrijver and Lasserre relaxations for $0-1$ programming. *Mathematics of Operations Research*, 28(3):470–496, 2003.

[18] László Lovász and Alexander Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.

[19] Ran Raz. A parallel repetition theorem. In *SIAM Journal of Computing*, 27(3):763–803, 1998.

[20] Anup Rao. Parallel Repetition in Projection Games and a Concentration Bound In *Proceedings of the ACM Symposium on the Theory of Computing*, number 40, 2008.

[21] Shmuel Safra and Oded Schwartz. On Parallel-Repetition, Unique-Game and Max-Cut. *Manuscript*, 2007.

[22] Luca Trevisan. Approximation algorithms for unique games. In *Annual Symposium on Foundations of Computer Science*, number 46, pages 197–205, 2005.

APPENDIX

A. DEFERRED PROOFS

A.1 Stronger relaxations

PROOF OF THEOREM 3.2. We consider a strengthening of the UNIQUE GAMES SDP that is obtained by adding constraints similar to the constraints (11), (11), and (12). Let $\mathcal{I}_0$ be the set of all possible subsets of $V \times [k]$ that have cardinality at most r. For every subset $I \in \mathcal{I}_0$, we introduce a new variable $\mathbf{y}_I$ to the SDP in Figure 1. We include the following additional constraints.

$$\forall\, I, J \in \mathcal{I}_0,\ \forall\, I', J' \in \mathcal{I}_0 \qquad \langle \mathbf{y}_I, \mathbf{y}_J \rangle = \langle \mathbf{y}_{I'}, \mathbf{y}_{J'} \rangle \quad \text{if } I \cup J = I' \cup J' \tag{17}$$

$$\forall\, u \in V,\ \forall\, i \in [k] \qquad \mathbf{y}_{(u,i)} = \tfrac{1}{\sqrt{k}} \mathbf{u}_i \tag{18}$$

$$\|\mathbf{y}_\emptyset\|^2 = 1 \tag{19}$$

The resulting SDP is still a relaxation for UNIQUE GAMES, since all additional constraints are valid for the case that all $\mathbf{u}_i$ are either $\mathbf{0}$ or $\sqrt{k}\mathbf{1}$ for some fixed unit vector $\mathbf{1}$. Since $\mathcal{U}$ has ΓMAX2LIN form, we can find a shift-invariant optimal vector solution $\{\mathbf{u}_i\}_{u\in V, i\in[k]}$ to the SDP. Note that an (approximately) optimal solution to the SDP can be computed in time $(kn)^{O(r)}$. Suppose the value of the solution is $1-\eta$. Let $\mathbf{V}_u := \mathbf{u}_{\Lambda_{\text{opt}}(u)}$ for some labeling Λ_{opt} with $\mathsf{val}(\Lambda_{\text{opt}}) \geq 1-\varepsilon$.

We claim that $\mathbf{E}_{u,v\in V} \|\mathbf{V}_u - \mathbf{V}_v\|^2 \leq \frac{1}{z_r(G)}\mathbf{E}_{(u,v)\in E} \|\mathbf{V}_u - \mathbf{V}_v\|^2$. In order to show the claim it is sufficient to show that the vectors $\mathbf{z}_u := \mathbf{V}_u$ can be extended to a solution for the SDP in Figure 2, that is, we need to exhibit vectors $\mathbf{x}_I$ for $I \in \mathcal{I}$ that satisfy the constraints (11) and (11). Note, the last constraint (12) is not essential, since it can always be enforced by scaling the solution by a suitable factor (this does not change the objective value (9)). We can choose the vectors $\mathbf{x}_I$ as $\mathbf{x}_I = \sqrt{k}\mathbf{y}_{\phi(I)}$, where $\phi(I) = \{(u, \Lambda_{\text{opt}}(u)) \mid u \in I\} \in \mathcal{I}_0$.

The claim, together with Lemma 3.1, implies

$$\begin{aligned}\mathbf{E}_{u,v\in V}[\,\rho(u,v)\,] &\leq \tfrac{1}{2}\mathbf{E}_{u,v\in V} \|\mathbf{V}_u - \mathbf{V}_v\|^2 \\ &\leq \tfrac{1}{2z_r(G)}\mathbf{E}_{(u,v)\in E} \|\mathbf{V}_u - \mathbf{V}_v\|^2 \\ &\leq (\eta + 2\varepsilon)/z_r(G).\end{aligned}$$

By Lemma 2.1, the rounding procedure $\mathcal{R}$ from Section 2.2 allows us to compute a labeling Λ such that $\mathsf{val}(\Lambda) \geq 1 - 6\eta - 12\mathbf{E}_{u,v\in V}[\,\rho(u,v)\,] \geq 1 - 50\varepsilon/z_r(G)$. Here it is important to note that the rounding procedure $\mathcal{R}$ does not depend on the vectors $\mathbf{V}_u$. The existence of these vectors is enough to conclude that the rounding procedure succeeds. □

A.2 Parallel Repetition

Proof of Lemma 4.2. Let $\{\mathbf{u}_i\}_{u\in V\cup W, i\in[k]}$ be an optimal solution to the SDP in Figure 3. We have $\mathbf{E}_{(v,w)\in E}\mathbf{E}_{i\in[k]} \langle \mathbf{v}_i, \mathbf{w}_{\pi_{vw}(i)} \rangle = 1-\eta$.

We first show how to obtain a set of vectors $\{\mathbf{u}'_i\}_{u\in V\cup W, i\in[k]}$ that satisfies constraints (2) and (3) and has objective value $\mathbf{E}_{(v,w)\in E}\mathbf{E}_{i\in[k]}\langle \mathbf{v}'_i, \mathbf{w}'_{\pi_{vw}(i)}\rangle = 1-\eta$. Let $s_u = 1 - \frac{1}{k}\sum_{i,j} |\langle \mathbf{u}_i, \mathbf{u}_j\rangle| \geq 0$. For every vertex $u \in V$ and label $i \in [k]$, let $\mathbf{f}_{u,i} \in \mathbb{R}^{[k]\times[k]}$ be the vector defined by

$$\mathbf{f}_{u,i}(i,i) = \sqrt{s_u}, \qquad \forall\, j, j' \neq i.\ \mathbf{f}_{u,i}(j,j') = 0,$$
$$\forall\, j \neq i.\ \mathbf{f}_{u,i}(i,j) = \tfrac{1}{\sqrt{2}}|\langle \mathbf{u}_i, \mathbf{u}_j\rangle|^{1/2},$$
$$\forall\, j \neq i.\ \mathbf{f}_{u,i}(j,i) = -\tfrac{1}{\sqrt{2}}\langle \mathbf{u}_i, \mathbf{u}_j\rangle / |\langle \mathbf{u}_i, \mathbf{u}_j\rangle|^{1/2}.$$

Note that $\langle \mathbf{f}_{u,i}, \mathbf{f}_{u,j}\rangle = -\langle \mathbf{u}_i, \mathbf{u}_j\rangle$ for $i \neq j$, and $\langle \mathbf{f}_{u,i}, \mathbf{f}_{u,i}\rangle = s_u + \sum_{j\neq i} |\langle \mathbf{u}_i, \mathbf{u}_j\rangle|$. Now let $\{\mathbf{x}_{u,i}\}_{u\in V, i\in[k]}$ be a set of vectors such that $\langle \mathbf{x}_{u,i}, \mathbf{x}_{u,j}\rangle = \langle \mathbf{f}_{u,i}, \mathbf{f}_{u,j}\rangle$ but $\langle \mathbf{x}_{u,i}, \mathbf{x}_{v,j}\rangle = 0$ for $u \neq v$. Let $\{\mathbf{u}'_i\}_{u\in V, i\in[k]}$ be the set of vectors with $\mathbf{u}'_i = \mathbf{u}_i \oplus \mathbf{x}_{u,i}$. The value of the objective function has not changed, since $\langle \mathbf{u}'_i, \mathbf{v}'_j\rangle = \langle \mathbf{u}_i, \mathbf{v}_j\rangle$ for distinct vertices u and v. On the other hand, vectors $\langle \mathbf{u}'_i, \mathbf{u}'_j\rangle = \langle \mathbf{u}_i, \mathbf{u}_j\rangle + \langle \mathbf{f}_{u,i}, \mathbf{f}_{u,j}\rangle = 0$ for $i \neq j$, and $\sum_{i=1}^k \|\mathbf{u}'_i\|^2 = \sum_{i=1}^k \|\mathbf{u}_i\|^2 + \|\mathbf{f}_{u,i}\|^2 = k$. Hence the vectors $\mathbf{u}'_i$ satisfy the constraints (2) and (3).

It remains to show how to obtain a set of vectors that satisfies the non-negativity constraint (4) and that has objective value at least $1-2\eta$. By the previous paragraph, we can assume that the vectors $\mathbf{u}_i$ already satisfy the constraints (2) and (3). Now consider the set of vectors $\{\tilde{\mathbf{u}}_i\}_{u\in V\cup W, i\in[k]}$ with $\tilde{\mathbf{u}}_i = \|\mathbf{u}_i\|\,\overline{\mathbf{u}}_i^{\otimes 2}$. The vectors $\tilde{\mathbf{u}}_i$ still satisfy the constraints (2) and (3). They also satisfy constraint (4), because $\langle \tilde{\mathbf{u}}_i, \tilde{\mathbf{v}}_j\rangle = \|\mathbf{u}_i\|\,\|\mathbf{v}_j\|\,\langle \overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j\rangle^2 \geq 0$. We can use the same reasoning as in the proof of Claim 2.8 to show that the objective value $\mathbf{E}_{(v,w)\in E}\mathbf{E}_{i\in[k]}\langle \tilde{\mathbf{v}}_i, \tilde{\mathbf{w}}_{\pi_{vw}(i)}\rangle \geq 1-2\eta$. Specifically, this lower bound follows from the fact $(\mathbf{E}_{i\in[k]}\langle \mathbf{v}_i, \mathbf{w}_{\pi_{vw}(i)}\rangle)^2 \leq \mathbf{E}_{i\in[k]}\langle \tilde{\mathbf{v}}_i, \tilde{\mathbf{w}}_{\pi_{vw}(i)}\rangle$. □

A.3 Integrality gap and UGC hardness

Proof of Theorem 2.5. (Sketch) In [15] it is shown that for every $\eta > 0$ small enough and for every n large enough, there is a Unique Games instance $\mathcal{U}_\eta$ on $k = \Theta(\log(n))$ labels and a constraint graph with n vertices, such that (1) $\mathsf{opt}(\mathcal{U}_\eta) \leq O(1/\log^\eta n)$, and (2) there is an SDP solution for $\mathcal{U}_\eta$ of value at least $1 - O(\eta)$. The instances $\mathcal{U}_\eta$ constructed in [15] are regular. We normalize the weights of the constraints such that the weighted degree of every vertex is equal to 1. Now consider the instance $\mathcal{U}'_\eta$ obtained by adding k additional constraints between every pair of vertices. Every new constraint has weight $\eta/k(n-1)$. For every vertex u, the new constraints contribute weight η to the weighted degree of u. For every pair (u,v) of vertices, we choose the permutations for the new constraints in such a way that every assignment satisfies exactly one of the new constraints between (u,v). For example, we can use the k cyclic shifts as permutations. Notice that the SDP solution for $\mathcal{U}_\eta$ still has value at least $1 - O(\eta)$ for our new instance $\mathcal{U}'_\eta$. Also, $\mathsf{opt}(\mathcal{U}'_\eta) \leq 1/\log^\eta n + \eta/k = O(\log^\eta n)$. On the other hand, the normalized eigenvalue gap λ of the constraint graph of $\mathcal{U}'_\eta$ is $\Omega(\eta)$. This can be seen from the variational characterization (7) of λ. The nominator is always at least an $\Omega(\eta)$ fraction of the denominator. □

Proof of Theorem 2.6. (Sketch) We can use essentially the same reduction as in the above proof of Theorem 2.5. Let $\varepsilon > 0$ be a small enough constant. The UGC asserts that there exists a $k = k(\varepsilon)$ such that it is NP-hard to distinguish between Unique Games instances of value at least $1-\varepsilon$ and instances of value less than ε. We show a reduction from this problem to the problem of distinguishing between Unique Games instances with value at least $1-2\eta$ and instances of value less than 2ε, under the additional promise of $\lambda \geq \eta/(1+\eta)$. Let $\mathcal{U}$ be a Unique Games instances on n vertices. By Lemma 1.5 in [5] we can assume that every vertex in $\mathcal{U}$ has weighted degree 1. As before, we obtain a new Unique Games instance $\mathcal{U}'$ by adding k new constraints between every pair of vertices, each of weight $\eta/k(n-1)$. We can choose the new constraints in such a way that only a $1/k$ fraction of them can be satisfied simultaneously. Notice that U' has $\lambda \geq \eta/(1+\eta)$. This follows from the characterization (7), since a $\eta/(1+\eta)$ fraction of the edges in $\mathcal{U}'$ form a complete graph. On the other hand, the reduction does not change soundness and completeness by too much. If $\mathcal{U}$ has value less than ε, then $\mathcal{U}'$ has value less than 2ε. If $\mathcal{U}$ has value at least $1-\varepsilon$, then $\mathcal{U}'$ has value at least $1-2\eta$. □