Electronic Colloquium on Computational Complexity, Report No. 41 (2010)

Improved Algorithms for Unique Games via Divide and Conquer

Sanjeev Arora* Russell Impagliazzo† William Matthews‡ David Steurer§

March 1, 2010

Abstract

We present two new approximation algorithms for Unique Games. The first generalizes the results of [2, 15] who give polynomial time approximation algorithms for graphs with high conductance. We give a polynomial time algorithm assuming only good *local* conductance, i.e. high conductance for small subgraphs. The second algorithm runs in mildly exponential time, $e^{\alpha n}$, but makes no assumptions about the underlying constraint graph. As the completeness approaches 1 (completeness $1 - \epsilon$), the constant α in the running time rapidly approaches 0 ($\alpha = \exp(-\Omega(1/\epsilon))$.) The value of the solutions returned by these algorithms depend only on the completeness of the Unique Game and either the local conductance or the allowed running time respectively. In particular, the performance of these algorithms does not depend on the number of labels in the Unique Game.

Both algorithms are based on new methods for partitioning graphs by cutting small fractions of edges when the graph can be embedded in a suitable metric space.

*Work supported by NSF grant CCF-0832797

†Work supported by the Simonyi Fund, The Bell Company Fellowship and Fund for Math, and NSF grants DMS-083573 and CCF-0832797. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

‡This research is supported by NSF grant CCF-0947262 from the Division of Computing and Communication Foundations. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

§Work supported by NSF grant CCF-0832797

ISSN 1433-8092

1 Introduction

Khot's unique game conjecture implies that for many optimization problems, the current best approximation algorithms cannot be improved if $P \neq NP$ ([12, 5, 13]). In addition to these important implications, the conjecture has become even more intriguing in the last couple of years because of its seeming close connection to semidefinite programming (and thus to high-dimensional geometry): finite-sized problem instances that have high integrality gaps for the SDP relaxation are used as gadgets in the reductions that show that the SDP relaxation is the best possible algorithm (assuming UGC is true) [16, 14]. Very recently [17], another close connection has emerged between UGC and detecting the local expansion of graphs (more specifically, the expansion profile of sets of different sizes).

Needless to say, the best algorithms for solving unique games also involve SDPs [12, 6, 19, 4, 2, 15]. Several of these algorithms are the best possible with respect to some of the problem parameters – improving upon them would refute UGC – and some algorithms refute stronger versions of the conjecture that had been earlier mentioned as plausible hypotheses. These algorithmic results can be viewed as a kind of stress test of the UGC, one which the original unique games conjecture appears to have passed.

However, UGC has yet to pass another important stress test: nobody has yet identified plausible families of unique games that are hard to solve. (The above-mentioned connection between UGC and local expansion [17] does not help because of a circularity: plausible hard instances for graph expansion problems are based upon unique games, and vice versa!) The general intuition for constraint satisfaction problems is that hard instances should involve a constraint graph that is strongly connected, so that a simple divide-and-conquer type algorithm fails. (Indeed, if the constraint graph is weakly connected, e.g. because it has low tree width, then constraint satisfaction problems become easy.) In fact, for many CSPs such as 3SAT, random sparse instances – which are of course highly connected – are believed to be quite hard. However, in case of unique games, the paper [2] turns this general intuition on its head by giving an algorithm that solves unique games whenever the constraint graph is a sufficiently strong expander (see also [15]). In particular, this shows that random graphs do not give hard instances for unique games. Thus, both extremes (e.g., graphs of low tree-width and strong expanders) are easy for unique games. Is there a middle ground of hard instances?

In this paper, we give two algorithmic results that further restrict and clarify this middle ground of hard instances. Our first algorithm implies that not only can hard instances not be random, they cannot be modified random graphs (so-called *semi-random graphs*) or be composed of large random-like sub-graphs. This algorithm is a generalization of the result of [2] from graphs with high conductance to graphs with high *local* conductance. A graph has high conductance if every cut has a number of cross edges at least a large proportion of the total degrees of the nodes in the smaller side. A graph has high α-local conductance if this is true when the smaller total degree of a side is at most an α fraction of all edges, i.e., unbalanced cuts have good conductance. A graph with large local conductance could be a disconnected or loosely connected union of smaller highly connected graphs, so large local conductance – though implied by large conductance – may be a significantly weaker condition than large conductance. Our algorithm stops working in polynomial time when the conductance falls below a certain threshold, and this threshold appears in two other recent manuscripts. The above-mentioned paper [17] conjectures that the unique game conjecture holds for graphs whose local conductance is "significant" but just a bit below our threshold. In addition, [17] independently gives a different algorithm for unique game on graphs with local conductance at least the same threshold. Another recent paper [18] gives algorithms for approximating local conductance, but it also stops working below our threshold. Mere coincidence?

While the exact trade-off in our algorithm is complicated, when a graph has constant α-local expansion, and the true optimal solution satisfies all but an ϵ of constraints, we find a solution that satisfies a constant fraction of constraints when $\epsilon \leq O\left(1/\log\frac{1}{\alpha}\right)$. We also show that this algorithm is robust under adding or subtracting a small constant fraction of the edges from a locally expanding graph. In particular, we show that *semi-random* graphs, where a random graph is modified by an adversary who may add or remove a constant fraction of edges, are also easy instances of unique games.

Our second algorithm solves an arbitrary unique game with optimal solution value $1-\epsilon$, in expected time $2^{\alpha n} poly(n)$ where $\alpha = \exp(-\Omega(1/\epsilon))$. For the unique games conjecture to hold, unique games should be hard for an arbitrarily small constant ϵ. This result shows that the complexity starts becoming mildly exponential as ϵ approaches 0, and turns polynomial when $\epsilon = O(1/\log n)$, essentially matching the threshold for easiness of Trevisan's algorithm in [19]. Our first algorithm has roughly the same threshold as well, since for $\alpha = 1/n$, every regular graph has conductance 1. So in some sense, we have independent confirmation that $1 - O(1/\log n)$ might be the critical threshold for completeness of hard unique games. (Since some NP-complete languages like 3SAT are believed to require strict exponential time, our result also suggests lower bounds on the efficiency of any putative reduction from 3SAT to UGC with completeness $1-\epsilon$. Namely, the reduction must increase instance size by at least a factor $\exp(\Omega(1/\epsilon))$.) Very recently, and subsequently to our work, [1] improve on our mildly exponential algorithm and give one that runs in time $\exp(n^{poly(\epsilon)})$.

Both of our algorithms are based on a divide-and-conquer strategy, utilized in conjunction with the rounding method for SDPs from [2]. We first quickly summarize their algorithm.

As most other unique games algorithms do, we first find a solution to the SDP solution. This maps each variable and value for that variable to a vector in $\mathbb{R}^{nk}$. However, [2] show how to combine the k vectors for a single variable into a single vector without losing much information. They give a rounding method where the fraction of violated constraints is roughly a function of the average distances between pairs of vectors. The quality of the SDP solution implies that average distances of pairs connected by edges is small. They then use the well-known fact that every labeling of a graph with good conductance by points in Euclidean space with small average distances on the edges also has most pairs of nodes close together. They conclude that their rounding method works well on graphs with good conductance.

However, if the graph is not globally an expander then the above analysis breaks down since these graphs need not even be connected. Our modification of this approach is to essentially use the fact that if the rounding method *fails* globally then the vectors must be clustered. Specifically, we show how to *partition* the high-dimensional sphere into regions where either distances between vectors in the region are small (call these regions with "low average radius") or the number of vectors is at most about the expected number (call these *small* regions). The partitioning ensures that only a small fraction of all constraints go between the regions, so we can ignore these. In each of the low-radius regions, the rounding algorithm will give a good solution for the variables in that region. We show that for local expanders, most variables must be in low-radius regions. For arbitrary graphs, this is not the case. However, since each region has a relatively small fraction of variables, we can use an exponential-time algorithm on each such region. (The existence of an exponential-time algorithm is slightly non-trivial, since we need it to be $O(2^n)$ time, not $O(k^n)$; see

Section 2.5.)

1.1 Definitions and results

Definition 1.1 (Unique Game)**.** A *unique game* consists of an undirected graph $G = (V, E)$, a set $[k]$ of k labels, and a set of bijective constraints $\pi_{uv} : [k] \to [k]$, one for each edge $\{u, v\}$. The goal of the game is to find a labeling $L : V \to [k]$ which maximizes the fraction of satisfied edges. An edge $\{u, v\}$ is satisfied by L when $L(v) = \pi_{uv}(L(u))$. The *value* of a labeling L is the fraction of edges satisfied by L, and the value of the unique game is the maximum value over all possible labelings.

Conjecture 1.1 (Unique Games Conjecture [12])**.** *For any arbitrarily small constants $\epsilon, \delta > 0$, there exists a $k > 0$ such that it is NP-hard to distinguish a unique game with k labels and value at least $1 - \epsilon$ from one with value at most δ.*

Definition 1.2 (Volume)**.** For any graph $G = (V, E)$ and any set $S \subseteq V$, define the volume of the set S as $\mathrm{Vol}(S) = \frac{1}{2|E|} \sum_{v \in S} d_v$, where d_v is the degree of vertex v.

Definition 1.3 (Local Conductance)**.** The local conductance ϕ_α of G on sets up to volume α is defined as

$$\phi_\alpha = \min_{S \subseteq V, \mathrm{Vol}(S) \leq \alpha} \frac{|E(S, \overline{S})|}{|E| \, \mathrm{Vol}(S)}$$

Any graph with local conductance ϕ_α on sets up to volume α is said to be an (α, ϕ_α)-locally conducting graph.

On a d-regular constraint graph, $\mathrm{Vol}(S) = |S|/|V|$ and ϕ_α is the edge expansion sets of size at most αn divided by $d/2$.

Theorem 1.1. *There exist a constants $0 < c_1, c_2 < 1$ such that given an (α, ϕ_α)-locally conducting Unique Game with value at least $1 - \epsilon$, where $\epsilon/\phi_\alpha^2 < c_1$ and $\alpha > e^{-c_2 \phi_\alpha^2/\epsilon}$, we can find a labeling with value at least*

$$1 - O\left(\sqrt[3]{R \ln \frac{1}{R}}\right) \qquad \text{where} \qquad R = \frac{\epsilon}{\phi_\alpha^2} \ln \frac{\phi_\alpha}{\alpha \epsilon}$$

in probabilistic polynomial time.

Corollary 1.1. *Given an (α, ϕ_α)-locally conducting Unique Game with value at least $1 - \epsilon$, we can find a labeling of value at least $1/2$ when $\phi_\alpha \geq \Omega\sqrt{\epsilon \ln \frac{1}{\alpha}}$.*

Theorem 1.2. *There exists a constant c such that given a Unique Game with value at least $1 - \epsilon$ and parameter $\alpha > e^{-c/\epsilon}$, in expected time $2^{\alpha n} poly(n)$ we can find a labeling with value*

$$1 - O\left(\sqrt[3]{R \ln \frac{1}{R}}\right) \qquad \text{where} \qquad R = \epsilon \ln \frac{1}{\alpha \epsilon}$$

Corollary 1.2. *Given a Unique Game with value at least $1 - \epsilon$ for all sufficiently small ϵ, we can find a labeling of value at least $1/2$ in expected time $2^{2^{-\Omega(1/\epsilon)} n}$.*

Definition 1.4 (Almost Local Conductance)**.** A graph $G = (V, E)$ is $(\alpha, \phi_\alpha, \gamma^+, \gamma^-)$-almost locally conducting if there exists an (α, ϕ_α)-locally conducting graph $G_0 = (V, E_0)$ such that $|E \setminus E_0| \le \gamma^+|E_0|$ and $|E_0 \setminus E| \le \gamma^-|E_0|$.

Models for semi-random graphs have been considered for maximum independent set [8, 9], k-coloring [3, 9], and other similar problems [9, 7]. We consider the following model for semi-random unique games and claim that the resulting graph is $(\alpha, \phi_\alpha, \gamma^+, \gamma^-)$-almost locally conducting.

1. Allow an adversary to choose an (α, ϕ_α)-locally conducting graph $G_0 = (V, E_0)$ (or begin with a random graph: roughly $(\alpha, (1-\alpha))$-locally conducting)
2. Allow an adversary to add any $\gamma^+|E_0|$ edges.
3. Allow an adversary to remove any $\gamma^-|E_0|$ edges.
4. Allow an adversary to choose constraints for all of the edges such that the game has value $1-\epsilon$.

Theorem 1.3. *There exist a constants $0 < c_1, c_2 < 1$ such that given an $(\alpha, \phi_\alpha, \gamma^+, \gamma^-)$-locally conducting Unique Game with value at least $1-\epsilon$, where $\epsilon/\phi_\alpha 2 < c_1$ and $\alpha > e^{-c_2\phi_\alpha 2/\epsilon}$, we can find a labeling with value at least*

$$1 - O\left(\sqrt[3]{R \ln \frac{1}{R}}\right) - \left(\frac{\gamma^-}{1-\gamma^-}\right)\left(1 + \frac{2}{\phi_\alpha}\right) - \left(\frac{\gamma^+}{1-\gamma^-}\right) \qquad \textit{where} \qquad R = \frac{\epsilon}{\phi_\alpha 2} \ln \frac{\phi_\alpha}{\alpha\epsilon}$$

in probabilistic polynomial time.

Since a random degree d graph has conductance $\phi_{1/2} \ge 1/2 - O\left(1/\sqrt{d}\right)$ [10] with probability $1 - n^{-\Omega(\sqrt{d})}$, allowing an adversary to modify such a graph gives the following corollary.

Corollary 1.3. *Given a random degree $d \ge 3$ graph $G_0 = (V, E_0)$, we allow an adversary to add and remove $\gamma^+ \le 1/2$ and $\gamma^- \le 1/2$ fraction of edges respectively and then choose constraints for all edges such that the resulting game has value at least $1-\epsilon$, then we can find a solution of value at least*

$$1 - O\left(\sqrt[3]{R \ln \frac{1}{R}} + \gamma^- + \gamma^+\right) \qquad \textit{where} \qquad R = \frac{\epsilon}{1-\sqrt{d}} \ln \frac{1}{\epsilon}$$

in probabilistic polynomial time for all but $n^{-\Omega(\sqrt{d})}$ probability in the choice of G_0.

2 Algorithm

Consider a unique game on a graph $G = (V, E)$ with constraints $\pi_{uv} : [k] \to [k]$ for each edge $\{u, v\} \in E$. We assume that the game has value at least $1-\epsilon$.

All of our algorithms for unique games follow the same general outline, given in Figure 1. We start by using the SDP solution to define a certain distance function ρ on the graph vertices due to Arora et al. [2], which can be viewed both as some kind of earth-mover distance, as well as a geometric ℓ_2^2 distance up to small distortion. Our divide-and-conquer approach tries to partition the graph into low-diameter regions with respect to ρ, and then apply the algorithm of Arora et

1. Run the standard SDP to get vectors $\{\mathbf{v_i}\}_{v\in V, i\in[k]}$. Use these to define a distance function, ρ, between vertices. (Section 2.1, due to [2])
2. Embed ρ into ℓ_2^2 to get a set of vectors $\{\mathbf{V_v}\}_{v\in V}$. (Section 2.2, due to [2])
3. Using $\{\mathbf{V_v}\}_{v\in V}$, partition V into: (Section 2.3, inspired by [6, 11])
    - A small set of outliers
    - "Low diameter" sets, where almost all pairs of vertices are close in terms of ρ
    - "Small" sets, each of which only contains a few vertices
4. Use propagation rounding for each "low diameter" set. (Section 2.4, due to [2])
5. For each "small" set do one of the following:
    - In the general case, run an exponential time essentially brute force algorithm (Sections 2.5, 3.5)
    - When the graph has good local conductance, just ignore the "small" sets (Section 3.4)

Figure 1: High level algorithm overview.

al. [2] on each region. Such a low-diameter decomposition doesn't exist in general metrics with the kind of parameters we are in interested in, where the diameter of each region is only a constant factor smaller than the overall diameter. However, since ρ has a geometric realization as ℓ_2^2 we can construct an *almost* low-diameter decomposition, where the average diameter is small once we throw away some outliers. This may be of independent interest.

Thus Steps 3 and 5 are the new contributions. Steps 1, 2, and 4 closely follow Arora et al. [2]. They argue that if the value of the SDP is at least $1-\epsilon$ then the average distance ρ along edges is at most ϵ. They then argue, using their embedding of ρ into ℓ_2^2, that if the underlying graph is an expander ($\lambda \gg 0$) then the average ρ distance between arbitrary pairs of vertices is approximately ϵ/λ. Finally they show that if the average distance between pairs of vertices is Δ, then they can find a labeling of the graph of value $1 - O(\Delta)$.

Since we do not assume that the graph is an expander, we cannot guaranteed that the average distance between arbitrary pairs of vertices is small. Instead, we partition the graph into sets and show that almost all vertices are in sets that are either "small" or have "low average radius" The "low average radius" sets can be labeled using the same approach as Arora et al. When we allow either mildly exponential time or assume the graph has good local conductance, we can handle the "small" sets.

We will now formalize these notions and then give the details of the algorithm in Section 2. Unless otherwise noted, throughout the paper random edges are chosen uniformly and random vertices are chosen with probability proportional to their degree.

Definition 2.1 (Small Sets)**.** We say that a set S is α-small if $\mathrm{Vol}(S) \le \alpha$.

Definition 2.2 (Low-radius Sets)**.** We say that a set S has *average radius* at most Δ if $\exists v \in S$

$$\text{Maximize} \quad \mathop{\mathrm{E}}_{\{u,v\}\in E}\ \mathop{\mathrm{E}}_{i\in[k]}\left[\langle \mathbf{u_i}, \mathbf{v}_{\pi_{\mathbf{uv}}(\mathbf{i})}\rangle\right] \tag{2.1}$$

Subject to

$$\forall u \in V \qquad \mathop{\mathrm{E}}_{i\in[k]}\left[\|\mathbf{u_i}\|^2\right] = 1 \tag{2.2}$$

$$\forall u \in V \quad \forall i \neq j \qquad \langle \mathbf{u_i}, \mathbf{u_j}\rangle = 0 \tag{2.3}$$

$$\forall u, v \in V \quad \forall i, j \qquad \langle \mathbf{u_i}, \mathbf{v_j}\rangle \geq 0 \tag{2.4}$$

Figure 2: The standard semidefinite program for unique games.

such that $\mathrm{E}_{u\in S}\left[\rho(u,v)\right] \leq \Delta$.

2.1 SDP and ρ distance

We begin with the standard SDP relaxation given in Figure 2. For every vertex $u \in V$, we associate a set of k vectors $\{\mathbf{u_1}, \ldots, \mathbf{u_k}\}$. Intuitively, assigning label i to vertex u could correspond to setting $\mathbf{u_i} = \sqrt{k}\mathbf{1}$ and $\mathbf{u_j} = \mathbf{0}$ for all $j \neq i$. Here $\mathbf{1}$ is an arbitrary unit vector, and $\mathbf{0}$ is the zero vector. If this intuition is followed, the value of the SDP objective function corresponds exactly to the value of the assignment for the unique game. However in general, the solution to the SDP will just be some set of vectors with optimal value satisfying the constraints. Thus, since we assume the value of the game is at least $1-\epsilon$, the value of the SDP objective function must be at least $1-\epsilon$.

Following [2], for every pair of vertices u, v, let $\sigma_{uv} : [k] \to [k]$ be a bijection that maximizes $\mathrm{E}_{i\in[k]}\left[\langle \mathbf{u_i}, \mathbf{v}_{\sigma_{\mathbf{uv}}(\mathbf{i})}\rangle\right]$. Define the distance $\rho(u,v)$ between a pair of vertices as

$$\rho(u,v) = \frac{1}{2}\mathop{\mathrm{E}}_{i\in[k]}\left[\|\mathbf{u_i} - \mathbf{v}_{\sigma_{\mathbf{uv}}(\mathbf{i})}\|^2\right] = 1 - \mathop{\mathrm{E}}_{i\in[k]}\left[\langle \mathbf{u_i}, \mathbf{v}_{\sigma_{\mathbf{uv}}(\mathbf{i})}\rangle\right] \tag{2.5}$$

This is a form of earth-mover distance.

Let $\epsilon_{uv} = \frac{1}{2}\mathrm{E}_{i\in[k]}\left[\|\mathbf{u_i} - \mathbf{v}_{\pi_{\mathbf{uv}}(\mathbf{i})}\|^2\right] = 1 - \mathrm{E}_{i\in[k]}\left[\langle \mathbf{u_i}, \mathbf{v}_{\pi_{\mathbf{uv}}(\mathbf{i})}\rangle\right]$. From the SDP and the assumption that the unique game has value $1-\epsilon$, we get $\mathrm{E}_{\{u,v\}\in E}\left[\epsilon_{uv}\right] \leq \epsilon$.

2.2 Embedding ρ into ℓ_2^2

The following lemma of Arora at al. [2] gives us a vector for each vertex where the ℓ_2^2 distance between each pair of vertices is approximately the ρ distance.

Lemma 2.1 ([2, Lemma 2.1]). *For every positive even integer t and every SDP solution, there exists a set of unit vectors $\{\mathbf{V_u}\}_{u\in V}$ such that for every pair u, v of vertices*

$$\frac{1}{2t}\|\mathbf{V_u} - \mathbf{V_v}\|^2 \leq \rho(u,v) \leq \|\mathbf{V_u} - \mathbf{V_v}\|^2 + O\left(2^{-t/2}\right)$$

Moreover, these vectors and be constructed in polynomial time from the SDP solution.

2.3 Partitioning via cones

The geometric partitioning used involves origin-centered cones of a certain volume whose central axis is uniformly distributed. A similar idea was used in the graph coloring algorithm of [11], and also in the UG approximation algorithm of [6].

In other words, given a graph $G = (V, E)$ and the set of vectors $\{\mathbf{V_v}\}_{v\in V}$ from the previous section, and parameter α, we will choose a threshold s approximately $\sqrt{\log 1/\alpha}$ and define a distribution D on subsets $X = \{v \in V \mid \langle \mathbf{V_v}, \mathbf{g}\rangle \geq s\}$ of vertices where $\mathbf{g}$ is a Gaussian random vector.

This distribution satisfies two properties, essentially proved in [6, 11]. There exists a constant c independent of s such that for any pair of vertices u and v the following hold when X is chosen according to the distribution:

$$\Pr\left[u \notin X \mid v \in X\right] \leq O\left(\|\mathbf{V_u} - \mathbf{V_v}\| s\right) \tag{2.6}$$

$$\Pr\left[u \in X \mid v \in X\right] \leq e^{-c\|\mathbf{V_u}-\mathbf{V_v}\|^2 s^2} \tag{2.7}$$

We define the sets in our partition by repeatedly sampling X_i from D and creating a set S_i containing all of the vertices in X_i that are not contained in some previous S_j. We repeat this process until all vertices of V are in some set S_i.

We allow the partitioning to fail in some sense on a small set of vertices. We refer to these vertices as outliers, and bound their total volume. For all vertices other than these outliers, we get good quantitative bounds on their sets. The following lemma quantifies the relevant properties of this partitioning.

Lemma 2.2. *Given a parameters α and β, we can choose the threshold s and define a notion of an "outlier" such that with probability at least $1/3$ the following conditions hold:*

1. *The set Z of outliers satisfies $\mathrm{Vol}(Z) \leq \Delta$.*
2. *For all i, $S_i \setminus Z$ is either α-small or has average radius Δ.*
3. *At most a Δ/β fraction of edges have endpoints in different sets.*

Where $\Delta = O\left(\sqrt[3]{R \ln \frac{1}{R}}\right)$ and $R = \beta^2 \epsilon \ln \frac{1}{\alpha\beta\epsilon}$.

In the general case, we will use Lemma 2.2 with $\beta = 1$. We prove Lemma 2.2 in Section 3.

In the special case where G is (α, ϕ_α)-locally conducting, we use the following stronger lemma:

Lemma 2.3. *Given an (α, ϕ_α)-locally conducting graph, we can choose the threshold s and define a notion of an outlier such that with probability at least $1/3$ the resulting partitioning satisfies the following conditions:*

1. *The set Z of outliers satisfies $\mathrm{Vol}(Z) \leq \Delta$.*
2. *For all i, $S_i \setminus Z$ has average radius Δ.*
3. *At most a Δ fraction of edges have endpoints in different sets.*

Where $\Delta = O\left(\sqrt[3]{R \ln \frac{1}{R}}\right)$ and $R = \frac{\epsilon}{\phi_\alpha^2} \ln \frac{\phi_\alpha}{\alpha\epsilon}$.

We prove Lemma 2.3 in Section 3.1.

If the graph G is $(\alpha, \phi_\alpha, \gamma^+, \gamma^-)$-almost locally conducting, we use the following variant of the locally conducting lemma:

Lemma 2.4. *Given an $(\alpha, \phi_\alpha, \gamma^+, \gamma^-)$-almost locally conducting graph, we can choose the threshold s and define a notion of an "outlier" such that with probability at least $1/3$ the resulting partitioning satisfies the following conditions:*

1. *The set Z of outliers satisfies $\mathrm{Vol}\,(Z) \leq O\,(\Delta) + \left(\frac{\gamma^-}{1-\gamma^-}\right)\left(1 + \frac{2}{\phi_\alpha}\right) + \left(\frac{\gamma^+}{1-\gamma^-}\right)$.*
2. *For all i, $S_i \setminus Z$ has average radius Δ.*
3. *At most a Δ fraction of edges have endpoints in different sets.*

Where $\Delta = O\left(\sqrt[3]{R \ln \frac{1}{R}}\right)$ and $R = \frac{\epsilon}{\phi_\alpha^2} \ln \frac{\phi_\alpha}{\alpha\epsilon}$.

We prove Lemma 2.4 in Section 3.2. Other than the additional outliers which are ignored, we treat almost locally conducting unique games the same as locally conducting unique games.

2.4 Propagation Rounding

The following Lemma follows directly from Arora et al. [2]. Rather than applying this lemma to the entire graph as [2] does, we apply it to each S_i that has average radius Δ.

Lemma 2.5 ([2])**.** *Given a unique game $G = (V, E)$ with value $1 - \epsilon$, vertex $v \in V$ and $\Delta > 0$ satisfying $\mathrm{E}_{u \in V}\,[\rho(u, v)] \leq \Delta$, we can find a labeling for G with value at least $1 - O\,(\Delta + \epsilon)$.*

Since the statement of this Lemma differs slightly from the statement in [2, Lemma 2.1], we give a brief proof in Section 3.3.

2.5 Exponential time algorithm with no dependence on number of labels

A unique game instance with n nodes and k labels appears at first glance to be a CSP that has $n \log k$ "nondeterministic bits", and therefore a trivial exhaustive search would require k^n time. We observe that one can do the exhaustive search more cleverly in $\exp(n)$ expected time. This does not seem to have been noticed before.

This algorithm will be applied to each α-small set S_i.

Lemma 2.6. *Given a unique game G on n vertices and m edges with value $1 - \epsilon$ and parameter $z > 0$, we can find a labeling with value at least $1 - (\epsilon + z)$ in expected time $(1 - \frac{\epsilon}{z})^{-n} poly(n)$.*

Proof. Consider an optimal labeling for G. At most ϵm edges are unsatisfied. We will call these edges "bad" edges and all other edges "good" edges.

We will randomly construct a forest F in G. While there are more than zm edges between connected components in F, choose one uniformly at random and add it to F. At each step, we add a "bad" edge with probability at most ϵ/z. We add at most $n-1$ edges to F, so with probability at least $(1 - \epsilon/z)^n$ we never add a "bad" edge.

If F consists of only "good" edges, we can easily find a labeling with value at least $1 - (\epsilon + z)$: There are zm edges between connected components, and at most all of these are unsatisfied. For each connected component, choose a vertex and then try all possible labels for that vertex. For each label, label the rest of the component consistently with the edges in F. Keep the labeling that satisfies the most edges. Since we have no "bad" edges in F, we have at most ϵm edges unsatisfied within connected components.

We repeat this process until we find a labeling of value at least $1-(\epsilon+z)$. The expected number of times we need to repeat it is at most $(1-\frac{\epsilon}{z})^{-n}$. $\square$

3 Partitioning Proofs

Proof of Lemma 2.2. Given a graph $G=(V,E)$, vectors $\{\mathbf{V_v}\}_{v\in V}$, and parameter $\alpha>0$, let $\eta = \mathrm{E}_{(u,v)\in E}\left[\|\mathbf{V_u}-\mathbf{V_v}\|^2\right]$. Let Δ' and s be parameters of ϵ and α to be determined. Let $S(v)$ refer to the set S_i containing v.

We will use the distribution on vertices where each is chosen with probability proportional to its degree, and we will use the uniform distribution on edges. Thus, choosing a random vertex and then an random edge adjacent to it is equivalent to choosing a random edge, and for any set $S\subseteq V$, $\Pr_v[v\in S]=\mathrm{Vol}(S)$.

Let i be the minimum index such that X_i contains u or v. Assume that it contains v (otherwise rename u and v).

$$\Pr[S(u)\neq S(v)] = \Pr[u\notin X_i \mid v\in X_i] \le O\left(\|\mathbf{V_u}-\mathbf{V_v}\| s\right)$$
$$\Pr[S(u)= S(v)] = \Pr[u\in X_i \mid v\in X_i] \le e^{-c\|\mathbf{V_u}-\mathbf{V_v}\|^2 s^2}$$

Let $F(v)=\{u\in S(v)\mid \|\mathbf{V_u}-\mathbf{V_v}\|^2>\Delta'\}$ denote the set of vertices that are in v's set, but far from v. $\mathrm{E}[\mathrm{Vol}(F(v))]\le \Pr_u\left[S(u)=S(v)\mid \|\mathbf{V_u}-\mathbf{V_v}\|^2>\Delta'\right]\le e^{-c\Delta' s^2}$

We define the set $Z=\{v\in V\mid \mathrm{Vol}(F(v))>3e^{-c\Delta' t^2}/\Delta'\}$ of "outliers", which are the vertices that are far from many of the vertices in their set. By Markov's inequality, $\Pr[\mathrm{Vol}(Z)>\Delta']<1/3$.

For all $v\in V\setminus Z$, with $\mathrm{Vol}(S(v))\ge\alpha$, we bound the average distance from v: $\Pr_{u\in S(v)}\left[\|\mathbf{V_u}-\mathbf{V_v}\|^2>\Delta'\right]=\mathrm{Vol}(F(v))/\mathrm{Vol}(S(v))\le\frac{3e^{-c\Delta' s^2}}{\alpha\Delta'}$.

Let $\gamma=|\{(u,v)\in E\mid S(u)\neq S(v)\}|/|E|$ be the fraction of edges with endpoints in different sets. $\mathrm{E}[\gamma]\le O\left(\mathrm{E}_{\{u,v\}\in E}\|\mathbf{V_u}-\mathbf{V_v}\|s\right)\le O\left(\sqrt{\eta}s\right)$. Thus, there exists c' such that $\Pr\left[\gamma>c'\sqrt{\eta}s\right]<1/3$.

We wish to choose s,Δ' to balance the fraction of outliers and the fraction of cut edges. Choosing $\Delta'=O\left(\sqrt[3]{\eta\ln\frac{1}{\alpha\eta}}\right)$ and $s=\Delta'/\sqrt{\eta}$ gives that with probability at least $1/3$, both the total volume of outliers Z and the total fraction of cut edges are at most Δ'.

Next, we will convert these bounds in terms of ℓ_2^2 distance to bounds in terms of ρ distance. We choose t to be $2\left\lceil\log_2\frac{1}{\sqrt[3]{\beta^2\epsilon\ln\frac{1}{\alpha\beta\epsilon}}}\right\rceil$ for Lemma 2.1, giving

$$\|\mathbf{V_u}-\mathbf{V_v}\|^2\le O\left(\rho(u,v)\ln\frac{1}{\beta\epsilon\ln\frac{1}{\alpha\beta\epsilon}}\right)$$
$$\rho(u,v)\le\|\mathbf{V_u}-\mathbf{V_v}\|^2+O\left(\sqrt[3]{\beta^2\epsilon\ln\frac{1}{\alpha\beta\epsilon}}\right)\le\|\mathbf{V_u}-\mathbf{V_v}\|^2+O\left(\Delta'\right)$$

Since we assume that the SDP has value at least $1-\epsilon$, we get that $\mathrm{E}_{(u,v)\in E}\,\rho(u,v)\le\epsilon$. Combined with the above bounds, we get $\eta\le O\left(\epsilon\ln\frac{1}{\beta\epsilon}\right)$.

If $\|\mathbf{V_u} - \mathbf{V_v}\|^2 < \Delta'$ then $\rho(u, v) < \Delta$, where $\Delta = \Delta' + O\left(\sqrt[3]{\beta^2\epsilon}\right) = O(\Delta')$. For any v where $\Pr_{u \in S(v)} \|\mathbf{V_u} - \mathbf{V_v}\|^2 > \Delta' < \Delta'$, we get $\Pr_{u \in S(v)} \rho(u, v) > \Delta < \Delta' < \Delta$.

Combining everything, we get $\Delta = O\left(\sqrt[3]{R \ln \frac{1}{R}}\right)$ where $R = \beta^2 \epsilon \ln \frac{1}{\alpha\beta\epsilon}$. □

3.1 Locally Conducting Case

Proof of Lemma 2.3. From Lemma 2.2 with $\beta = 2/\phi_\alpha$ we get sets S_i and a set Z of outliers. In addition, we get that at most $r\phi_\alpha/2$ fraction of edges go between sets.

From the local conductance of G, every α-small set $S = S_i \setminus Z$ has at least $|E|\phi_\alpha \operatorname{Vol}(S)$ edges leaving it. Each edge crossing between partitions may be counted at most twice this way. Since the total number of edges crossing between sets is at most $|E|\Delta\phi_\alpha/2$ the total volume in α-moderate sets S is at most Δ. We add all such sets to Z, increasing its volume by at most Δ. □

3.2 Partitioning Almost Locally Conducting Unique Games

Given a $(\alpha, \phi_\alpha, \gamma^+, \gamma^-)$-almost locally conducting graph, $G = (V, E)$, let $G_0 = (V, E_0)$ denote the underlying (α, ϕ_α)-locally conducting graph. Let $E^+ = E \setminus E_0$ and $E^- = E_0 \setminus E$ denote the sets of edges added and removed respectively when transforming G_0 into G. We are guaranteed that $|E^+| \le \gamma^+|E_0|$ and $|E^-| \le \gamma^-|E_0|$.

Intuitively we would like to argue in much the same was as in the locally conducting case. We would like to say that the total volume of "moderate" sets can be bounded in terms of ϕ_α and the number of cut edges plus the number of edges that we must add back to G to make it locally conducting. This is complicated by the fact that for any set of vertices, S, the volume of S in G and the volume of S in G_0 may be very different.

Proof of Lemma 2.4. We use Lemma 2.2 with $\beta = 1/\phi_\alpha$ and $\hat{\alpha} = \frac{\alpha}{2+\gamma^+}$ to produce sets $S_1, \ldots, S_\ell$, and a set of outliers Z. At most a $\Delta\phi_\alpha$ fraction of edges go between sets.

Let $d_X(v), \operatorname{Vol}_X S$, and $E_X(S, \overline{S})$ denote the degree of v, volume of S, and edges from S to $\overline{S}$ when restricted to X, where X is a graph or a set of edges. When X is a set of edges, restrict to the corresponding graph (V, X).

Claim: For all but a f volume of vertices v, $\operatorname{Vol}_G S(v) < \hat{\alpha} \implies \operatorname{Vol}_{G_0} S(v) \le \alpha$, where $f = \frac{|E^-|}{|E_0|\alpha/\hat{\alpha} - |E|}$.

Let $D = \{S_i \mid \operatorname{Vol}_G S_i < \hat{\alpha} \wedge \operatorname{Vol}_{G_0} S_i > \alpha\}$. The following holds for all $S_i \in D$:

$$\sum_{v \in S_i} d_E(v) < 2|E|\hat{\alpha} \qquad \sum_{v \in S_i} d_{E_0}(v) > 2|E_0|\alpha$$

$$d_E(v) = d_{E_0}(v) + d_{E^+}(v) - d_{E^-}(v)$$

$$\sum_{v \in S_i} d_{E^-}(v) > \sum_{v \in S_i} d_{E^-}(v) - d_{E^+}(v) > 2|E_0|\alpha - 2|E|\hat{\alpha}$$

Now considering all of D

$$
\begin{aligned}
|D|\big(2|E_0|\alpha - 2|E|\hat{\alpha}\big) &< \sum_{S_i \in D} \sum_{v \in S_i} d_{E^-}(v) \leq 2|E^-| \\
|D| &< \frac{|E^-|}{|E_0|\alpha - |E|\hat{\alpha}} \\
\mathrm{Vol}\left(\cup_{S_i \in D} S_i \right) &< |D|\hat{\alpha} < \frac{|E^-|}{|E_0|\alpha/\hat{\alpha} - |E|}
\end{aligned}
$$

We can then bound the total volume of vertices in small sets

$$
\begin{aligned}
\frac{1}{E_0} \sum_{S_i, \mathrm{Vol}_{G_0} S_i \leq \alpha} |E_{G_0}(S_i, \overline{S_i})| &\geq \phi_\alpha \sum_{S_i, \mathrm{Vol}_{G_0} S_i \leq \alpha} \mathrm{Vol}_{G_0} S_i \\
(1+\gamma^+)|E_0| &\geq |E| \geq (1-\gamma^-)|E_0| \\
\#\text{cut edges in } G &\leq |E|\Delta\phi_\alpha/2 \\
\#\text{cut edges in } G_0 &\leq \#\text{cut edges in } G + |E^-| \leq |E^-| + |E|\Delta\phi_\alpha \\
\frac{2(|E^-| + \Delta\phi_\alpha|E|)}{\phi_\alpha|E_0|} &\geq \sum_{S_i, \mathrm{Vol}_{G_0} S_i \leq \alpha} \mathrm{Vol}_{G_0} S_i
\end{aligned}
$$

Combining this with our claim above

$$
\begin{aligned}
\sum_{S_i, \mathrm{Vol}_G S_i < \hat{\alpha}} \mathrm{Vol}_{G_0} S_i &\leq \frac{|E^-|}{|E_0|\alpha/\hat{\alpha} - |E|} + \sum_{S_i, \mathrm{Vol}_{G_0} S_i \leq \alpha} \mathrm{Vol}_{G_0} S_i \\
&\leq \frac{|E^-|}{|E_0|\alpha/\hat{\alpha} - |E|} + \frac{2(|E^-| + \Delta\phi_\alpha|E|)}{\phi_\alpha|E_0|} \\
\text{Let } Z &= \cup_{S_i, \mathrm{Vol}_G S_i \leq \hat{\alpha}} S_i \\
\mathrm{Vol}_{G_0} Z &\leq \frac{|E^-|}{|E_0|\alpha/\hat{\alpha} - |E|} + \frac{2(|E^-| + \Delta\phi_\alpha|E|)}{\phi_\alpha|E_0|} \\
\mathrm{Vol}_G Z &\leq \frac{|E_0| \, \mathrm{Vol}_{E_0} Z + |E^+|}{|E|} \\
&\leq \frac{\frac{|E^-|}{\alpha/\hat{\alpha} - |E|/|E_0|} + 2|E^-|/\phi_\alpha + \Delta|E| + |E^+|}{|E|} \\
&= \frac{|E^-|/|E|}{\alpha/\hat{\alpha} - |E|/|E_0|} + \frac{2|E^-|}{|E|\phi_\alpha} + \Delta + |E^+|/|E| \\
&\leq \Delta + \frac{\gamma^-/(1-\gamma^-)}{\alpha/\hat{\alpha} - (1+\gamma^+)} + \frac{2\gamma^-}{(1-\gamma^-)\phi_\alpha} + \frac{\gamma^+}{1-\gamma^-} \\
&= \Delta + \frac{\gamma^-}{1-\gamma^-}\left(1 + \frac{2}{\phi_\alpha}\right) + \frac{\gamma^+}{1-\gamma^-}
\end{aligned}
$$

$\square$

3.3 Propagation Rounding

Proof of Lemma 2.5. Label vertex v with label i with probability $\frac{1}{k}\|\mathbf{v_i}\|^2$. For every other vertex $u \in V$, assign u label $\sigma_{vu}(i)$.

Each edge $\{u, w\}$ is satisfied when $\pi_{uw} \circ \sigma_{vu}(i) = \sigma_{vw}(i)$, or equivalently when $\sigma_{wv} \circ \pi_{uw} \circ \sigma_{vu}(i) = i$. Let $\hat{\pi} = \sigma_{wv} \circ \pi_{uw} \circ \sigma_{vu}$. The probability that the edge is not satisfied is $\mathrm{E}_{i\in[k]} \|\mathbf{v_i}\|^2 \mathbf{1}_{i\neq\hat{\pi}(i)}$, where $\mathbf{1}_E$ is the indicator random variable for the event E. Since the vectors $\mathbf{v_i}$ and $\mathbf{v_j}$ are orthogonal when $i \neq j$, we get

$$\begin{aligned}
\mathop{\mathrm{E}}_{i\in[k]} \|\mathbf{v_i}\|^2 \mathbf{1}_{i\neq\hat{\pi}(i)} &= \frac{1}{2} \mathop{\mathrm{E}}_{i\in[k]} (\|\mathbf{v_i}\|^2 + \|\mathbf{v_{\hat{\pi}(i)}}\|^2)\mathbf{1}_{i\neq\hat{\pi}(i)} \\
&= \frac{1}{2} \mathop{\mathrm{E}}_{i\in[k]} \|\mathbf{v_i} - \mathbf{v_{\hat{\pi}(i)}}\|^2 \\
&\leq \frac{1}{2} \mathop{\mathrm{E}}_{i\in[k]} \Big[\big(\|\mathbf{v_i} - \mathbf{u_{\sigma_{vu}(i)}}\| \\
&\qquad + \|\mathbf{u_{\sigma_{vu}(i)}} - \mathbf{w_{\pi_{uw}\circ\sigma_{vu}(i)}}\| \\
&\qquad + \|\mathbf{w_{\pi_{uw}\circ\sigma_{vu}(i)}} - \mathbf{v_{\sigma_{wv}\circ\pi_{uw}\circ\sigma_{vu}(i)}}\|\big)^2\Big] \\
&\leq \frac{1}{2} \mathop{\mathrm{E}}_{i\in[k]} \Big[3\big(\|\mathbf{v_i} - \mathbf{u_{\sigma_{vu}(i)}}\|^2 \\
&\qquad + \|\mathbf{u_{\sigma_{vu}(i)}} - \mathbf{w_{\pi_{uw}\circ\sigma_{vu}(i)}}\|^2 \\
&\qquad + \|\mathbf{w_{\pi_{uw}\circ\sigma_{vu}(i)}} - \mathbf{v_{\sigma_{wv}\circ\pi_{uw}\circ\sigma_{vu}(i)}}\|^2\big)\Big] \\
&= 3(\rho(v, u) + \epsilon_{uw} + \rho(w, v))
\end{aligned}$$

Thus, $\{u, w\}$ is satisfied with probability $1 - O(\rho(v, u) + \epsilon_{uw} + \rho(w, v))$. Averaging over all edges, we get a value $1 - O(\mathrm{E}_{u\in V}\, \rho(u, v) + \epsilon) = 1 - O(\Delta + \epsilon)$. □

3.4 Details of algorithm for locally conducting graphs

Proof of Theorem 1.1. Given an (α, ϕ_α)-locally conducting unique games instance G, begin by solving the SDP (Figure 2) and computing distance functions $\rho(u, v)$ (2.5). Partitioning G using Lemma 2.3 a set Z of outliers and sets S_i such that for all i, $S_i \setminus Z$ has average radius Δ.

Since the SDP has value at least $1 - \epsilon$, $\mathrm{E}_{\{u,v\}}\, \epsilon_{uv} < \epsilon$. Thus, at most a total volume $\sqrt{\epsilon}$ of sets S have $\epsilon_S > \sqrt{\epsilon}$ where we define $\epsilon_S = \mathrm{E}_{\{u,v\}\in E, u,v\in S}\, \epsilon_{uv}$. We may assume that we do not satisfy the at most $\sqrt{\epsilon}$ fraction of edges in such partitions.

For each $S = S_i \setminus Z$ that has average radius Δ set where $\epsilon_S \leq \sqrt{\epsilon}$, we use Lemma 2.5 to label S. This gives a labeling of value at least $1 - O(\Delta + \sqrt{\epsilon}) = 1 - O(\Delta)$ for S.

There is at most a $\Delta + \sqrt{\epsilon} = O(\Delta)$ total volume of partitions that do not get labelled. Thus we get a labeling for G of value at least $1 - O(\Delta)$. □

3.5 Details of Mildly Exponential Time Algorithm

Proof of Theorem 1.2. Begin by solving the SDP for G and computing distance functions $\rho(u, v)$. Use Lemma 2.2 with $\beta = 1$ to construct sets S_i and a set of outliers Z.

Since the SDP has value at least $1-\epsilon$, $\mathrm{E}_{\{u,v\}}\,\epsilon_{uv} < \epsilon$. Thus, at most a total volume $\sqrt[3]{\epsilon}$ of sets S have $\epsilon_S > \epsilon^{2/3}$ where we define $\epsilon_S = \mathrm{E}_{\{u,v\}\in E, u,v\in S}\,\epsilon_{uv}$. We may assume that we do not satisfy the at most $\sqrt[3]{\epsilon}$ fraction of edges in such sets.

For each set $S = S_i \setminus Z$ with $\epsilon_S \le \epsilon^{2/3}$, the following holds. If $\epsilon_S < \epsilon^{2/3}$, we still do the following, but make no guarantees on the value of the solution produced.

- If S is α-small:

  We begin by removing all vertices of degree less than $\sqrt[3]{\epsilon}d$, where $d = 2m/n$ is the average degree in G. Each vertex removed causes fewer than $\sqrt[3]{\epsilon}d$ edges to be removed. A total of at most n vertices are removed from all sets, so at most a total of $\sqrt[3]{\epsilon}dn = 2\sqrt[3]{\epsilon}m$ edges are removed. Let S' denote S with these edges removed. Since each vertex in S' has degree at least $\sqrt[3]{\epsilon}d$, and S' has volume at most α in G, there are at most $\alpha n/\sqrt[3]{\epsilon}$ vertices in S'.

  We will use Lemma 2.6 with $z = \sqrt[3]{\epsilon}$ to label S as if S had value $1-\epsilon^{2/3}$. This will take expected time $(1-\frac{\epsilon^{2/3}}{\sqrt[3]{\epsilon}})^{-\alpha n/\sqrt[3]{\epsilon}} poly(n) \le e^{\alpha n} poly(n)$.

- If S has average radius Δ and $\epsilon_S \le \epsilon^{2/3}$:

  We use Lemma 2.5 to label S. This gives a labeling of value at least $1-O\left(\Delta+\sqrt[3]{\epsilon}\right) = 1-O\left(\Delta\right)$

- Otherwise, we leave S unlabeled, but there is at most a total volume $O\left(\Delta+\sqrt[3]{\epsilon}\right)$ of such sets.

At most a $O\left(\sqrt[3]{\epsilon}+\Delta\right)$ total fraction of vertices and edges are ignored or unlabeled. We assume that all of these edges are unsatisfied. Thus, we get a labeling of value $1-O\left(\Delta\right)$. $\square$

4 Conclusions

All algorithms for unique games so far can be viewed as end-results of failed attempts to disprove UGC. Arguably, each failure brings us closer to understanding the UGC. This paper points to the need to understand unique game instances that have low but nontrivial local expansion. The recent paper [17] reminds us that understanding local expansion may also be difficult and in turn be related to UGC itself.

Our algorithm also gives an intuition about what SDP solutions of hard instances must look like. Since the rounding works well whenever we find a region with many more vectors than expected, the hard instances are those where the vectors are more or less uniformly spread around the sphere. (This will be made precise in the final version, since making it very precise actually involves changing the SDP relaxation.) Somewhat surprisingly, this “randomness” of distribution of vectors implies that the instance is *highly structured*, in that the uniformity gives a way to partition the graph into roughly equal sized components with few cut edges.

References

[1] Sanjeev Arora, Boaz Barak, and David Steurer. A $2^{n^{poly(\epsilon)}}$-time algorithm for unique games. Manuscript, 2010.

[2] Sanjeev Arora, Subhash A. Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 21–28, New York, NY, USA, 2008. ACM.

[3] Avrim Blum and Joel Spencer. Coloring random and semi-random k-colorable graphs. *J. Algorithms*, 19(2):204–234, 1995.

[4] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for unique games. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 205–214, New York, NY, USA, 2006. ACM Press.

[5] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. In *CCC '05: Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, pages 144–153, Washington, DC, USA, 2005. IEEE Computer Society.

[6] Eden Chlamtac, Konstantin Makarychev, and Yury Makarychev. How to play unique games using embeddings. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 687–696, Washington, DC, USA, 2006. IEEE Computer Society.

[7] Amin Coja-Oghlan. Solving np-hard semirandom graph problems in polynomial expected time. *J. Algorithms*, 62(1):19–46, 2007.

[8] Uriel Feige and Joe Kilian. Heuristics for finding large independent sets, with applications to coloring semi-random graphs. In *FOCS '98: Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, page 674, Washington, DC, USA, 1998. IEEE Computer Society.

[9] Uriel Feige and Joe Kilian. Heuristics for semirandom graph problems. *J. Comput. Syst. Sci.*, 63(4):639–673, 2001.

[10] J. Friedman, J. Kahn, and E. Szemerédi. On the second eigenvalue of random regular graphs. In *STOC '89: Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 587–598, New York, NY, USA, 1989. ACM.

[11] David Karger, Rajeev Motwani, and Madhu Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998.

[12] Subhash Khot. On the power of unique 2-prover 1-round games. In *STOC '02: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 767–775, New York, NY, USA, 2002. ACM Press.

[13] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? In *FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04)*, pages 146–154, Washington, DC, USA, 2004. IEEE Computer Society.

[14] Subhash Khot, Prasad Raghavendra, Rishi Saket, and David Steurer. Integrality gaps for strong sdp relaxations of unique games. In *Proc. IEEE FOCS 2009*.

[15] Konstantin Makarychev and Yury Makarychev. How to play unique games on expanders. Technical Report arXiv:0903.0367, March 2009.

[16] Prasad Raghavendra. Optimal algorithms and inapproximability results for every csp? In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 245–254, New York, NY, USA, 2008. ACM.

[17] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. Manuscript, 2009.

[18] Prasad Raghavendra, David Steurer, and Prasad Tetali. Approximations for the isoperimetric and spectral profile of graphs and for restricted eigenvalues of diagonally-dominant matrices. Manuscript, 2009.

[19] Luca Trevisan. Approximation algorithms for unique games. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 197–205, Washington, DC, USA, 2005. IEEE Computer Society.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de