How to Play Unique Games on Expanders

Konstantin Makarychev[1] and Yury Makarychev[2]

[1] IBM T.J. Watson Research Center
[2] Toyota Technological Institute at Chicago

Abstract. In this paper, we improve a result by Arora, Khot, Kolla, Steurer, Tulsiani, and Vishnoi on solving the Unique Games problem on expanders. Given a $(1-\varepsilon)$-satisfiable instance of Unique Games with the constraint graph G, our algorithm finds an assignment satisfying at least a $1 - C\varepsilon/h_G$ fraction of all constraints if $\varepsilon < c\lambda_G$ where h_G is the edge expansion of G, λ_G is the second smallest eigenvalue of the Laplacian of G, and C and c are some absolute constants.

1 Introduction

In this paper, we study Unique Games on expander graphs.

Definition 1 (Unique Games Problem) *Given a constraint graph $G = (V, E)$ and a set of permutations π_{uv} on the set $[k] = \{1, \dots, k\}$ (for all edges (u, v)), the goal is to assign a label (or state) x_u from $[k]$ to each vertex u so as to satisfy the maximum number of constraints of the form $\pi_{uv}(x_u) = x_v$. The value of a solution is the fraction of satisfied constraints.*

The famous Unique Games Conjecture (UGC) of Khot [8] states that for every positive ε and δ, there exists k such that it is NP-hard to distinguish between the case where a $1-\varepsilon$ fraction of all constraints is satisfiable and the case where at most a δ fraction of all constraints is satisfiable. This conjecture has attracted a lot of attention since it implies strong inapproximability results for such fundamental problems as MAX CUT [9], Vertex Cover [10], Maximum Acyclic Subgraph [6], k-CSP [7] [11], which are not known to follow from more standard complexity assumptions. Several approximation algorithms for Unique Games were developed in a series of papers by Khot [8], Trevisan [12], Gupta and Talwar [5], Charikar, Makarychev and Makarychev [3], and Chlamtac, Makarychev and Makarychev [4]. These papers, however, did not disprove the Unique Games Conjecture.

In order to better understand Unique Games, we need to identify which instances of Unique Games are easy, and which instances are potentially hard (the quantitative measure of hardness of a family of instances is the "approximation guarantee" of the "optimal" algorithm for this family). That motivates the study of specific families of Unique Games. Arora, Khot, Kolla, Steurer, Tulsiani, and Vishnoi [1] disproved the UGC for Unique Games on spectral expanders. Specifically, they showed how given a $(1-\varepsilon)$ satisfiable instance of Unique Games (i.e. an instance in which the optimal solution satisfies at least a $(1-\varepsilon)$ fraction of constraints), one can obtain a solution of value

$$1 - C\frac{\varepsilon}{\lambda_G}\log\left(\frac{\lambda_G}{\varepsilon}\right)$$

in polynomial time, here C is an absolute constant and λ_G is the second smallest eigenvalue of the Laplacian of G (see Section 2 for definitions).

In this paper, we improve their result and show that, if the ratio ε/λ_G is less than some universal positive constant c, one can obtain a solution of value

$$1 - C'\frac{\varepsilon}{h_G}$$

in polynomial time, here h_G is the edge expansion of G. In general, λ_G can be significantly smaller than h_G, then our result gives much better approximation guarantee. For example, if Cheeger's inequality (see below) is tight for a graph G, then $\lambda_G \approx h_G^2/8$; and

$$1 - C'\frac{\varepsilon}{h_G} \gg 1 - 8C\frac{\varepsilon}{h_G^2}\log\left(\frac{h_G^2}{8\varepsilon}\right) \approx 1 - C\frac{\varepsilon}{\lambda_G}\log\left(\frac{\lambda_G}{\varepsilon}\right).$$

Say, if $\varepsilon \approx \lambda_G$, the algorithm of Arora, Khot, Kolla, Steurer, Tulsiani, and Vishnoi satisfies only a small constant fraction of all constraints, while our algorithm satisfies almost all constraints. However, even if $\lambda_G \approx h_G$, our bound is asymptotically stronger, since

$$1 - C'\frac{\varepsilon}{h_G} \geq 1 - C'\frac{\varepsilon}{\lambda_G} > 1 - C'\frac{\varepsilon}{\lambda_G}\log\left(\frac{\lambda_G}{\varepsilon}\right)$$

(i.e., our bound does not have a $\log(\lambda_G/\varepsilon)$ factor).

1.1 Overview

In this section, we give an informal overview of the algorithm. The algorithm uses the standard SDP relaxation for Unique Games (see Section 2.2). The SDP solution gives a vector u_i for every vertex u and label i. For simplicity, let us consider so-called uniform case when all vectors u_i

have the same length. Then by scaling all vectors, we can assume that they are unit vectors, and thus vectors $u_1, \ldots, u_k$ (corresponding to one vertex) form an orthonormal frame.

For every two vertices u and v, we say that labels i and j are *matched* if $\|u_i - v_j\|^2 < r$, where $r < 1$ is a small threshold value. Note that for every two labels j_1 and j_2,

$$\|u_i - v_{j_1}\|^2 + \|u_i - v_{j_2}\|^2 \geq \|v_{j_1} - v_{j_2}\|^2 = 2 > 2r.$$

Therefore, each label i is matched with at most one label j for fixed vertices u and v. We denote this j by $\sigma_{uv}(i)$ (if it exists).

Now we use a simple prorogation algorithm. We choose a random vector u and assign it a random label i. Then we label each vertex v with the label $\sigma_{uv}(i)$, if it is defined; and with an arbitrary label, otherwise. Let X be the set of vertices v s.t. $\sigma_{uv}(i)$ is defined. We prove (see Lemma 9) that for every edge (v, w) with $v, w \in X$, our assignment satisfies the constraint between v and w w.h.p. if the contribution of the edge (v, w) to the SDP objective is small. Intuitively, that happens because both vectors $v_{\sigma_{uv}(i)}$ and $w_{\sigma_{uw}(i)}$ are close to u_i, and therefore they are close to each other. On the other hand, the SDP contribution of the edge (v, w) equals

$$\frac{1}{k}\sum_{j=1}^{k} \|v_j - w_{\pi_{vw}(j)}\|^2.$$

Thus if the SDP contribution is small then the vector v_j should be close to $w_{\pi_{vw}}(j)$ for most labels j. Since each $v_{\sigma_{uv}(i)}$ is close only to $w_{\sigma_{uw}(i)}$, we have $\sigma_{uw}(i) = \pi_{vw}(\sigma_{uv}(i))$ w.h.p., that is, the constraint between v and w is satisfied.

The crucial step now is to prove that the set X contains almost all vertices, and so we can ignore edges with one or two endpoints outside of X. First, we prove that the set X is not very small in Lemma 5 (using a "global correlation" result of Arora, Khot, Kolla, Steurer, Tulsiani, and Vishnoi). Using a standard region growing argument we then show that the cut between X and $V \setminus X$ is very small (if we choose the threshold r randomly; see Lemma 7). Since the graph G is an expander, that implies that either X or $V \setminus X$ is very small. But we know that X is not very small. We conclude that in fact $V \setminus X$ is very small (Lemma 8).

To deal with the general case — when vectors u_i have different lengths — we use the vector normalization machinery developed by Chlamtac, Makarychev and Makarychev [4].

In Section 2, we give basic definitions and describe the semideifnite relaxation for Unique Games. In Section 3, we present the algorithm and its analysis.

2 Preliminaries

2.1 Expanders: Second Eigenvalue and Edge Expansion

We assume that the underlying constraint graph $G = (V, E)$ is a d-regular expander. The two key parameters of the expander G are the edge expansion h_G and the second eigenvalue of the Laplacian λ_G. The edge expansion gives a lower bound on the size of every cut: for every subset of vertices $X \subset V$, the size of the cut between X and $|V \setminus X|$ is at least

$$h_G \times \frac{\min(|X|, |V \setminus X|)}{|V|} |E|.$$

It is formally defined as follows:

$$h_G = \min_{X \subset V} \left(\frac{|\delta(X, V \setminus X)|}{|E|} \Big/ \frac{\min(|X|, |V \setminus X|)}{|V|} \right),$$

here $\delta(X, V \setminus X)$ denotes the cut — the set of edges going from X to $V \setminus X$. One can think of the second eigenvalue of the Laplacian matrix

$$L_G(u, v) = \begin{cases} 1, & \text{if } u = v \\ -1/d, & \text{if } (u, v) \in E \\ 0, & \text{otherwise.} \end{cases}$$

as of continuous relaxation of the edge expansion. Note that the smallest eigenvalue of L_G is 0; and the corresponding eigenvector is a vector of all 1's, denoted by $\mathbf{1}$. Thus

$$\lambda_G = \min_{x \perp \mathbf{1}} \frac{\langle x, L_G x \rangle}{\|x\|^2}.$$

Cheeger's inequality,

$$h_G^2/8 \leq \lambda_G \leq h_G,$$

shows that h_G and λ_G are closely related; however λ_G can be much smaller than h_G (the lower bound in the inequality is tight).

2.2 Semidefinite Relaxation for Unique Games

We use the standard SDP relaxation for the Unique Games problem.

$$\text{minimize } \frac{1}{2|E|} \sum_{(u,v)\in E} \sum_{i=1}^{k} \|u_i - v_{\pi_{uv}(i)}\|^2$$

subject to

$$\forall u \in V\ \forall i,j \in [k], i \neq j \quad \langle u_i, u_j \rangle = 0 \tag{1}$$

$$\forall u \in V \quad \sum_{i=1}^{k} \|u_i\|^2 = 1 \tag{2}$$

$$\forall u,v,w \in V\ \forall i,j,l \in [k] \quad \|u_i - w_l\|^2 \leq \|u_i - v_j\|^2 + \|v_j - w_l\|^2 \tag{3}$$

$$\forall u,v \in V\ \forall i,j \in [k] \quad \|u_i - v_j\|^2 \leq \|u_i\|^2 + \|v_j\|^2 \tag{4}$$

$$\forall u,v \in V\ \forall i,j \in [k] \quad \|u_i\|^2 \leq \|u_i - v_j\|^2 + \|v_j\|^2 \tag{5}$$

For every vertex u and label i we introduce a vector u_i. In the intended integral solution $u_i = 1$, if u is labeled with i; and $u_i = 0$, otherwise. All SDP constraints are satisfied in the integral solution; thus this is a valid relaxation. The objective function of the SDP measures what fraction of all Unique Games constraints is *not satisfied*.

3 Algorithm

We define the *earthmover distance* between two sets of orthogonal vectors $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$ as follows:

$$\Delta(\{u\}_i, \{v\}_i) \equiv \min_{\sigma(i)\in\mathcal{S}_k} \sum_{i=1}^{k} \|u_i - v_{\sigma(i)}\|^2,$$

here $\mathcal{S}_k$ is the symmetric group, the group of all permutations on the set $[k] = \{1, \dots, k\}$. Given an SDP solution $\{u_i\}_{u,i}$ we define the earthmover distance between vertices in a natural way:

$$\Delta(u,v) = \Delta(\{u_1, \dots, u_k\}, \{v_1, \dots, v_k\}).$$

Arora et al. [1] proved that if an instance of Unique Games on an expander is almost satisfiable, then the average earthmover distance between two vertices (defined by the SDP solution) is small. We will need the following corollary from their results:

For every $R \in (0,1)$, there exists a positive c, such that for every $(1-\varepsilon)$ satisfiable instance of Unique Games on an expander graph G, if $\varepsilon/\lambda_G < c$, then the expected earthmover distance between two random vertices is less than R i.e.

$$\mathbb{E}_{u,v \in V}\left[\Delta(u,v)\right] \leq R.$$

In fact, Arora et al. [1] showed that $c \geq \Omega(R/\log(1/R))$, but we will not use this bound. Moreover, in the rest of the paper, we fix the value of $R < 1/4$. We pick c_R, so that if $\varepsilon/\lambda_G < c_R$, then

$$\mathbb{E}_{u,v \in V}\left[\Delta(u,v)\right] \leq R/4. \tag{6}$$

Our algorithm transforms vectors $\{u_i\}_{u,i}$ in the SDP solution to vectors $\{\tilde{u}_i\}_{u,i}$ using a *vector normalization* technique introduced by Chlamtac, Makarychev and Makarychev [4]:

Lemma 1 *[4] For every SDP solution $\{u_i\}_{u,i}$, there exists a set of vectors $\{\tilde{u}_i\}_{u,i}$ satisfying the following properties:*

1. *Triangle inequalities in ℓ_2^2: for all vertices u, v, w in V and all labels i, p, q in $[k]$,*

$$\|\tilde{u}_i - \tilde{v}_p\|_2^2 + \|\tilde{v}_p - \tilde{w}_q\|_2^2 \geq \|\tilde{u}_i - \tilde{w}_q\|_2^2.$$

2. *For all vertices u, v in V and all labels i, j in $[k]$,*

$$\langle \tilde{u}_i, \tilde{v}_j \rangle = \frac{\langle u_i, v_j \rangle}{\max(\|u_i\|^2, \|v_j\|^2)}.$$

3. *For all non-zero vectors u_i, $\|\tilde{u}_i\|_2^2 = 1$.*
4. *For all u in V and $i \neq j$ in $[k]$, the vectors $\tilde{u}_i$ and $\tilde{u}_j$ are orthogonal.*
5. *For all u and v in V and i and j in $[k]$,*

$$\|\tilde{v}_j - \tilde{u}_i\|_2^2 \leq \frac{2\,\|v_j - u_i\|^2}{\max(\|u_i\|^2, \|v_j\|^2)}.$$

The set of vectors $\{\tilde{u}_i\}_{u,i}$ can be obtained in polynomial time.

Now we are ready to describe the rounding algorithm. The algorithm given an SDP solution, outputs an assignment of labels to the vertices.

Approximation Algorithm

Input: an SDP solution $\{u_i\}_{u,i}$ of cost ε.

Initialization

1. Pick a random vertex u (uniformly distributed) in V. We call this vertex *the initial vertex*.
2. Pick a random label $i \in [k]$ for u; choose label i with probability $\|u_i\|^2$. Note that $\|u_1\|^2 + \cdots + \|u_k\|^2 = 1$. We call i *the initial label*.
3. Pick a random number t uniformly distributed in the segment $[0, \|u_i\|^2]$.
4. Pick a random r in $[R, 2R]$.

Normalization

5. Obtain vectors $\{\tilde{u}_i\}_{u,i}$ as in Lemma 1.

Propagation

6. For every vertex v,
   - Find all labels $p \in [k]$ such that $\|v_p\|^2 \geq t$ and $\|\tilde{v}_p - \tilde{u}_i\|^2 \leq r$. Denote the set of p's by S_v:
   $$S_v = \left\{p : \|v_p\|^2 \geq t \text{ and } \|\tilde{v}_p - \tilde{u}_i\|^2 \leq r\right\}.$$
   - If S_v contains exactly one element p, then assign the label p to v.
   - Otherwise, assign an arbitrary (say, random) label to v.

Denote by σ_{vw} the partial mapping from $[k]$ to $[k]$ that maps p to q if $\|\tilde{v}_p - \tilde{w}_q\|^2 \leq 4R$. Note that σ_{vw} is well defined i.e. p cannot be mapped to different labels q and q': if $\|\tilde{v}_p - \tilde{w}_q\|^2 \leq 4R$ and $\|\tilde{v}_p - \tilde{w}_{q'}\|^2 \leq 4R$, then, by the ℓ_2^2 triangle inequality (see Lemma 1, item 1), $\|\tilde{w}_q - \tilde{w}_{q'}\|^2 \leq 8R$, but $\tilde{w}_q$ and $\tilde{w}_{q'}$ are orthogonal unit vectors, so

$$\|\tilde{w}_q - \tilde{w}_{q'}\|^2 = 2 > 8R.$$

Clearly, σ_{vw} defines a partial matching between labels of v and labels of w: if $\sigma_{vw}(p) = q$, then $\sigma_{wv}(q) = p$.

Lemma 2 *If $p \in S_v$ and $q \in S_w$ with non-zero probability, then $q = \sigma_{vw}(p)$.*

Proof. If $p \in S_v$ and $q \in S_w$ then for some vertex u and label i, $\|\tilde{v}_p - \tilde{u}_i\|^2 \leq 2R$ and $\|\tilde{w}_q - \tilde{u}_i\|^2 \leq 2R$, thus by the triangle inequality $\|\tilde{v}_p - \tilde{w}_q\|^2 \leq 4R$ and by the definition of σ_{vw}, $q = \sigma_{vw}(p)$.

Corollary 3 *Suppose, that $p \in S_v$, then the set S_w either equals $\{\sigma_{vw}(p)\}$ or is empty (if $\sigma_{vw}(p)$ is not defined, then S_w is empty). Particularly, if u and i are the initial vertex and label, then the set S_w either equals $\{\sigma_{uw}(i)\}$ or is empty. Thus, every set S_w contains at most one element.*

Lemma 4 *For every choice of the initial vertex u, for every $v \in V$ and $p \in [k]$ the probability that $p \in S_v$ is at most $\|v_p\|^2$.*

Proof. If $p \in S_v$, then $i = \sigma_{vu}(p)$ is the initial label of u and $t \le \|v_p\|^2$. The probability that both these events happen is

$$\Pr(i \in S_u) \times \Pr\left(t \le \|v_p\|^2\right) = \|u_i\|^2 \times \min(\|v_p\|^2/\|u_i\|^2, 1) \le \|v_p\|^2$$

(recall that t is a random real number on the segment $[0, \|u_i\|^2]$).

Denote the set of those vertices v for which S_v contains exactly one element by X. First, we show that on average X contains a constant fraction of all vertices (later we will prove a much stronger bound on the size of X).

Lemma 5 *If $\varepsilon/\lambda_G \le c_R$, then the expected size of X is at least $|V|/4$.*

Proof. Consider an arbitrary vertex v. Estimate the probability that $p \in S_v$ given that u is the initial vertex. Suppose that there exists q such that $\|v_p - u_q\|^2 \le \|v_p\|^2 \cdot R/2$, then

$$\|\tilde{u}_q - \tilde{v}_p\|^2 \le \frac{2\|u_q - v_p\|^2}{\max(\|u_q\|^2, \|v_p\|^2)} \le R.$$

Thus, $q = \sigma_{vu}(p)$ and $\|\tilde{u}_q - \tilde{v}_p\|^2 \le r$ with probability 1. Hence, if q is chosen as the initial label and $\|v_p\|^2 \ge t$, then $v_p \in S_v$. The probability of this event is $\|u_q\|^2 \times \min(\|v_p\|^2/\|u_q\|^2, 1)$. Notice that

$$\|u_q\|^2 \times \min(\|v_p\|^2/\|u_q\|^2, 1) = \min(\|v_p\|^2, \|u_q\|^2) \ge \|v_p\|^2 - \|u_q - v_p\|^2 \ge \frac{\|v_p\|^2}{2}.$$

Now, consider all p's for which there exists q such that $\|v_p - u_q\|^2 \le \|v_p\|^2 \cdot R/2$. The probability that one of them belongs to S_v, and thus $v \in X$, is at least

$$\begin{aligned}
\frac{1}{2} \sum_{p:\min_q(\|u_q - v_p\|^2) \le \|v_p\|^2 \cdot R/2} \|v_p\|^2 &= \frac{1}{2}\sum_{p=1}^{k} \|v_p\|^2 - \frac{1}{2} \sum_{p:\min_q(\|u_q - v_p\|^2) > \|v_p\|^2 \cdot R/2} \|v_p\|^2 \\
&\ge \frac{1}{2} - \frac{1}{2} \times \sum_{p=1}^{k} \frac{2}{R} \min_q(\|u_q - v_p\|^2) \\
&\ge \frac{1}{2} - \frac{\Delta(\{u\}_q, \{v\}_p)}{R}.
\end{aligned}$$

Since the average value of $\Delta(\{u\}_q, \{v\}_p)$ over all pairs (u, v) is at most $R/4$ (see (6)), the expected size of X (for a random initial vertex u) is at least $|V|/4$.

Corollary 6 *If $\varepsilon/\lambda_G \le c_R$, then*

$$\Pr(|X| > |V|/8) > \frac{1}{8}.$$

Lemma 7 *The expected size of the cut between X and $V \setminus X$ is at most $6\varepsilon/R|E|$.*

Proof. We show that the size of the cut between X and $V \setminus X$ is at most $6\varepsilon/R|E|$ in the expectation for any choice of the initial vertex u. Fix an edge (v, w) and estimate the probability that $v \in X$ and $w \in V \setminus X$. If $v \in X$ and $w \in V \setminus X$, then S_v contains a unique label p, but S_w is empty (see Corollary 3) and, particularly, $\pi_{vw}(p) \notin S_w$. This happens in two cases:

- There exists p such that $i = \sigma_{vu}(p)$ is the initial label of u and $\|w_{\pi_{vw}(p)}\|^2 < t \le \|v_p\|^2$. The probability of this event is at most

$$\sum_{p=1}^{k} \|u_{\sigma_{vu}(p)}\|^2 \times \left| \frac{\|v_p\|^2 - \|w_{\pi_{vw}(p)}\|^2}{\|u_{\sigma_{vu}(p)}\|^2} \right| \le \sum_{p=1}^{k} \|v_p - w_{\pi_{vw}(p)}\|^2.$$

- There exists p such that $i = \sigma_{vu}(p)$ is the initial label of u, $t \le \|v_p\|^2$ and $\|\tilde{u}_i - \tilde{v}_p\|^2 < r \le \|\tilde{u}_i - \tilde{w}_{\pi_{vw}(p)}\|^2$. The probability of this event is at most

$$\sum_{p=1}^{k} \|u_{\sigma_{vu}(p)}\|^2 \times \frac{\|v_p\|^2}{\|u_{\sigma_{vu}(p)}\|^2} \times \left| \frac{\|\tilde{u}_{\sigma_{vu}(p)} - \tilde{w}_{\pi_{vw}(p)}\|^2 - \|\tilde{u}_{\sigma_{vu}(p)} - \tilde{v}_p\|^2}{R} \right|$$
$$\le \sum_{p=1}^{k} \|v_p\|^2 \times \frac{\|\tilde{v}_p - \tilde{w}_{\pi_{vw}(p)}\|^2}{R} \le \sum_{p=1}^{k} \|v_p\|^2 \times \frac{2\|v_p - w_{\pi_{vw}(p)}\|^2}{R \cdot \max(\|v_p\|^2, \|w_{\pi_{vw}(p)}\|^2)}$$
$$\le \frac{2}{R} \sum_{p=1}^{k} \|v_p - w_{\pi_{vw}(p)}\|^2.$$

Note that the probability of the first event is zero, if $\|w_{\pi_{vw}(p)}\|^2 \ge \|v_p\|^2$; and the probability of the second event is zero, if $\|\tilde{u}_{\sigma_{vu}(p)} - \tilde{v}_p\|^2 \ge \|\tilde{u}_{\sigma_{vu}(p)} - \tilde{w}_{\pi_{vw}(p)}\|^2$.

Since the SDP value equals

$$\frac{1}{2|E|}\sum_{(v,w)\in E}\sum_{p=1}^{k}\|v_p - w_{\pi_{vw}(p)}\|^2 \le \varepsilon.$$

The expected fraction of cut edges is at most $6\varepsilon/R$.

Lemma 8 *If $\varepsilon \le \min(c_R\lambda_G, h_G R/1000)$, then with probability at least 1/16 the size of X is at least*

$$\left(1-\frac{100\varepsilon}{h_G R}\right)|V|.$$

Proof. The expected size of the cut $\delta(X, V\setminus X)$ between X and $V\setminus X$ is less than $6\varepsilon/R|E|$. Hence, since the graph G is an expander, one of the sets X or $V\setminus X$ must be small:

$$\mathbb{E}\left[\min(|X|,|V\setminus X|)\right] \le \frac{1}{h_G}\times\frac{\mathbb{E}\left[|\delta(X,V\setminus X)|\right]}{|E|}\times|V| \le \frac{6\varepsilon}{h_G R}|V|.$$

By Markov's Inequality,

$$\Pr\left(\min(|X|,|V\setminus X|)\le\frac{100\varepsilon}{h_G R}|V|\right)\ge 1-\frac{1}{16}.$$

Observe, that $100\varepsilon/(h_G R)|V| < |V|/8$. However, by Corollary 6, the size of X is greater than $|V|/8$ with probability greater than $1/8$. Thus

$$\Pr\left(|V\setminus X|\le\frac{100\varepsilon}{h_G R}|V|\right)\ge\frac{1}{16}.$$

Lemma 9 *The probability that for an arbitrary edge (v,w), the constraint between v and w is not satisfied, but v and w are in X is at most $4\varepsilon_{vw}$, where*

$$\varepsilon_{vw}=\frac{1}{2}\sum_{i=1}^{k}\|v_i - w_{\pi_{vw}(i)}\|^2.$$

Proof. We show that for every choice of the initial vertex u the desired probability is at most $4\varepsilon_{vw}$. Recall, that if $p\in S_v$ and $q\in S_w$, then $q=\sigma_{vw}(p)$. The constraint between v and w is not satisfied if $q\ne\pi_{vw}(p)$. Hence, the probability that the constraint is not satisfied is at most,

$$\sum_{p:\pi_{vw}(p)\ne\sigma_{vw}(p)}\Pr\left(p\in S_v\right).$$

If $\pi_{vw}(p) \neq \sigma_{vw}(p)$, then

$$\|\tilde{v}_p - \tilde{w}_{\pi_{vw}(p)}\|^2 \geq \|\tilde{w}_{\pi_{vw}(p)} - \tilde{w}_{\sigma_{vw}(p)}\|^2 - \|\tilde{v}_p - \tilde{w}_{\sigma_{vw}(p)}\|^2 \geq 2 - 4R \geq 1.$$

Hence, by Lemma 1 (5),

$$\|v_p - w_{\pi_{vw}(p)}\|^2 \geq \|v_p\|^2/2.$$

Therefore, by Lemma 4,

$$\sum_{p:\pi_{vw}(p)\neq\sigma_{vw}(p)} \Pr(p \in S_v) \leq \sum_{p:\pi_{vw}(p)\neq\sigma_{vw}(p)} \|v_p\|^2 \leq 2\sum_{p=1}^{k} \|v_p - w_{\pi_{vw}(p)}\|^2 = 4\varepsilon_{vw}.$$

Theorem 10 *There exists a polynomial time approximation algorithm that given a $(1-\varepsilon)$ satisfiable instance of Unique Games on a d-expander graph G with $\varepsilon/\lambda_G \leq c$, the algorithm finds a solution of value*

$$1 - C\frac{\varepsilon}{h_G},$$

where c and C are some positive absolute constants.

Proof. We describe a randomized polynomial time algorithm. Our algorithm may return a solution to the SDP or output a special value *fail*. We show that the algorithm outputs a solution with a constant probability (that is, the probability of failure is bounded away from 1); and conditional on the event that the algorithm outputs a solution its expected value is

$$1 - C\frac{\varepsilon}{h_G}. \tag{7}$$

Then we argue that the algorithm can be easily derandomized — simply by enumerating all possible values of the random variables used in the algorithm and picking the best solution. Hence, the deterministic algorithm finds a solution of value at least (7).

The randomized algorithm first solves the SDP and then runs the rounding procedure described above. If the size of the set X is more than

$$\left(1 - \frac{100\varepsilon}{h_G R}\right)|V|,$$

the algorithm outputs the obtained solution; otherwise, it outputs *fail*.

Let us analyze the algorithm. By Lemma 8, it succeeds with probability at least 1/16. The fraction of edges having at least one endpoint in $V \setminus X$ is at most $100\varepsilon/(h_G R)$ (since the graph is d-regular). We conservatively assume that the constraints corresponding to these edges are violated. The expected number of violated constraints between vertices in X, by Lemma 9 is at most

$$\frac{4\sum_{(u,v)\in E}\varepsilon_{uv}}{\Pr\left(|X| \geq 100\varepsilon/(h_G R)\right)} \leq 64 \times \left(\frac{1}{2}\sum_{(u,v)\in E}\|u_i - v_{\pi_{vw}(i)}\|^2\right) \leq 64\varepsilon|E|.$$

The total fraction of violated constraints is at most $100\varepsilon/(h_G R) + 64\varepsilon$.

References

1. S. Arora, S. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. Vishnoi. *Unique games on expanding constraint graphs are easy.* In Proceedings of the 40th ACM Symposium on Theory of Computing, pp. 21–28, 2008.
2. P. Austrin and E. Mossel. *Approximation resistant predicates from pairwise independence.* In Proceedings of the 2008 IEEE 23rd Annual Conference on Computational Complexity. IEEE Computer Society, 249–258.
3. M. Charikar, K. Makarychev, and Y. Makarychev. *Near-Optimal Algorithms for Unique Games.* In Proceedings of the 38th ACM Symposium on Theory of Computing, pp. 205–214, 2006.
4. E. Chlamtac, K. Makarychev, and Y. Makarychev. *How to Play Unique Games Using Embeddings.* In Proceedings of the 47th IEEE Symposium on Foundations of Computer Science, pp. 687–696, 2006.
5. A. Gupta and K. Talwar. *Approximating Unique Games.* In Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms, pp. 99–106, 2006.
6. V. Guruswami, R. Manokaran, and P. Raghavendra. *Beating the Random Ordering is Hard: Inapproximability of Maximum Acyclic Subgraph.* In Proceedings of the 49th IEEE Symposium on Foundations of Computer Science, pp. 573–582, 2008.
7. V. Guruswami and P. Raghavendra. *Constraint satisfaction over a non-boolean domain: Approximation algorithms and unique-games hardness.* In Proceedings of APPROX-RANDOM, 77–90, 2008.
8. S. Khot. *On the power of unique 2-prover 1-round games.* In Proceedings of the 34th ACM Symposium on Theory of Computing, pp. 767–775, 2002.
9. S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. *Optimal inapproximability results for MAX-CUT and other two-variable CSPs?* SIAM Journal of Computing 37(1), pp. 319–357, 2007.
10. S. Khot and O. Regev. *Vertex cover might be hard to approximate to within* $2-\varepsilon$.. In Proceedings of the 18th IEEE Annual Conference on Computational Complexity, 2003.
11. A. Samorodnitsky and L. Trevisan. *Gowers uniformity, influence of variables, and PCPs.* In Proceedings of the 38th annual ACM symposium on Theory of computing, pp. 11–20, 2006.
12. L. Trevisan. *Approximation Algorithms for Unique Games.* In Proceedings of the 46th IEEE Symposium on Foundations of Computer Science, pp. 197–205, 2005.