Improved Rounding for Parallel Repeated Unique Games

David Steurer*

Princeton University, Princeton NJ 08544, USA.

Abstract. We show a tight relation between the behavior of unique games under parallel repetition and their semidefinite value. Let G be a unique game with alphabet size k. Suppose the semidefinite value of G, denoted $\mathrm{sdp}(G)$, is at least $1-\varepsilon$. Then, we show that the optimal value $\mathrm{opt}(G^\ell)$ of the ℓ-fold repetition of G is at least $1-O(\sqrt{\ell\varepsilon\log k})$. This bound confirms a conjecture of Barak et al. (2008), who showed a lower bound that was worse by $\sqrt{\ell\varepsilon\log(1/\varepsilon)}$. A consequence of our bound is the following tight relation between the semidefinite value of G and the amortized value $\overline{\mathrm{opt}}(G) := \sup_{\ell\in\mathbb{N}}\mathrm{opt}(G^\ell)^{1/\ell}$,

$$\mathrm{sdp}(G)^{O(\log k)} \le \overline{\mathrm{opt}}(G) \le \mathrm{sdp}(G)\,.$$

The proof closely follows the approach of Barak et al. (2008). Our technical contribution is a natural orthogonalization procedure for nonnegative functions. The procedure has the property that it preserves distances up to an absolute constant factor. In particular, our orthogonalization avoids the additive increase in distances caused by the truncation step of Barak et al. (2008).

1 Introduction

A *unique game* G with *vertex set* V and *alphabet* Σ consists of a list of constraints encoded by triples (u,v,π), where $u,v\in V$ are vertices and $\pi\in \mathbb{S}_\Sigma$ is a permutation of Σ. An *assignment* $x\in\Sigma^V$ satisfies a constraint (u,v,π) if $x_v=\pi(x_u)$. The *(optimal) value* of G, denote $\mathrm{opt}(G)$, is defined as the maximum fraction of constraints of G that can be satisfied simultaneously, that is, $\mathrm{opt}(G) := \max_{x\in\Sigma^V}\mathbf{P}_{(u,v,\pi)\sim G}\{x_v=\pi(x_u)\}$. (Here, $(u,v,\pi)\sim G$ means that (u,v,π) is a random constraint of G.)

Khot's Unique Games Conjecture [9] asserts that it is NP-hard to approximate the value of a unique game in a certain regime. (According to this conjecture, for every constant $\varepsilon>0$, there exists $k\in\mathbb{N}$ such that given a unique game G with alphabet size at most k, it is NP-hard to distinguish the cases $\mathrm{opt}(G)\ge 1-\varepsilon$ and $\mathrm{opt}(G)\le\varepsilon$.)

A sequence of recent works showed that this conjecture implies (often optimal) hardness results for many basic combinatorial optimization problems [9,11,10,14,12,4,1,15,13,7]. Most strikingly, Raghavendra [15] showed that the

* Supported by NSF Grants CCF-0832797, 0830673, and 0528414.

Unique Games Conjecture, if true, implies that for every constraint satisfaction problem, it is NP-hard to achieve a strictly better approximation guarantee than the one obtained by a simple generic semidefinite programming relaxation.

Due to its potential consequences, the Unique Games Conjecture is one of the central open questions about approximation algorithms and hardness of approximation.

In contrast to the numerous implications of a positive resolution, only few consequence of a refutation of the conjecture are known (some consequence for the approximability of graph expansion in a certain regime were recently shown [16]). The most compelling question in this context is whether an algorithm refuting the Unique Games Conjecture would lead to better approximation algorithms for Max Cut. (Or equivalently: Would the optimality of current algorithms for Max Cut, imply the Unique Games Conjecture?)

A natural candidate for such a reduction from Max Cut to Unique Games (the problem of computing the value of a unique game) is *parallel repetition*. (For simplicity, we can identify Max Cut with the problem of computing the value of a unique game with alphabet size 2.) In general, ℓ-fold parallel repetition takes a unique game G with vertex V and alphabet Σ and outputs a unique game, denoted G^ℓ, with vertex set V^ℓ and alphabet Σ^ℓ. For every ℓ-tuple $(u_1, v_1, \pi_1), \ldots, (u_\ell, v_\ell, \pi_\ell)$ of constraints in G, the game G^ℓ contains a constraint $(\underline{u}, \underline{v}, \underline{\pi})$, where $\underline{u} = (u_1, \ldots, u_\ell)$, $\underline{v} = (v_1, \ldots, v_\ell)$, and $\underline{\pi}$ is the permutation of Σ^ℓ obtained by applying π_i to the i^{th} coordinate. By construction, $\operatorname{opt}(G^\ell) \geq \operatorname{opt}(G)^\ell$. However, this lower bound is not always tight. Raz [18] showed the first strong upper bound on the value of parallel-repeated games: If $\operatorname{opt}(G) = 1 - \varepsilon$, then $\operatorname{opt}(G^\ell) \leq (1 - \varepsilon^c)^{\Omega(\ell/s)}$ for some constant c and $s = O(\log k)$. Subsequently, Holenstein [8] simplified this proof and showed an improved parallel repetition bound with $c = 3$. Rao [17] further improved this result and showed a parallel repetition bound with $c = 2$ and $s = O(1)$. We remark that these bounds also hold for *projection games*, where π need not be a permutation but can be an arbitrary function. (In fact, the bounds of Raz and Holenstein hold in an even more general setting.)

Feige et al. [5] noted that a parallel repetition bound with $c < 2$ would imply a reduction from Max Cut to Unique Games that achieves the goal above (an algorithm refuting the Unique Games Conjecture would lead to a better approximation algorithm for Max Cut). Raz [19] ruled out this possibility and showed that for a simple family of unique games (odd-cycle games) Rao's parallel repetition bound is optimal.

A related, more general question is whether parallel repeated games could be hard instances for Unique Games. Concretely, we could consider the following strengthening of the Unique Games Conjecture: For every constant $\varepsilon > 0$, there exists $\ell \in \mathbb{N}$ such that given a unique game G with alphabet size 2 it is NP-hard to distinguish the cases $\operatorname{opt}(G^\ell) \geq 1 - \varepsilon$ and $\operatorname{opt}(G^\ell) \leq \varepsilon$. We remark that the analogous conjecture for projection games is known to be true (a consequence of the PCP Theorem and Raz's parallel repetition bound). Extending the techniques of Raz's analysis of odd-cycle games, Barak et al. [2] showed that this variant of the Unique Games Conjecture is false (unless P = NP). The authors give a polynomial time algorithm that given a unique game G with alphabet size 2 and $\mathrm{opt}(G^\ell) \geq 1-\varepsilon$, computes an assignment for G^ℓ of value $1-O(\sqrt{\varepsilon})$. (In particular, for every small enough $\varepsilon > 0$, this algorithm can distinguish between $\mathrm{opt}(G^\ell) \geq 1-\varepsilon$ and $\mathrm{opt}(G^\ell) \leq \varepsilon$.)

The algorithm of Barak et al. is based on a semidefinite programming relaxation for $\mathrm{opt}(G)$. They show that the behavior of $\mathrm{opt}(G^\ell)$ is closely characterized by the optimal value of this semidefinite relaxation, denoted $\mathrm{sdp}(G)$. For the case that G has alphabet size 2, the bound of Barak et al. is tight up to constant factors. For the case of larger alphabets, the authors show bounds that are tight up to logarithmic factors. In this work, we improve the analysis of Barak et al. and show a relation between $\mathrm{sdp}(G)$ and the behavior of $\mathrm{opt}(G^\ell)$, which is tight up to constant factors even for larger alphabets.

1.1 Results

Let G be a unique game with alphabet size k. We are interested in the behavior of the value of G under parallel repetition. Barak et al. [2] show that if the semidefinite value of G is at least $1-\varepsilon$, then the optimal value of G^ℓ is at least $1-O(\sqrt{s\ell\varepsilon})$, where $s = \log k + \log(1/\varepsilon)$. It was conjectured that the $\log(1/\varepsilon)$ term in this bound is not necessary. (The $\log k$ term is known to be necessary.) We confirm this conjecture and show the following lower bound on $\mathrm{opt}(G^\ell)$ in terms of $\mathrm{sdp}(G)$.

Theorem 1. *For every unique game G with alphabet size k and* $\mathrm{sdp}(G) \geq 1-\varepsilon$*,*

$$\mathrm{opt}(G^\ell) \geq 1 - O(\sqrt{\ell\varepsilon\log k})\,.$$

As a consequence of this theorem and results of Feige and Lovász [6] and Charikar, Makarychev, and Makarychev [3], we obtain the following tight relation between the *amortized value* $\overline{\mathrm{opt}}(G) := \sup_{\ell\in\mathbb{N}} \mathrm{opt}(G^\ell)^{1/\ell}$ and the semidefinite value of G. (See Section A.1 for a proof of this theorem.)

Theorem 2. *For every unique game G with alphabet size k,*

$$\mathrm{sdp}(G)^{O(\log k)} \leq \overline{\mathrm{opt}}(G) \leq \mathrm{sdp}(G)\,.$$

(We remark that for every $k \in \mathbb{N}$, there exist unique games G that achieve the above lower bound on $\overline{\mathrm{opt}}(G)$. See [2] for details.) The approach of [2] for proving lower bounds on the value of repeated games G^ℓ involves an intermediate relaxation, denoted here $\mathrm{sdp}_+(G)$, which is "sandwiched" between $\mathrm{opt}(G)$ and $\mathrm{sdp}(G)$. (This approach is also implicit in Raz's counterexample to strong parallel repetition [19].) Using Holenstein's correlated sampling technique [8], it is straightforward to derive lower bounds on $\mathrm{opt}(G^\ell)$ in terms of $\mathrm{sdp}_+(G)$ (see Section 2.2 for the definition of $\mathrm{sdp}_+(G)$ and [2] for more discussion about the correlated sampling technique).

The key result in [2] is a lower bound on $\mathrm{sdp}_+(G)$ in terms of the semidefinite value of G. We prove the following improved bound (which is optimal up to the constants hidden in the $O(\cdot)$-notation).

Theorem 3. *For every unique G with alphabet size k and* $\mathrm{sdp}(G) \geq 1-\varepsilon$*,*

$$\mathrm{sdp}_+(G) \geq 1 - O(\varepsilon \log k)\,.$$

Assuming this theorem and some properties of the intermediate relaxation $\mathrm{sdp}_+(G)$ (which are presented in Section 2.2), we can prove Theorem 1.

Proof (Theorem 1). Let G be a unique game with alphabet size k and $\mathrm{sdp}(G) \geq 1-\varepsilon$. Theorem 3 shows that $\mathrm{sdp}_+(G) \geq 1-O(\varepsilon \log k)$. The intermediate relaxation satisfies $\mathrm{sdp}_+(G^\ell) \geq \mathrm{sdp}_+(G)^\ell$ (Lemma 9). Hence, $\mathrm{sdp}_+(G^\ell) \geq 1-O(\ell\varepsilon \log k)$. On the other hand, Theorem 8 implies that $\mathrm{opt}(G^\ell) \geq 1-O(\sqrt{\eta})$ if $\mathrm{sdp}_+(G^\ell) \geq 1-\eta$. Thus, we can conclude that $\mathrm{opt}(G^\ell) \geq 1 - O(\sqrt{\ell\varepsilon \log k})$.

In the next section (Section 1.2), we give a sketch of the proof of Theorem 3 and compare it to the proof of the previous bounds by [2]. We present a detailed proof of Theorem 3 in Section 3.

We record another consequence of Theorem 3, which shows that the semidefinite value of parallel repeated unique games is a good approximation of the optimal value (assuming that the alphabet size of the underlying unique game is not too large). (We omit the proof.)

Theorem 4. *Let $H = G^\ell$ be a parallel repetition of a unique game G with alphabet size k. Suppose* $\mathrm{sdp}(H) \geq 1-\varepsilon$*. Then,* $\mathrm{opt}(H) \geq 1 - O(\sqrt{\varepsilon \log k})$*.*

We remark that [3] show the same bound with k being replaced by the alphabet size of H. In our setting, H has alphabet size k^ℓ. (The bound of [3] is for general unique games H and does not assume that H is obtained by parallel repetition.) For $\ell = 1$, the two bounds agree. For more repetitions, our bound is strictly stronger. In this sense, parallel repeated unique games are easier to approximate than general unique games.

1.2 Proof Overview and Techniques

Our proof of Theorem 3 closely follows the approach of [2]. It is illustrative to start with an outline of this approach. (This discussion assumes that the reader is somewhat familiar with unique games and the relaxation $\mathrm{sdp}(G)$. See Section 2.1 and Section 2.2 for formal definitions.) Let G be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$. The unique game G is represented as a list of triples (u, v, π), where $u, v \in V$ and π is a permutation of Σ. The triple (u, v, π) encodes the constraint that the labels of u and v satisfy the relation π. In other words, an assignment $L \in \Sigma^V$ satisfies the constraint (u, v, π) if $L_v = \pi(L_u)$.

In the relaxation $\mathrm{sdp}(G)$, we consider vector-valued assignments instead of the usual assignments. More precisely, we assign orthogonal vectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k$

to every vertex $u \in V$. The vectors are normalized such that $\sum_i \|\boldsymbol{u}_i\|^2 = 1$ for every $u \in V$. Such a vector-valued assignment satisfies a constraint (u, v, π) if $\boldsymbol{u}_i$ is close to $\boldsymbol{v}_{\pi(i)}$ for most labels $i \in \Sigma$. Formally, the violation of the constraint (u, v, π) is measured by $\sum_i \|\boldsymbol{u}_i - \boldsymbol{v}_{\pi(i)}\|^2$.

To prove Theorem 3, we need to transform an optimal solution for $\mathrm{sdp}(G)$ to a good solution for $\mathrm{sdp}_+(G)$. As the notation suggests, the relaxation $\mathrm{sdp}_+(G)$ is quite similar to $\mathrm{sdp}(G)$. Instead of assigning arbitrary orthogonal vectors to a vertex $u \in V$, we ask for orthogonal vectors with only nonnegative coordinates. Note that nonnegative vectors are orthogonal only if they are supported on disjoint sets of coordinates. It is notationally more convenient to talk about nonnegative functions instead of vectors with only nonnegative coordinates.

To summarize, the relaxation $\mathrm{sdp}_+(G)$ asks us to assign nonnegative functions $f_{u,1}, \ldots, f_{u,k}$ with disjoint supports to every vertex $u \in V$. As before, the functions are normalized such that $\sum_i \|f_{u,i}\|^2 = 1$ for every $u \in V$, and the violation of a constraint (u, v, π) is measured by $\sum_i \|f_{u,i} - f_{v,\pi(i)}\|^2$.

Let $\{\boldsymbol{u}_i\}_{u \in V, i \in \Sigma} \subseteq \mathbb{R}^d$ be an optimal solution for $\mathrm{sdp}(G)$. To avoid some technical details, let us assume that all vectors have the same length so that $\|\boldsymbol{u}_i\|^2 = 1/k$ for every $u \in V$ and $i \in \Sigma$. To construct a solution for $\mathrm{sdp}_+(G)$, Barak et al. consider the distributions $N(\bar{\boldsymbol{u}}_i, \sigma^2 I)$ on $\mathbb{R}^d$. (Here, $N(\bar{\boldsymbol{u}}_i, \sigma^2 I)$ denotes the standard d-dimensional Gaussian distribution centered at the unit vector in direction $\boldsymbol{u}_i$, with standard deviation σ in each coordinate.) A tentative solution for $\mathrm{sdp}_+(G)$ is constructed by letting $f_{u,i}$ be the square root of the density function of $N(\boldsymbol{u}_i, \sigma^2 I)$ (suitably normalized). It turns out that the new violations $\sum_i \|f_{u,i} - f_{v,\pi(i)}\|^2$ exceed the original violations $\sum_i \|\boldsymbol{u}_i - \boldsymbol{v}_{\pi(i)}\|^2$ by at most a factor $O(\sigma^{-2})$. However, the supports of the functions $f_{u,1}, \ldots, f_{u,k}$ are far from disjoint (in fact, all of them have the same support, namely $\mathbb{R}^d$). Hence, the idea is to massage the functions such that their supports become disjoint. The approach taken by Barak et al. is to restrict $f_{u,i}$ to the Voronoi cell of the vector $\boldsymbol{u}_i$ (the set of points of $\mathbb{R}^d$ that are closer to $\boldsymbol{u}_i$ than to any other vector $\boldsymbol{u}_j$). Since the vectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k$ are pairwise orthogonal, only a small portion of the L_2-mass of $f_{u,i}$ is outside of the Voronoi cell of $\boldsymbol{u}_i$. Concretely, if $f'_{u,i}$ is the restriction of $f_{u,i}$ to the Voronoi cell of $\boldsymbol{u}_i$, the truncated L_2-mass is bounded by $\sum_i \|f_{u,i} - f'_{u,i}\|^2 \leq k e^{-\Omega(1/\sigma^2)}$. By choosing σ appropriately, we can balance the additional violation due to this truncation and the initial multiplicative increase of the violations. The reason why this approach falls short of proving Theorem 3 is that the truncation causes an additive increase in the violations, whereas for Theorem 3 we can only afford a multiplicative increase of the violation. (We remark that this additive increase is not an artifact of the analysis but a property of the construction.)

Our contribution is a construction that avoids this additive increase of the violations. We refer to this construction as smooth orthogonalization for nonnegative functions. We construct the functions $f'_{u,i}$ in the following way: As before, $f'_{u,i}$ is identically 0 outside of the Voronoi cell of $\boldsymbol{u}_i$. Inside the Voronoi cell of $\boldsymbol{u}_i$, we define $f'_{u,i}(\boldsymbol{x}) = f_{u,i}(\boldsymbol{x}) - f_{u,j}(\boldsymbol{x})$, where j is the label of the vector $\boldsymbol{u}_j$ that is second-smallest distance to $\boldsymbol{x}$. (In other words, we consider a refinement of the Voronoi partition according to second nearest neighbors.) With this construction, we can write $f'_{u,i}(\boldsymbol{x})$ as a piecewise-linear function of the values $f_{u,1}(\boldsymbol{x}), \ldots, f_{u,k}(\boldsymbol{x})$, i.e.,

$$f'_{u,i}(\boldsymbol{x}) = \max \left\{ f_{u,i}(\boldsymbol{x}) - f_{u,1}(\boldsymbol{x}), \ldots, f_{u,i}(\boldsymbol{x}) - f_{u,k}(\boldsymbol{x}) \right\} .$$

Since such piecewise-linear functions are Lipschitz, it follows that there exists a number L such that $\sum_i \|f'_{u,i} - f'_{v,\pi(i)}\|^2 \leq L \sum_i \|f_{u,i} - f_{v,\pi(i)}\|^2$ for all $u, v \in V$ and permutations π of Σ. It follows that this construction causes only a multiplicative increase of the violations. The remaining problem is to bound the Lipschitz constant L. A priori, L could grow with k. In our case, the Lipschitz constant L is bounded by an absolute constant independent of k. The reason is roughly that at every point $\boldsymbol{x} \in \mathbb{R}^d$, at most four of the values $f_{u,1}(\boldsymbol{x}), \ldots, f_{u,k}(\boldsymbol{x}), f_{v,1}(\boldsymbol{x}), \ldots, f_{v,k}(\boldsymbol{x})$ contribute to the distance $\sum_i \|f'_{u,i} - f'_{v,\pi(i)}\|^2$. See Section 3.1 for details.

2 Preliminaries

2.1 Unique Games and Parallel Repetition

Let V and Σ be two (finite) sets. A *unique game* G with *vertex set* V and *alphabet* Σ is defined by a distribution over triples (u, v, π), where $u, v \in V$ are vertices and $\pi\colon \Sigma \to \Sigma$ is a permutation of the alphabet Σ. We refer to the ℓ-fold product distribution G^ℓ as the *ℓ-fold repetition* of G. Note that G^ℓ is a unique game with vertex set V^ℓ and alphabet Σ^ℓ.

The *(optimal) value of* G is defined by

$$\operatorname{opt}(G) \stackrel{\text{def}}{=} \max_{L \in \Sigma^V} \mathop{\mathbf{P}}_{(u,v,\pi)\sim G} \left\{ L_v = \pi(L_u) \right\} . \tag{2.1}$$

The *amortized value of* G is defined by

$$\overline{\operatorname{opt}}(G) \stackrel{\text{def}}{=} \sup_{\ell \in \mathbb{N}} \operatorname{opt}(G^\ell)^{1/\ell} . \tag{2.2}$$

Theorem 5 ([17]). *If G is a unique game with* $\operatorname{opt}(G) \leq 1 - \eta$*, then* $\overline{\operatorname{opt}}(G) \leq 1 - \Omega(\eta^2)$*.*

2.2 Semidefinite and Nonnegative Relaxation

The *semidefinite value* of a unique game G with vertex set V and alphabet Σ is defined by

$$\operatorname{sdp}(G) \stackrel{\text{def}}{=} \max \mathop{\mathbf{E}}_{(u,v,\pi)\sim G} \textstyle\sum_{i\in\Sigma} \langle \boldsymbol{u}_i, \boldsymbol{v}_{\pi(i)} \rangle , \tag{2.3}$$

where we maximize over all collections $\{\boldsymbol{u}_i\}_{u \in V, i \in \Sigma}$ of vectors that satisfy

$$\textstyle\sum_{i\in\Sigma} \|\boldsymbol{u}_i\|^2 = 1 \qquad (u \in V) , \tag{2.4}$$

$$\langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle = 0 \qquad (u \in V,\ i \neq j \in \Sigma) . \tag{2.5}$$

The optimization problem $\mathrm{sdp}(G)$ is a relaxation of $\mathrm{opt}(G)$. Hence, $\mathrm{sdp}(G) \geq \mathrm{opt}(G)$ for every unique game G.

Theorem 6 ([3]). *Suppose G is a unique game with alphabet size k and $\mathrm{sdp}(G) \geq 1-\varepsilon$. Then, $\mathrm{opt}(G) \geq 1 - O(\sqrt{\varepsilon \log k})$.*

Theorem 7 ([6]). *For every unique game G and number $\ell \in \mathbb{N}$, we have $\mathrm{sdp}(G^\ell) = \mathrm{sdp}(G)^\ell$.*

Theorem 7 implies that $\mathrm{sdp}(G) \geq \overline{\mathrm{opt}}(G)$, i.e., $\mathrm{sdp}(G)$ is also a relaxation for $\overline{\mathrm{opt}}(G)$. Combining this fact with Theorem 6, we get that $\overline{\mathrm{opt}}(G) \leq 1 - \Omega(\eta^2/\log k)$ if G is a unique game with alphabet size k and $\mathrm{opt}(G) \leq 1-\eta$. Note that Theorem 5 shows that the $\log k$ factor in this bound is not necessary.

The *Hellinger*[1] *value* of G is defined as

$$\mathrm{sdp}_+(G) \stackrel{\mathrm{def}}{=} \max \mathop{\mathbf{E}}_{(u,v,\pi)\sim G} \textstyle\sum_{i\in\Sigma} \langle f_{u,i}, f_{v,\pi(i)} \rangle \,, \tag{2.6}$$

where we maximize over all collections $\{f_{u,i}\}_{u\in V, i\in\Sigma}$ of nonnegative functions on Ω such that

$$\textstyle\sum_{i\in\Sigma} \|f_{u,i}\|^2 = 1 \qquad (u \in V)\,, \tag{2.7}$$

$$\mathrm{supp}(f_{u,i}) \cap \mathrm{supp}(f_{u,j}) = \emptyset \qquad (u \in V,\ i \neq j \in \Sigma)\,. \tag{2.8}$$

Here, (Ω, μ) is some probability space and the norms and inner products for functions $f, g \colon \Omega \to \mathbb{R}$ are defined as $\langle f, g\rangle := \int_\Omega fg \,\mathrm{d}\mu$ and $\|f\| := \langle f, f\rangle^{1/2}$. Without loss of generality, we could assume $\Omega = [0,1]$ and that μ is the usual Lebesgue measure.

Notice that (2.8) is equivalent to the constraint that $f_{u,i}$ and $f_{u,j}$ are orthogonal for all $i \neq j$ and $u \in V$ (at least if Ω is finite and every atom has positive probability mass). Hence, the optimization problem $\mathrm{sdp}_+(G)$ is equivalent to $\mathrm{sdp}(G)$ except for the constraint that the value of the functions (coordinates of the vectors) are nonnegative.

Our result relies on the following theorem of [2].

Theorem 8 ([2]). *If G is a unique game with $\mathrm{sdp}_+(G) \geq 1-\varepsilon$, then $\mathrm{opt}(G) \geq 1 - 2\sqrt{2\varepsilon}$.*

Furthermore, we can lower bound $\mathrm{sdp}_+(G^\ell)$ in terms of $\mathrm{sdp}_+(G)$ as expected. This lemma follows from the fact that we can construct a solution for $\mathrm{sdp}_+(G^\ell)$ by taking appropriate tensor products of the functions $f_{u,i}$ that form an optimal solution for $\mathrm{sdp}_+(G)$.

Lemma 9. For every unique game G, we have $\mathrm{sdp}_+(G^\ell) \geq \mathrm{sdp}_+(G)^\ell$.

[1] The relaxation $\mathrm{sdp}_+(G)$ has an alternative description in terms minimizing squared Hellinger distances of jointly distributed random variables subject to a certain set of constraints (called distributional strategies in [2]). See [2] for details about this alternative description. The two formulations of $\mathrm{sdp}_+(G)$ are equivalent.

2.3 Mapping Vectors to Nonnegative Functions

Let $N(0,\sigma^2)^d$ be the Gaussian measure on $\mathbb{R}^d$ with mean 0 and covariance $\sigma^2 I$ (each coordinate is independent Gaussian with mean 0 and standard deviation σ). Let $\phi_\sigma \colon \mathbb{R}^d \to \mathbb{R}_+$ be the density of the measure $N(0,\sigma^2)^d$ with respect to the usual Lebesgue measure λ^d on $\mathbb{R}^d$,

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

Let $L_2(\mathbb{R}^d)$ be the (Hilbert) space of functions $f \colon \mathbb{R}^d \to \mathbb{R}$ such that $\int_{\mathbb{R}^d} f^2 \, \mathrm{d}\lambda^d$ is bounded. The inner product in $L_2(\mathbb{R}^d)$ is given $\langle f, g\rangle = \int fg \, \mathrm{d}\lambda^d$.

Let $T_{\boldsymbol{u}}$ be the translation operator on $L_2(\mathbb{R}^d)$, so that $T_{\boldsymbol{u}} f(\boldsymbol{x}) = f(\boldsymbol{x} - \boldsymbol{u})$. Barak et al. [2] consider the following mapping M_σ from $\mathbb{R}^d$ to nonnegative functions in $L_2(\mathbb{R}^d)$,

$$M_\sigma(\boldsymbol{u}) \stackrel{\text{def}}{=} \|\boldsymbol{u}\| \sqrt{T_{\bar{\boldsymbol{u}}} \phi_\sigma} \,. \tag{2.9}$$

Here, $\bar{\boldsymbol{u}}$ denotes the unit vector in the direction of $\boldsymbol{u}$. The mapping M_σ preserves norms, that is, $\|M_\sigma(\boldsymbol{u})\| = \|\boldsymbol{u}\|$ for every vector $\boldsymbol{u} \in \mathbb{R}^n$. We need the following additional properties of M_σ. (See Section A.2 for a proof of this lemma.)

Lemma 10 ([2]). For any two vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^d$,

$$\|M_\sigma(\boldsymbol{u}) - M_\sigma(\boldsymbol{v})\|^2 \le O(\sigma^{-2}) \cdot \|\boldsymbol{u} - \boldsymbol{v}\|^2 .$$

Furthermore, $\langle M_\sigma(\boldsymbol{u}), M_\sigma(\boldsymbol{v})\rangle = \|\boldsymbol{u}\| \, \|\boldsymbol{v}\| \cdot e^{-1/4\sigma^2}$ if $\boldsymbol{u}$ and $\boldsymbol{v}$ are orthogonal.

3 Improved Rounding of Repeated Unique Games

In this section, we will prove the following theorem — our main result.

Theorem (Restatement of Theorem 3). For every unique G with alphabet size k and $\mathrm{sdp}(G) \ge 1 - \varepsilon$,

$$\mathrm{sdp}_+(G) \ge 1 - O(\varepsilon \log k) .$$

The key ingredients of the proof of Theorem 3 are the mapping M_σ from vectors to nonnegative functions (Section 2.3) and the following smooth orthogonalization procedure for nonnegative functions (or vectors).

Lemma 11 (Smooth Nonnegative Orthogonalization). There exists a mapping $Q \colon L_2(\mathbb{R}^d)^k \to L_2(\mathbb{R}^d)^k$ with the following properties: Let $f_1, \dots, f_k$ and $g_1, \dots, g_k$ be nonnegative functions in $L_2(\mathbb{R}^d)$ such that $\sum_i \|f_i\|^2 = \sum_i \|g_i\|^2 = 1$ and $\sum_{i \ne j} \langle f_i, f_j \rangle + \sum_{i \ne j} \langle g_i, g_j \rangle \le \gamma$. Suppose $(f'_1, \dots, f'_k) = Q(f_1, \dots, f_k)$ and $(g'_1, \dots, g'_k) = Q(g_1, \dots, g_k)$. Then, for every permutation π of $[k]$,

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

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

We remark that Lemma 11 does not depend in any way on the fact that the functions f_i and g_i are defined on $\mathbb{R}^d$. (In fact, we could formulate the same lemma for finite dimensional vectors with nonnegative coordinates. We choose to state and prove the lemma in this form because the proof of Theorem 3 naturally leads to nonnegative functions defined on $\mathbb{R}^d$.) We prove Lemma 11 at the end of this section. Assuming this lemma, we can prove Theorem 3 as follows:

Proof (Theorem 3). Let G be a unique game with vertex set V and alphabet $\Sigma = [k]$. Suppose $\operatorname{sdp}(G) \geq 1 - \varepsilon$. For a parameter $\sigma > 0$, which we determine later, let $f_{u,i} := M_\sigma(\boldsymbol{u}_i)$ be the nonnegative functions in $L_2(\mathbb{R}^d)$ obtained by applying M_σ to a collection of vectors $\{\boldsymbol{u}_i\}_{u \in V, i \in \Sigma}$ corresponding to an optimal solution for $\operatorname{sdp}(G)$. Since M_σ preserves norms, we have for every $u \in V$,

$$\sum\nolimits_{i \in \Sigma} \|f_{u,i}\|^2 = \sum\nolimits_{i \in \Sigma} \|\boldsymbol{u}_i\|^2 = 1\,. \tag{3.1}$$

Since $\langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle = 0$ for $i \neq j$, Lemma 10 also shows that for every $u \in V$,

$$\begin{aligned} \sum_{i \neq j} \langle f_{u,i}, f_{u,j} \rangle &\leq \sum_{i \neq j} \|\boldsymbol{u}_i\|\,\|\boldsymbol{u}_j\| \cdot e^{-1/4\sigma^2} \\ &\leq e^{-1/4\sigma^2} \Big(\sum_{i \in \Sigma} \|\boldsymbol{u}_i\| \Big)^2 \leq k \cdot e^{-1/4\sigma^2} \sum_{i \in \Sigma} \|\boldsymbol{u}_i\|^2 \leq k \cdot e^{-1/4\sigma^2}\,. \end{aligned} \tag{3.2}$$

Hence, for $\sigma^2 = 1/(4 \log(k/\gamma))$, we can make sure that $\sum_{i \neq j} \langle f_{u,i}, f_{u,j} \rangle \leq \gamma$ for every $u \in V$. By Lemma 10, we have

$$\mathop{\mathbf{E}}_{(u,v,\pi) \sim G} \|f_{u,i} - f_{v,\pi(i)}\|^2 \leq O(\sigma^{-2}) \mathop{\mathbf{E}}_{(u,v,\pi) \sim G} \|\boldsymbol{u}_i - \boldsymbol{v}_{\pi(i)}\|^2 \leq O(\sigma^{-2})\varepsilon\,. \tag{3.3}$$

For the last inequality, we used the assumption $\operatorname{sdp}(G) \geq 1 - \varepsilon$ of Theorem 3 and the fact that the vectors $\{\boldsymbol{u}_i\}_{u \in V, i \in \Sigma}$ form an optimal solution for $\operatorname{sdp}(G)$. The functions $\{f_{u,i}\}_{u \in V, i \in \Sigma}$ form an approximate solution for $\operatorname{sdp}_+(G)$, in the sense of (3.1)–(3.3). Using the smooth nonnegative orthogonalization $Q\colon L_2(\mathbb{R}^d)^k \to L_2(\mathbb{R}^d)^k$ from Lemma 11, we obtain nonnegative functions $f'_{u,i} = Q(f_{u,1}, \ldots, f_{u,k})_i$ that form a feasible solution for $\operatorname{sdp}_+(G)$. Combining (3.3) and Lemma 11 also shows that for small enough γ (say $\gamma = 1/8$),

$$\mathop{\mathbf{E}}_{(u,v,\pi) \sim G} \|f'_{u,i} - f'_{v,\pi(i)}\|^2 \leq O(1) \mathop{\mathbf{E}}_{(u,v,\pi) \sim G} \|f_{u,i} - f_{v,\pi(i)}\|^2 \leq O(\sigma^{-2})\varepsilon\,.$$

Since we chose $\sigma^2 = \Omega(1/\log k)$, it follows that $\operatorname{sdp}_+(G) \geq 1 - O(\varepsilon \log k)$.

3.1 Proof of Lemma 11 (Smooth Nonnegative Orthogonalization)

First, we consider the following orthogonalization step,

$$Q^{(1)}\colon L_2(\mathbb{R}^d)^k \to L_2(\mathbb{R}^d)^k, \quad (f_1, \ldots, f_k) \mapsto (f'_1, \ldots, f'_k)\,, \tag{3.4}$$

$$f'_i(\boldsymbol{x}) = \begin{cases} f_i(\boldsymbol{x}) - \max_{j \neq i} f_j(\boldsymbol{x}) & \text{if } f_i(\boldsymbol{x}) > \max_{j \neq i} f_j(\boldsymbol{x})\,, \\ 0 & \text{otherwise.} \end{cases} \tag{3.5}$$

Next, we consider the following renormalization step,

$$Q^{(2)}\colon L_2(\mathbb{R}^d)^k \to L_2(\mathbb{R}^d)^k\,, \tag{3.6}$$

$$\Big(f_1,\ldots,f_k\Big) \mapsto \left(\tfrac{1}{\lambda}f_1,\ldots,\tfrac{1}{\lambda}f_k\right)\,, \tag{3.7}$$

$$\text{where } \lambda^2 = \textstyle\sum_i \|f_i\|^2\,. \tag{3.8}$$

To prove Lemma 11, we choose Q as the composition of $Q^{(1)}$ and $Q^{(2)}$. Let $f_1,\ldots,f_k$ and $g_1,\ldots,g_k$ be nonnegative functions in $L_2(\mathbb{R}^d)$ as in Lemma 11, i.e., $\sum_{i\neq j}\langle f_i,f_j\rangle + \sum_{i\neq}\langle g_i,g_j\rangle \le \gamma$ and $\sum_i\|f_i\|^2 = \sum_i\|g_i\|^2 = 1$.

Let $(f_1^{(1)},\ldots,f_k^{(1)}) = Q^{(1)}(f_1,\ldots,f_k)$ and $(f_1^{(2)},\ldots,f_k^{(2)}) = Q^{(2)}(f_1^{(1)},\ldots,f_k^{(1)})$. Similarly, let $(g_1^{(1)},\ldots,g_k^{(1)}) = Q^{(1)}(g_1,\ldots,g_k)$ and $(g_1^{(2)},\ldots,g_k^{(2)}) = Q^{(2)}(g_1^{(1)},\ldots,g_k^{(1)})$. In the rest of the section, we first establish several properties of these functions (see Claim 12 and Claim 13) and then use these properties to prove Lemma 11.

Claim 12 (Properties of $Q^{(1)}$).

1. For every permutation π of $[k]$,

$$\sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2 \le 8\sum_i \|f_i - g_{\pi(i)}\|^2\,.$$

2. For all $i\neq j$,

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

3.

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

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

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

To prove Item 1, we will show that for every $x\in\mathbb{R}^d$,

$$\sum_i \left(f_i^{(1)}(\boldsymbol{x}) - g_{\pi(i)}^{(1)}(\boldsymbol{x})\right)^2 \le 8\sum_i \left(f_i(\boldsymbol{x}) - g_{\pi(i)}(\boldsymbol{x})\right)^2\,. \tag{3.9}$$

Since $Q^{(1)}$ is invariant under permutation of its inputs, we may assume π is the identity permutation. At this point, we can verify (3.9) by an exhaustive case distinction. Fix $\boldsymbol{x}\in\mathbb{R}^d$. Let i_f be the index i that maximizes $f_i(\boldsymbol{x})$. (We may assume the maximizer is unique.) Let j_f be the index such that $f_{j_f}(\boldsymbol{x}) = \max_{j\neq i_f} f_j(\boldsymbol{x})$. Similarly, define i_g and j_g such that $g_{i_g}(\boldsymbol{x}) = \max_i g_i(\boldsymbol{x})$ and $g_{j_g}(\boldsymbol{x}) = \max_{j\neq i_g} g_j(\boldsymbol{x})$. We may assume that $i_f = 1$ and $j_f = 2$. Furthermore, we may assume $i_g, j_g \in \{1,2,3,4\}$. Notice that the sum on the left-hand side of

(3.9) has at most two non-zero terms (corresponding to the indices $i \in \{i_f, i_g\} \subseteq \{1,\ldots,4\}$). Hence, to verify (3.9), it is enough to show

$$\max_{i\in\{1,\ldots,4\}} \left|f_i^{(1)}(\boldsymbol{x}) - g_i^{(1)}(\boldsymbol{x})\right| \le 4 \max_{i\in\{1,\ldots,4\}} \left|f_i(\boldsymbol{x}) - g_i(\boldsymbol{x})\right| . \tag{3.10}$$

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

$$\left|f_i^{(1)}(\boldsymbol{x}) - g_i^{(1)}(\boldsymbol{x})\right| = \left|q_i\big(f_1(\boldsymbol{x}),\ldots,f_4(\boldsymbol{x})\big) - q_i\big(g_1(\boldsymbol{x}),\ldots,g_4(\boldsymbol{x})\big)\right| \le 4\varepsilon .$$

Claim 13 (Properties of $Q^{(2)}$).

1. For every permutation π of $[k]$,
$$\sum_i \|f_i^{(2)} - g_{\pi(i)}^{(2)}\|^2 \le \tfrac{4}{1-4\gamma} \sum_i \|f_i^{(1)} - g_{\pi(i)}^{(1)}\|^2 .$$
2. For all $i \in \Sigma$,
$$\operatorname{supp}(f_i^{(2)}) = \operatorname{supp}(f_i^{(1)}) .$$
3. $$\sum_i \|f_i^{(2)}\|^2 = 1 .$$

Proof. Again Item 2 and Item 3 follow immediately by definition of the mapping $Q^{(2)}$. To prove Item 1, let $\lambda_f, \lambda_g > 0$ be the multipliers such that $f_i^{(2)} = f_i^{(1)}/\lambda_f$ and $g_i^{(2)} = g_i^{(1)}/\lambda_g$ for all $i \in [k]$. Item 1 of Claim 12 shows that λ_f^2 and λ_g^2 lie in the interval $[1-2\gamma, 1]$. We estimate the distances between $f_i^{(2)}$ and $g_{\pi(i)}^{(2)}$ as follows,

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

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

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

Combining the previous two estimates, we get as desired

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

Combining Claim 12 and Claim 13 yields Lemma 11.

Lemma (Restatement of Lemma 11). There exists a mapping $Q\colon L_2(\mathbb{R}^d)^k \to L_2(\mathbb{R}^d)^k$ with the following properties: Let $f_1, \ldots, f_k$ and $g_1, \ldots, g_k$ be nonnegative functions in $L_2(\mathbb{R}^d)$ such that $\sum_i \|f_i\|^2 = \sum_i \|g_i\|^2 = 1$ and $\sum_{i \neq j} \langle f_i, f_j \rangle + \sum_{i \neq j} \langle g_i, g_j \rangle \leq \gamma$. Suppose $(f_1', \ldots, f_k') = Q(f_1, \ldots, f_k)$ and $(g_1', \ldots, g_k') = Q(g_1, \ldots, g_k)$. Then, for every permutation π of $[k]$,

$$\sum_i \|f_i' - g_{\pi(i)}'\|^2 \leq \tfrac{32}{1-4\gamma} \sum_i \|f_i - g_{\pi(i)}\|^2 .$$

Furthermore, $\sum_i \|f_i'\|^2 = \sum_i \|g_i'\|^2 = 1$ and $\mathrm{supp}(f_i') \cap \mathrm{supp}(f_j') = \mathrm{supp}(g_i') \cap \mathrm{supp}(g_j') = \emptyset$ for all $i \neq j$.

Proof. We choose $Q = Q^{(2)} \circ Q^{(1)}$ as the composition of $Q^{(1)}$ and $Q^{(2)}$. In this case, $f_i' = f_i^{(2)}$ and $g_i' = g_i^{(2)}$, where the functions $f_1^{(2)}, \ldots, f_k^{(2)}$ and $g_1^{(2)}, \ldots, g_k^{(2)}$ are constructed as in the beginning of Section 3.1. Combining Item 1 of Claim 12 and Item 1 of Claim 13 gives the desired upper bound on $\sum_i \|f_i' - g_{\pi(i)}'\|^2$. Similarly, the remaining properties desired of f_i' and g_i' also follow by combining Claim 12 and Claim 13.

Acknowledgments

Thanks to Oded Regev for discussions and helpful comments on an earlier version of this manuscript. Thanks to Sanjeev Arora, Rajat Mittal, and Mario Szegedy for valuable discussions.

References

1. Austrin, P.: Towards sharp inapproximability for any 2-CSP. In: FOCS. pp. 307–317 (2007)
2. Barak, B., Hardt, M., Haviv, I., Rao, A., Regev, O., Steurer, D.: Rounding parallel repetitions of unique games. In: FOCS. pp. 374–383 (2008)
3. Charikar, M., Makarychev, K., Makarychev, Y.: Near-optimal algorithms for unique games. In: STOC. pp. 205–214 (2006)
4. Chawla, S., Krauthgamer, R., Kumar, R., Rabani, Y., Sivakumar, D.: On the hardness of approximating multicut and sparsest-cut. Computational Complexity 15(2), 94–114 (2006)
5. Feige, U., Kindler, G., O'Donnell, R.: Understanding parallel repetition requires understanding foams. In: IEEE Conference on Computational Complexity. pp. 179–192 (2007)
6. Feige, U., Lovász, L.: Two-prover one-round proof systems: Their power and their problems (extended abstract). In: STOC. pp. 733–744 (1992)

7. Guruswami, V., Manokaran, R., Raghavendra, P.: Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In: FOCS. pp. 573–582 (2008)
8. Holenstein, T.: Parallel repetition: Simplification and the no-signaling case. Theory of Computing 5(1), 141–172 (2009)
9. Khot, S.: On the power of unique 2-prover 1-round games. In: STOC. pp. 767–775 (2002)
10. Khot, S., Kindler, G., Mossel, E., O'Donnell, R.: Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? SIAM J. Comput. 37(1), 319–357 (2007)
11. Khot, S., Regev, O.: Vertex cover might be hard to approximate to within $2-\varepsilon$. J. Comput. Syst. Sci. 74(3), 335–349 (2008)
12. Khot, S., Vishnoi, N.K.: The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In: FOCS. pp. 53–62 (2005)
13. Manokaran, R., Naor, J., Raghavendra, P., Schwartz, R.: SDP gaps and UGC-hardness for multiway cut, 0-extension, and metric labeling. In: STOC. pp. 11–20 (2008)
14. Mossel, E., O'Donnell, R., Oleszkiewicz, K.: Noise stability of functions with low influences invariance and optimality. In: FOCS. pp. 21–30 (2005)
15. Raghavendra, P.: Optimal algorithms and inapproximability results for every CSP? In: STOC. pp. 245–254 (2008)
16. Raghavendra, P., Steurer, D.: Graph expansion and the unique games conjecture. In: STOC (2010), to appear
17. Rao, A.: Parallel repetition in projection games and a concentration bound. In: STOC. pp. 1–10 (2008)
18. Raz, R.: A parallel repetition theorem. SIAM J. Comput. 27(3), 763–803 (1998)
19. Raz, R.: A counterexample to strong parallel repetition. In: FOCS. pp. 369–373 (2008)

A Further Proofs

A.1 Relation of Amortized Value and Semidefinite Value

Theorem (Restatement of Theorem 2). For every unique game G with alphabet size k,

$$\mathrm{sdp}(G)^{O(\log k)} \leq \overline{\mathrm{opt}}(G) \leq \mathrm{sdp}(G)\,.$$

Proof. Feige and Lovász [6] show the upper bound on $\overline{\mathrm{opt}}(G)$. (In particular, they show $\mathrm{opt}(G^\ell) \leq \mathrm{sdp}(G^\ell)$ and $\mathrm{sdp}(G^\ell) = \mathrm{sdp}(G)^\ell$.) Charikar, Makarychev, and Makarychev [3] show that $\mathrm{opt}(G) \geq \mathrm{sdp}(G)^{-C\log k}/C$ for some absolute constant $C \geq 1$. (This constant C is necessarily larger than 1.) Hence, we are done if $\mathrm{sdp}(G)^{C'\log k} \leq 1/C$ for some absolute constant $C' \geq 1$. (In this case, we would have $\mathrm{opt}(G) \geq \mathrm{sdp}(G)^{-(C+C')\log k}$.)

On the other hand, if $\mathrm{sdp}(G) \geq 1-\varepsilon$, then Theorem 1 shows that $\mathrm{opt}(G^\ell) \geq 1 - C''\sqrt{\ell\varepsilon\log k}$ for some absolute constant $C'' \geq 1$. Hence, if $\varepsilon \log k \leq 1/(2C'')^2$, then we can find a natural number $\ell = \Omega(1/(\varepsilon\log k))$ such that $\mathrm{opt}(G^\ell) \geq 1/2$, which implies $\overline{\mathrm{opt}}(G) \geq 2^{-1/\ell} \geq (1-\varepsilon)^{O(\log k)}$. It is straight-forward to check that there exists an absolute constant $C' \geq 1$ (depending on C and C'') such that $\mathrm{sdp}(G)^{C'\log k} > 1/C$ implies that $\varepsilon\log k \leq 1/(2C'')^2$. We conclude that in all cases $\overline{\mathrm{opt}}(G) \geq \mathrm{sdp}(G)^{O(\log k)}$.

A.2 Mapping Unit Vectors to Nonnegative Functions

Recall the definition of the mapping $M_\sigma\colon \mathbb{R}^d \to L_2(\mathbb{R}^d)$, $M_\sigma(\boldsymbol{u}) := \|\boldsymbol{u}\|\sqrt{T_{\bar{\boldsymbol{u}}}\phi_\sigma}$. Here, $T_{\boldsymbol{u}}$ is the translation operator on $L_2(\mathbb{R}^d)$, so that $T_{\boldsymbol{u}} f(\boldsymbol{x}) = f(\boldsymbol{x}-\boldsymbol{u})$, and ϕ_σ is the density of the Gaussian measure $N(0,\sigma^2)^d$ with respect to the usual Lebesgue measure λ^d on $\mathbb{R}^d$, $\phi_\sigma(\boldsymbol{x}) := (\sigma\sqrt{2\pi})^{-d} e^{-\|\boldsymbol{x}\|^2/2\sigma^2}$. From the definition of M_σ, it follows that the mapping preserves norms, so that $\|M_\sigma(\boldsymbol{u})\| = \|\boldsymbol{u}\|$ for every $\boldsymbol{u} \in \mathbb{R}^d$. The following fact about the (Hellinger) affinity of translated Gaussians shows that the mapping M_σ also preserves angles (at least approximately).

Lemma 14 ([2]). Let $\boldsymbol{u}$ and $\boldsymbol{v}$ be two unit vectors in $\mathbb{R}^d$. Then

$$\int_{\mathbb{R}^d} \sqrt{T_{\boldsymbol{u}}\phi_\sigma \cdot T_{\boldsymbol{v}}\phi_\sigma}\,\mathrm{d}\lambda^d = e^{-\|\boldsymbol{u}-\boldsymbol{v}\|^2/8\sigma^2}\,.$$

Proof. Immediate from the identity $\sqrt{T_{\boldsymbol{u}}\phi_\sigma \cdot T_{\boldsymbol{v}}\phi_\sigma} = e^{-\|\boldsymbol{u}-\boldsymbol{v}\|^2/8\sigma^2} T_{\frac{1}{2}(\boldsymbol{u}+\boldsymbol{v})}\phi_\sigma$.

The following technical fact shows that in order for a mapping to preserve distances it is enough to preserve lengths and distances of unit vectors (angles).

Fact 15. For any two vectors $\boldsymbol{u},\boldsymbol{v} \in \mathbb{R}^n$, we have

$$\|\boldsymbol{u}-\boldsymbol{v}\|^2 = (\|\boldsymbol{u}\| - \|\boldsymbol{v}\|)^2 + \|\boldsymbol{u}\|\,\|\boldsymbol{v}\| \cdot \|\bar{\boldsymbol{u}} - \bar{\boldsymbol{v}}\|^2\,.$$

Combining Fact 15 and Lemma 14 yields Lemma 10.

Lemma (Restatement of Lemma 10). For any two vectors $\boldsymbol{u},\boldsymbol{v} \in \mathbb{R}^d$,

$$\|M_\sigma(\boldsymbol{u}) - M_\sigma(\boldsymbol{v})\|^2 \le O(\sigma^{-2}) \cdot \|\boldsymbol{u}-\boldsymbol{v}\|^2\,.$$

Furthermore, $\langle M_\sigma(\boldsymbol{u}), M_\sigma(\boldsymbol{v})\rangle = \|\boldsymbol{u}\|\,\|\boldsymbol{v}\| \cdot e^{-1/4\sigma^2}$ if $\boldsymbol{u}$ and $\boldsymbol{v}$ are orthogonal.

Proof. If $\boldsymbol{u}$ and $\boldsymbol{v}$ are orthogonal, then Lemma 14 shows that $\langle M_\sigma(\boldsymbol{u}), M_\sigma(\boldsymbol{v})\rangle = \|\boldsymbol{u}\|\,\|\boldsymbol{v}\| e^{-1/4\sigma^2}$, because $\|\bar{\boldsymbol{u}} - \bar{\boldsymbol{v}}\|^2 = 2$ for any two orthogonal vectors $\boldsymbol{u}$ and $\boldsymbol{v}$. It remains to show the upper bound on $\|M_\sigma(\boldsymbol{u}) - M_\sigma(\boldsymbol{v})\|^2$. By construction, $M_\sigma(\boldsymbol{u}) = \|\boldsymbol{u}\| M_\sigma(\bar{\boldsymbol{u}})$ and $M_\sigma(\boldsymbol{v}) = \|\boldsymbol{v}\| M_\sigma(\bar{\boldsymbol{v}})$. By Fact 15,

$$\|M_\sigma(\boldsymbol{u}) - M_\sigma(\boldsymbol{v})\|^2 = (\|\boldsymbol{u}\| - \|\boldsymbol{v}\|)^2 + \|\boldsymbol{u}\|\,\|\boldsymbol{v}\| \cdot \|M_\sigma(\bar{\boldsymbol{u}}) - M_\sigma(\bar{\boldsymbol{v}})\|^2\,.$$

On the other, Lemma 14 implies that

$$\tfrac{1}{2}\|M_\sigma(\bar{\boldsymbol{u}}) - M_\sigma(\bar{\boldsymbol{v}})\|^2 = 1 - e^{-\|\bar{\boldsymbol{u}}-\bar{\boldsymbol{v}}\|^2/8\sigma^2} \le \|\bar{\boldsymbol{u}} - \bar{\boldsymbol{v}}\|^2/8\sigma^2\,.$$

(Here, we used the approximation $e^{-x} \ge 1-x$.) Combining these bounds, yields as desired

$$\|M_\sigma(\boldsymbol{u}) - M_\sigma(\boldsymbol{v})\|^2 \le (\|\boldsymbol{u}\| - \|\boldsymbol{v}\|)^2 + \|\boldsymbol{u}\|\,\|\boldsymbol{v}\| \cdot \|\bar{\boldsymbol{u}} - \bar{\boldsymbol{v}}\|^2/4\sigma^2 \le \tfrac{1}{4\sigma^2}\|\boldsymbol{u}-\boldsymbol{v}\|^2\,.$$

(The last inequality follows from Fact 15.)