THE RELATIONSHIP BETWEEN EXPANSION AND THE UNIQUE GAMES CONJECTURE

DOW LAFEVERS

ABSTRACT. In this paper, we introduce the Unique Games conjecture of Khot, survey a handful of its many implications in the hardness of approximation, and attempt to characterize the unexpected relationship between Unique Games and graph expansion. In particular, we provide a full proof of the first-known sufficient condition for the Unique Games Conjecture, the Small-Set Expansion Hypothesis, introduced by Raghavendra and Steurer. We only assume a working familiarity with classical complexity theory, discrete probability, and the theory of algorithms.

CONTENTS

1. INTRODUCTION

The *Unique Games Conjecture* (UGC) of Khot [Kho02] is one of the central open problems in complexity theory and in the hardness of approximation. The conjecture asserts that a certain constraint satisfaction problem is hard to approximate in a very strong sense. Recently, the Unique Games Conjecture has been shown to imply optimal inapproximability results for classic problems like MAX-CUT [KKMO07], VERTEX COVER [KR08], and SPARSEST CUT [KV05], and, perhaps most impressively, Raghavendra [Rag08] showed that, if the UGC is true, then every constraint satisfaction problem has a so-called *sharp approximation threshold*

Date: December 14, 2023.

τ: for all $\epsilon > 0$, one can achieve a $\tau - \epsilon$ approximation in polynomial (in fact quasilinear by a result of Steurer [Ste10b]) time, but obtaining a $\tau + \epsilon$ approximation is **NP**-hard. Beyond approximation algorithms, the UGC has been linked to the theory of metric embeddings and computational geometry, (discrete) Fourier analysis, and the study of parallel repetition. So the UGC certainly has profound implications. Of course, necessary conditions for the UGC do not yield any information about the truth of the conjecture. Indeed, unlike most of the major open problems used as hypotheses in theoretical computer science like $\mathbf{P} \neq \mathbf{NP}$ or $\mathbf{NP} \not\subset \bigcap_{\gamma>0} \mathrm{DTIME}(2^{n^\gamma})$ there is no consensus among researchers of the veracity of the conjecture.

Until recently, there were no known formal consequences of an algorithmic refutation of the UGC. In this paper, we demonstrate a first-of-its-kind "reverse" reduction, from the problem of approximating the expansion of small sets to Unique Games due to Steurer and Raghavendra [RS10]. The hypothesis that approximating the expansion of small sets is computationally hard turns out to imply the Unique Games Conjecture by way of this reduction. The crucial consequence of this reduction (stated formally in Theorem 5.1 is that a refutation will obtain an algorithm for approximating the edge expansion of graphs in a certain regime—a fundamental optimization problem. This paper details this reduction and its implications, and outlines the reduction used to show that the UGC with an additional assumption of mild expansion on the constraint graph is actually equivalent to the hardness of approximating small set expansion.

The theorems and proofs in this paper mostly follow those of [RS10], [Ste10a], and [RST10a].

2. Preliminaries

We consider undirected weighted graphs with self-loops allowed. Thus, we can identify such a *graph* G with vertex set V as a symmetric distribution over pairs ij with $i, j \in V$, where the edges of G are those pairs ij in the support of this distribution. The distribution is symmetric since our graphs are undirected.

We write $i \sim G$ to denote a random vertex of of G obtained by sampling according to the distribution given by the degrees, $ij \sim G$ to denote a (uniformly[1]) random edge of G and for a vertex $i \in V$, we write $j \sim G(i)$ to denote a random neighbor of i in G, which can be obtained by sampling a random edge of G conditioned on the event that the first endpoint of the edge is i and outputting the second endpoint of that edge.

2.1. Expansion and expansion profile.

For $S, T \subset V$, we define $E_G(S,T)$ as the fraction of edges going from S to T,

$$E_G(S,T) \stackrel{\text{def}}{=} \mathbb{P}_{ij\sim G}[i \in S, j \in T].$$

For a vertex set $S \subset V$, we define its *edge boundary* $\partial_G(S)$ as the fraction of edges leaving S (and going to $V - S$),

$$\partial_G(S) \stackrel{\text{def}}{=} E_G(S, V - S).$$

[1] Unless otherwise noted, in this paper random means sampled according the uniform distribution.

We define its *volume* $\mu_G(S)$ as the fraction of edges with a vertex in S,

$$\mu(S) \stackrel{\text{def}}{=} E_G(S, V).$$

The (edge) *expansion* $\Phi_G(S)$ for $S \neq \emptyset$ is the ratio between these two quantities,

$$\Phi_G(S) \stackrel{\text{def}}{=} \frac{\partial_G(S)}{\mu_G(S)}.$$

We set $\Phi_G(\varnothing) = \infty$ as we are ultimately trying to minimize the expansion.

The *expansion* (also called the conductance or the Cheeger's constant) Φ_G of a graph G is the minimum expansion over all sets with volume at most $\frac{1}{2}$,

$$\Phi_G \stackrel{\text{def}}{=} \min_{\substack{S \subset V \\ \mu_G(S) \leq \frac{1}{2}}} \Phi_G(S).$$

More generally, for $\delta \in [0, \frac{1}{2}]$, the *expansion at volume* δ, denoted by $\Phi_G(\delta)$, is the minimum expansion over sets with volume at most δ,

$$\Phi_G(\delta) \stackrel{\text{def}}{=} \min_{\substack{S \subset V \\ \mu_G(S) \leq \delta}} \Phi_G(S).$$

The curve $\delta \mapsto \Phi_G(\delta)$ is called the *expansion profile of* G.

We note two alternative characterizations of the expansion (equivalent to which we gave):

$$\begin{aligned}\Phi_G(S) &= \mathbb{P}_{ij \sim G}[j \notin S \mid i \in S] \\ &= \mathbb{E}_{i \sim G}[\mathbb{P}_{j \sim G(i)}[j \notin S] \mid i \in S]\end{aligned}$$

We will always drop the subscript G in the above quantities when the graph is obvious.

We can now state the Small-Set Expansion problem:

Problem 2.1 (GAP-SMALL-SET EXPANSION (η, δ))**.** Given a graph G and constants $\eta, \delta > 0$, distinguish whether

$$\Phi_G(\delta) \geq 1 - \eta \qquad \text{or} \qquad \Phi_G(\delta) \leq \eta.$$

We also state the associated conjecture, which we will see in Section 3 (Theorem 3.1) implies the Unique Games Conjecture:

Conjecture 2.1 (Small-Set Expansion Hypothesis (SSEH))**.** *Given $\eta > 0$, there exists $\delta = \delta(\epsilon) > 0$ such that the problem GAP-SMALL-SET EXPANSION (η, δ) is **NP**-hard.*

It's worth remarking that the best known polynomial-time approximation algorithm for Problem 2.1 (given in [RST10b]) is insufficient to refute this conjecture.

A little more should be said about what **NP**-hardness means in the context of this gap promise problem. Since SAT is **NP**-hard (i.e., the Cook-Levin theorem), we can express Conjecture 2.1 as the following: for every constant $\eta > 0$, there exists $\delta = \delta(\eta) > 0$ and a polynomial time (Cook) reduction from SAT to SMALL-SET EXPANSION such that

- Every satisfiable SAT instance reduces to a SMALL-SET EXPANSION instance with optimal value at least $1 - \eta$,
- Every unsatisfiable SAT instance reduces to a SMALL-SET EXPANSION instance with optimal value at most η.

2.2. **Unique Games and partial games.** For historical reasons, we state the version of Unique Games as originally formulated by Khot [Kho02]:

Definition 2.2. An *instance* of Unique Games represented as $\mathcal{U} = (\mathcal{V}, \mathcal{E}, \Pi, [R])$ consists of the following.

(1) A graph over vertex set $\mathcal{V}$ with edges $\mathcal{E}$ between them
(2) A set of label $[R] = \{1, \dots, R\}$
(3) A set of permutations $\pi_{u\leftarrow v} : [R] \to [R]$ for each edge $e = (u, v) \in \mathcal{E}$

An assignment $\mathcal{A} : \mathcal{V} \to [R]$ of labels to vertices is said to *satisfy* an edge $(u, v) \in \mathcal{E}$ if $\pi_{v\leftarrow u}(\mathcal{A}(w)) = \mathcal{A}(v)$. The objective is to find an assignment $\mathcal{A}$ that satisfies the maximum number of edges.

As is convenient in hardness of approximation, one defines a gap-version of the Unique Games problem as follows:

Problem 2.2 (UNIQUE GAMES $(R, 1-\epsilon, \eta)$)**.** Given a Unique Games instance $\mathcal{U} = (\mathcal{V}, \mathcal{E}, \Pi = \{\pi_{u\leftarrow v} : [R] \to [R] \mid e = (u, v) \in \mathcal{E}\}, [R])$ and constants $\epsilon, \eta, R > 0$, distinguish whether

(1) There exists an assignment $\mathcal{A}$ of labels that satisfies a $1-\epsilon$ fraction of the edges, or
(2) No assignment satisfies more than an η-fraction of the edges.

For the purposes of our reductions, it will be more convenient to interpret an instance of Unique Games as given by a distribution over the constraints:

Definition 2.3 (Unique Games)**.** A *Unique Game* $\mathcal{U}$ with vertex set $\mathcal{V}$ and alphabet $[R]$ is a distribution over *constraints* $(u, v, \pi) \in \mathcal{V} \times \mathcal{V} \times \mathbb{S}_{[R]}$, where $\mathbb{S}_{[R]}$ is the set of permutations of $[R]$. An *assignment* $x \in [R]^{\mathcal{V}}$ *satisfies* a constraint (u, v, π) if $x_v = \pi(x_u)$, i.e., π maps labels for u to labels for v. The *value* $\mathcal{U}(x)$ of an assignment x for $\mathcal{U}$ is the fraction of constraints of $\mathcal{U}$ satisfied by the assignment x, i.e.,

$$\mathcal{U}(x) \stackrel{\text{def}}{=} \mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[\pi(x_u) = x_v].$$

Finally, the *optimal value* $\operatorname{opt}(\mathcal{U})$ is defined as the maximum of $\mathcal{U}(x)$ over all x, i.e.,

$$\operatorname{opt}(\mathcal{U}) \stackrel{\text{def}}{=} \max_{x\in[R]^{\mathcal{V}}} \mathcal{U}(x).$$

We will assume that the distribution over constraints is symmetric in the sense that a constraint (u, v, π) has the same probability as (v, u, π^{-1}).

To this version of a Unique Game, we can associate two graphs. The *constraint graph* $G(\mathcal{U})$ is a graph with vertex set $\mathcal{V}$ with edge distribution given by the following sampling procedure:

(1) Sample a random constraint $(u, v, \pi) \sim \mathcal{U}$,
(2) Output the edge uv.

The *label-extended graph* $\hat{G}(\mathcal{U})$ is a graph with vertex set $\mathcal{V} \times [R]$ with edge distribution given by the following sampling procedure:

(1) Sample a random constraint $(u, v, \pi) \sim \mathcal{U}$,
(2) Sample a random label $i \in [R]$,
(3) Output an edge connecting (u, i) and $(v, \pi(i))$.

The rationale for using this probabilistic interpretation of Unique Games is the following: an assignment $x \in [R]^{\mathcal{V}}$ naturally corresponds to a set $S \subset \mathcal{V} \times [R]$

with cardinality $|S| = |\mathcal{V}|$ (and therefore volume $\mu(S) = 1/R$). The value of the assignment x for the Unique Game $\mathcal{U}$ corresponds exactly to the expansion of the corresponding set S in the label-extended graph $\hat{G}(\mathcal{U})$,

$$\mathcal{U}(x) = 1 - \Phi(S).$$

This correspondence between expansion and the value of an assignment is precisely the basis of the connection between UNIQUE GAMES and SMALL-SET EXPANSION discussed in the following section.

It will be useful to consider a more general version of Unique Games where partial assignments are allowed (for those familiar with 2-prover games, this situation corresponds to a prover refusing to answer one of the referees questions):

Definition 2.4. Let $\mathcal{U}$ be a Unique Game with vertex set $\mathcal{V}$ and alphabet $[R]$. An assignment $x \in ([R] \cup \{\perp\})^{\mathcal{V}}$ is *α-partial* if at least an α fraction of the vertices are labeled (with symbols from $[R]$), i.e.,

$$\mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[x_u \neq \perp] \geq \alpha.$$

A partial assignment x satisfies a constraint (u, v, π) if both vertices are labeled and their labels satisfy the constraint $x_v = \pi(x_u)$. For concision, we write $x_v = \pi(x_u) \in [R]$ to denote the event where the partial assignment x satisfies the constraint (u, v, π). The *value* $\mathcal{U}(x)$ of a partial assignment x is the fraction of constraints satisfied by x,

$$\mathcal{U}(x) \stackrel{\text{def}}{=} \mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[x_v = \pi(x_u) \in [R]].$$

The *α-partial value* $\mathrm{opt}_\alpha(\mathcal{U})$ is the maximum value of an α-partial assignment normalized by the fraction of labeled vertices,

$$\mathrm{opt}_\alpha(\mathcal{U}) \stackrel{\text{def}}{=} \max_{x\in([R]\cup\{\perp\})^{\mathcal{V}}} \left\{ \frac{\mathcal{U}(x)}{\mathbb{P}_{u\in\mathcal{V}}[x_u \neq \perp]} \;\middle|\; \mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[x_u \neq \perp] \geq \alpha \right\}.$$

Note that $\mathrm{opt}(\mathcal{U}) = \mathrm{opt}_1(\mathcal{U})$ and $\mathrm{opt}_\alpha(\mathcal{U}) \leq \mathrm{opt}_\beta(\mathcal{U})$ whenever $\alpha \geq \beta$.

2.3. **Functions on graphs.** For a graph G with vertex set V, we write $L_2(V)$ to denote the function space $\{f : V \to \mathbb{R}\}$ equipped with the (natural) inner product

$$\langle f, g\rangle \stackrel{\text{def}}{=} \mathbb{E}_{i\sim V} f(i)g(i).$$

This inner product induces the norm $\|f\| := \langle f, f\rangle^{1/2}$. We will also be interested in the norm $\|f\|_1 := \mathbb{E}_{i\in V}|f(i)|$.

We identify the graph G wth the following linear (Markov) operator on V:

$$Gf(i) \stackrel{\text{def}}{=} \mathbb{E}_{j\sim G(i)} f(j).$$

The matrix corresponding to this operator is the (weighted) adjacency matrix of G normalized so that every row sums to 1 (so that Gf is a (right) stochastic matrix). The operator G is self-adjoint with respect to the inner product on $L_2(V)$ since $\langle f, Gg\rangle = \mathbb{E}_{ij\sim G} f(i)g(j)$ so its eigenvalues are real and its eigenfunctions form an orthogonal basis of $L_2(V)$. We note the following identities, which hold for all vertex sets $S, T \subset V$

$$E(S,T) = \langle \mathbb{1}_S, G\mathbb{1}_T\rangle \quad \mu(S) = \|\mathbb{1}_S\|^2 \quad \partial(S) = \langle \mathbb{1}_S, L_G\mathbb{1}_S\rangle \quad \Phi(S) = \frac{\langle \mathbb{1}_S, L_G\mathbb{1}_S\rangle}{\|\mathbb{1}_S\|^2}$$

where $L_G := I - G$ is the *Laplacian* of G.

3. Small-Set Expansion implies Unique Games

In this section, we prove that the Unique Games Conjecture is true if the Small-Set Expansion Hypothesis holds:

Theorem 3.1. *The Small-Set Expansion Hypothesis implies the Unique Games Conjecture.*

This result is remarkable for a number of reasons. First, it presents the first non-trivial "reverse" reduction from a natural combinatorial optimization problem to Unique Games. Prior to this result, inapproximability results only demonstrated reductions of Unique Games to a certain combinatorial problem. This result also connects the UGC to the more well-studied problem of approximating graph expansion, and makes concrete the conspicuous presence of small set expansion in semi-definite programming (SDP) integrality gap instances for Unique Games and related problems (see [KV05] and [ABS10] for an in-depth discussion on this). Finally, as a consequence of Theorem 3.1, an algorithmic refutation of the UGC now obtains a (polynomial time) algorithm for approximating edge expansion in a certain regime (see Theorem 5.1).

The proof of this theorem is predicated on a reduction from SMALL-SET EXPANSION to UNIQUE GAMES (the composition of Reduction 3.1 and Reduction 3.2). We prove this theorem at the end of this section.

The reduction decomposes into two parts: first, we reduce Small-Set Expansion to Partial Unique Games (i.e. partial assignments allowed). This is the content of Reduction 3.1 and Theorem 3.3. Second, we show how to reduce an arbitrary instance of Partial Unique Games to Unique Games; see Reduction 3.2 and Theorem 3.11. Theorem 3.1 then follows by instantiating Theorem 3.3 and Theorem 3.11 with an appropriate choice of parameters.

3.1. From Small-Set Expansion to Partial Unique Games.

Reduction 3.1 (SMALL-SET EXPANSION to PARTIAL UNIQUE GAMES)**.**

> **Input:** A regular graph G with vertex set V and parameters $\epsilon > 0$ and $R \in \mathbb{N}$ (satisfying $\epsilon R \in \mathbb{N}$).
>
> **Output:** A unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ with vertex set $V^{R'}$ and alphabet $[R']$ where $R' := (1+\epsilon)R$.

The unique game $\mathcal{U}$ corresponds to the following probabilistic verifier for an assignment $F : V^{R'} \to [R']$:

(1) Sample R random vertices $a_1, \ldots, a_R \sim G$. Let $A := (a_1, \ldots, a_R) \in V^R$,
(2) Sample two random neighbors $b_i, b'_i \sim G(a_i)$ for every $i \in [R]$,
(3) Sample $2\epsilon R$ random vertices $b_{R+1}, b'_{R+1}, \ldots, b_{R+\epsilon R}, b'_{R+\epsilon R} \sim V$. Let $B := (b_1, \ldots, b_{R'}), B' := (b'_1, \ldots, b'_{R'}) \in V^{R'}$,
(4) Sample two permutations $\pi, \pi' \in \mathcal{S}_{[R']}$.
(5) Verify that $\pi^{-1}(F(\pi.B)) = (\pi')^{-1}(F(\pi'.B'))$. (Here $\pi.B$ refers to the tuple obtained by permuting the coordinates of B according to π.)

Reduction 3.1 has the following approximation guarantees:

Theorem 3.2. *Given a regular graph G with n vertices and parameters $R \in \mathbb{N}$ and $\epsilon > 0$, Reduction 3.1 computes in time $n^{O(R)}$ a unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ such that the following assertions hold (for every $\epsilon' > \epsilon$):*

Completeness: *If the graph G contains a vertex set with volume $\delta := \frac{1}{2R}$ and expansion at most ϵ, then the unique game $\mathcal{U}$ has α-partial value $\mathrm{opt}_\alpha(\mathcal{U}) \geq 1 - 5\epsilon$, where $\alpha \geq \frac{1}{10}$.*

Soundness I: *If every vertex set of G with volume $\delta := \frac{1}{2R}$ has expansion at least $1 - \epsilon$, then for all $\alpha \geq \frac{1}{10}$, the unique game $\mathcal{U}$ has α-partial value $\mathrm{opt}_\alpha(\mathcal{U}) \leq O(\epsilon^{1/5})$.*

Soundness II: *If every vertex set of G with volume between $\Omega(\epsilon^2/R)$ and $O(1/\epsilon^2 R)$ has expansion at least ϵ' and half of the edges of every vertex of G are self-loops, then for all $\alpha \geq \frac{1}{10}$, the unique game $\mathcal{U}$ has α-partial value $\mathrm{opt}_\alpha(\mathcal{U}) \leq 1 - \frac{\epsilon'}{4}$.*

Remark 3.3. The original reduction from Small-Set Expansion to Partial Unique Games of [RS10] introduced random noise into the sampling procedure to make the reduction more efficient by sampling fewer vertices. We present a less efficient reduction to make the proof of the soundness of the reduction more succinct and less opaque.

3.1.1. *Completeness.* Let G be a regular graph with vertex set $V = [n]$. Note that the size of the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ produced by Reduction 3.1 has size $n^{O(R)}$. Thus the size of $\mathcal{U}$ is polynomial in the size of G for every $R = O(1)$.

The following lemma shows that Reduction 3.1 is complete in the sense that if G contains a set of volume δ with expansion close to 0, then the unique game has a partial assignment that satisfies almost all constraints with labeled vertices.

Lemma 3.4 (Completeness)**.** *For every set $S \subset V$ with $\mu(S) = \delta$ and $\Phi(S) = \eta$, there exists a partial assignment $F = F_S$ for the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ (as defined in Reduction 3.1) satisfying*

$$\mathcal{U}(F) \geq (1 - \epsilon - 4\eta)\alpha,$$

where $\alpha \geq R'\delta(1 - R'\delta)$ is the fraction of vertice of $\mathcal{U}$ labeled by F.

Proof. We may assume that $R' \leq 1/\delta$, for otherwise $(1 - \epsilon - 4\eta)\alpha$ is close to zero so that the lemma becomes trivial. Consider the following partial assignment $F : V^{R'} \to [R'] \cup \{\perp\}$ for $\mathcal{U}$ defined by

$$X := (x_1, \ldots, x_{R'}) \longmapsto \begin{cases} i & \text{if } \{i\} = \{i' \in [R'] \mid x_{i'} \in S\}, \\ \perp & \text{otherwise.} \end{cases}$$

In terms of 2-prover games, this (partial) strategy amounts to the provers answering with vertex i iff it is the only vertex of S in their tuple, and otherwise refusing to answer.

We compute the number of vertices labeled by F:

$$\alpha = \mathbb{P}_{X \sim V^{R'}}[F(X) \neq \perp] = \binom{R'}{1}\delta(1-\delta)^{R'-1} = R'\delta(1-\delta)^{R'-1} \geq R'\delta(1 - R'\delta)$$

where we use the fact that $\mu(S) = \delta$.

Next, we estimate the fraction of constraints satisfied by F. Sample A, B, B' as specified in Reduction 3.1. Observe that F behaves nicely under permutations in the sense that $F(\pi.X) = \pi(F(X))$. Using this fact and the fact that a constraint is satisfied only if both vertices (of the constraint) are labelled, we see that

$$\begin{aligned}\mathcal{U}(F) &= \mathbb{P}_{A,B,B'}[F(B) = F(B') \in [R']] \\ &\geq \mathbb{P}_{A,B,B'}[F(B) = F(B') \in [R]] \\ &= \mathbb{P}_{X \sim V^{R'}}[F(X) \in [R]] \cdot \mathbb{P}_{A,B,B'}[F(B) = F(B') \mid F(B) \in [R]].\end{aligned}$$

We observe that this inequality is basically tight because the event $F(X) \in [R'] \setminus [R]$ is very unlikely:

$$\begin{aligned}\mathbb{P}_{X \sim V^{R'}}[F(X) \in [R]] &= \mathbb{P}_{X \sim V^{R'}}[F(X) \in [R']]/(1+\epsilon) \\ &\geq (1-\epsilon)\mathbb{P}_{X \sim V^{R'}}[F(X) \in [R']] \\ &\geq (1-\epsilon)\mathbb{P}_{X \sim V^{R'}}[F(X) \neq \perp] \\ &= (1-\epsilon)\alpha\end{aligned}$$

where we use the inequality $\frac{1}{1+\epsilon} \geq 1-\epsilon$ (with equality if and only if $\epsilon = 1$). Next we relate the probability of the event $F(B) = F(B')$ conditioned on $F(B) \in [R]$ to the expansion of S. It turns out that this probability is more directly related to the expansion of S in G^2. Put $\eta' = \Phi_{G^2}(S) = \mathbb{P}_{a \sim V, b, b' \sim G(a)}[b' \notin S \mid b \in S]$. Then

$$\begin{aligned}\mathbb{P}_{A,B,B'}[F(B) = F(B') \mid F(B) \in [R]] &= \mathbb{P}_{A,B,B'}[F(B') = 1 \mid F(B) = 1] \\ &= (1-\eta')(1-\eta'\delta/(1-\delta))^{R-1} \\ &\geq 1 - \eta'\left(1 + \frac{1-1/R}{1-\delta}R\delta\right) \\ &\geq 1 - 2\eta'\end{aligned}$$

where we use the symmetry of F in the first line and in the second line that $\mathbb{P}_{a \sim V, b, b' \sim G(a}[b' \in S | b \notin S] = \eta'\delta/(1-\delta)$ (this is the expansion of $V - S$ in G^2). The first inequality is the normal estimate on the exponent and the last inequality uses that $R' \leq 1/\delta$ and thus $1 - 1/R \leq 1 - \delta$ and $R\delta \leq 1$.

Now we relate η' to η (the expansion of S in G^2 to that in G):

$$\begin{aligned}\eta'\delta &= \mathbb{P}_{a \sim V, b, b' \sim G(a)}[b \in S \wedge b' \notin S] \\ &\leq \mathbb{P}_{a \sim V, b, b' \sim G(a)}[(b \in S \wedge a \notin S) \vee (a \in S \wedge b' \notin S)] \\ &\leq \mathbb{P}_{a \sim V, b \sim G(a)}[b \in S \wedge a \notin S] + \mathbb{P}_{a \sim V, b' \sim G(a)}[a \in S \wedge b' \notin S] \\ &= 2\delta\eta.\end{aligned}$$

Combining the previous inequalities now obtains the lower bound on $\mathcal{U}(F)$:

$$\begin{aligned}\mathcal{U}(F) &= \mathbb{P}_{A,B,B'}[F(B) = F(B') \in [R']] \\ &\geq \mathbb{P}_{X \sim V^{R'}}[F(X) \in [R]]\mathbb{P}_{A,B,B'}[F(B') = 1 \mid F(B) = 1] \\ &\geq (1-\epsilon)\alpha \cdot (1-4\eta) \\ &\geq (1-\epsilon-4\eta)\alpha.\end{aligned}$$

□

3.1.2. *Soundness.* Let $F : V^{R'} \to [R'] \cup \{\perp\}$ be a partial assignment for the unique game $\mathcal{U} = \mathcal{U}_{\epsilon,R}(G)$ obtained by applying Reduction 3.1 to a regular graph G on $V = [n]$.

For a tuple $U \in V^{R'-1}$ and a vertex $x \in V$, we let $f(U, x)$ be the probability that F selects the coordinate of x after we place it at a random position of U and permute the tuple randomly, i.e.,

$$f(U,x) \stackrel{\text{def}}{=} \mathbb{P}_{i\in[R'],\pi\in\mathcal{S}_{R'}}[F(\pi.(U +_i x) = \pi(i)]$$

where $U +_i x$ denote the tuple obtained from U by inserting x at the i-th coordinate and moving the original coordinates $i,\dots,R'-1$ of U one to the the right. Notice that fixing $i = R'$ doesn't change the above probability since the permutation π is random.

For $U \in V^{R'-1}$, define the function $f_U : V \to [0,1]$ by

$$x \mapsto \mathbb{E}_{W\sim G^{\otimes(R-1)}(U),Z\sim V^{\epsilon R}} f(W,Z,x) \tag{3.5}$$

where W,Z is the concatenation of the tuples W and Z, $G^{\otimes(R-1)}$ is the $(R-1)$-fold tensor product (this is the Kronecker product from linear algebra of the adjacency matrix of G with itself), and $W \sim G^{\otimes(R-1)}(U)$ means that W is a random neighbor of the vertex U in $G^{\otimes(R-1)}$. For $U = (u_1,\dots,u_{R-1})$, the distribution $G^{\otimes(R-1)}(U)$ is the product of the distributions $G(u_1),\dots,G(U_{R-1})$.

Before continuing, we will need some properties of the functions $\{f_U\}$. The following lemma is proven in Appendix A.

Lemma 3.6. *Let $\alpha = \mathbb{P}_{X\sim V^{R'}}[F(X) \neq \perp]$ be the fraction of vertices of $\mathcal{U}$ labeled by the partial assignment F.*

1. *The typical L_1-norm of f_U equals $\mathbb{E}_{U\sim V^{R-1}}\|f_U\|_1 = \frac{\alpha}{R}$,*
2. *For every $U \in V^{R-1}$, the L_1 norm of f_U satisfies $\|f_U\|_1 \leq \frac{1}{\epsilon R}$,*
3. *The typical squared L_2-norm of Gf_U relates to the fraction of constraints satisfied by F in $\mathcal{U}$ as follows,*
$$\mathbb{E}_{U\sim V^{R-1}}\|Gf_U\|^2 \geq \frac{1}{R'}(\mathcal{U}(F) - \frac{1}{\epsilon R})$$
(Here, we identify G with its stochastic adjacency matrix.)

The following lemma is then a consequence of the previous Lemma 3.2 and a Markov-type inequality argument. In rough terms, the lemma shows that given a good partial assignment for $\mathcal{U}$, we can extract a function $f : V \to [0,1]$ with $\|f\|_1 \approx 1/R$ and, at the same time, the squared L_2-norm of Gf is comparable to the L^1 norm of f. In Lemma 3.8, we will see that given such a function we can find a non-expanding set for G^2 with volume roughly $1/R$.

Lemma 3.7. *Let F be a partial assignment for the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ (as specified in Reduction 3.1). Suppose $\alpha = \mathbb{P}_{X\sim V^{R'}}[F(X) \neq \perp]$ is the fraction of vertices of $\mathcal{U}$ labeled by F. Then, for every $\beta > 0$, there exists $U \in V^{R-1}$ such that the function $f_U : V \to [0,1]$ (as defined in equation (3.5)) satisfies*

$$\|Gf_U\|^2 \geq \left(\frac{\mathcal{U}(F)}{\alpha} - \beta - \frac{1}{\alpha\epsilon R}\right)\|f_U\|_1 \qquad \textit{and} \qquad \frac{\alpha\beta}{R'} \leq \|f_U\|_1 \leq \frac{1}{\epsilon R}$$

Proof. Since G is regular and $0 \leq f_U \leq 1$, we have $\|Gf_U\|^2 \leq \|f_U\|_1$ for every $U \in V^{R-1}$ so that we can lower bound the expected square L_2-norm of Gf_U conditioned on $\|f_U\|_1 \geq \alpha\beta/R'$:

$$\begin{aligned}
\mathbb{E}_{U\sim V^{R-1}}\|Gf_U\|^2\mathbb{1}_{\|f_U\|_1\geq\alpha\beta/R'} &= \mathbb{E}_{U\sim V^{R-1}}\|Gf_U\|^2 - \mathbb{E}_{U\sim V^{R-1}}\|Gf_U\|^2\mathbb{1}_{\|f_U\|_1\leq\alpha\beta/R'} \\
&\geq \mathbb{E}_{U\sim V^{R-1}}\|Gf_U\|^2 - \alpha\beta/R' \\
&\geq \frac{\alpha}{R'}(\mathcal{U}(F)/\alpha - \frac{1}{\alpha\epsilon R} - \beta)
\end{aligned}$$

where in the last line we use that $\mathbb{E}_{U\sim V^{R-1}}\|Gf_U\|^2 \geq (\mathcal{U}(F) - \frac{1}{\epsilon R})R'$ (Lemma 3.2, item 3). Since $\mathbb{E}_{U\sim V^{R-1}}\|f_U\|_1 = \alpha/R'$ (Lemma 3.2, item 1), there exists a tuple $U \in V^{R-1}$ such that

$$\|Gf_U\|^2 \mathbb{1}_{\|f_U\|_1 \geq \alpha\beta/R'} \geq \|f_U\|_1(\mathcal{U}(F)/\alpha - \frac{1}{\alpha\epsilon R} - \beta).$$

This function f_U satisfies

$$\|f_U\|_1 \geq \alpha\beta/R' \qquad \text{and} \qquad \|Gf_U\|^2 \geq \|f_U\|_1(\mathcal{U}(F)/\alpha - \frac{1}{\alpha\epsilon R} - \beta).$$

Since $\|f_U\|_1 \leq 1/\epsilon R$ by (Lemma 3.2, item 2), the tuple U is as desired. □

The following lemma is similar in spirit to Cheeger's inequality. Given a function f on a vertex set V taking values in $[0,1]$, we can find $S \subset V$ with volume roughly $\|f\|_1$ and expansion roughly $1 - \|Gf\|^2/\|f\|_1$ in the graph G^2.

The proof, however, is markedly simpler than that of Cheeger's inequality. It is enough to analyze the distribution over vertex sets S obtained by including x in S with probability $f(x)$ independently for every vertex $x \in V$. We will later apply this lemma to the function obtained in the previous lemma (Lemma 3.4).

Lemma 3.8. *Suppose $f : V \to \mathbb{R}$ satisfies $0 \leq f(x) \leq 1$ for every $x \in V$. Then, for every $\beta > 0$, there exists a set $S \subset V$ such that*

$$\beta\|f\|_1 \leq \mu(S) \leq \beta^{-1}\|f\|_1,$$
$$\Phi_{G^2}(S) \leq 1 - \frac{\|Gf\|^2}{\|f\|_1} + 2\beta + \beta/(n\|f\|_1).$$

Proof. Consider the following distribution over level sets $S \subset V$ of f: for every vertex $x \in V$, include $x \in S$ with probability $f(x)$ independently for every vertex.

Then the expected volume of S is $\mathbb{E}_S\mu(S) = \mathbb{E}_S\|\mathbb{1}_S\|_1 = \|f\|_1$ and $\mathbb{E}_S\mu(S)^2 \leq \|f\|_1^2 + 1/n$ using the identity $\mu(S) = \|\mathbb{1}_S\|_1$. On the other hand, we have the lower bound $\mathbb{E}_S\|G\mathbb{1}_S\|^2 \geq \|Gf\|^2$.

Next, we bound the expected squared L_2-norm of $G\mathbb{1}_S$ conditioned on the event $\beta \leq \mu(S)/\|f\|_1 \leq 1/\beta$ from below:

$$\begin{aligned}
&\mathbb{E}_S\|G\mathbb{1}_S\|^2\mathbb{1}_{\{\beta\leq\mu(S)/\|f\|_1\leq 1/\beta\}} \\
&\geq \mathbb{E}_S\|G\mathbb{1}_S\|^2 - \mathbb{E}_S\mu(S)\mathbb{1}_{\{\mu(S)>\|f\|_1/\beta\}} - \mathbb{E}_S\mu(S)\mathbb{1}_{\{\mu(S)<\beta\|f\|_1\}} \\
&\geq \|Gf\|^2 - \mathbb{E}_S\beta\mu(S)^2/\|f\|_1 - \beta\|f\|_1 \\
&\geq \|Gf\|^2 - 2\beta\|f\|_1 - \beta/(n\|f\|_1).
\end{aligned}$$

It follows that there exists a set $S^* \subset V$ with $\beta \leq \mu(S^*)/\|f\|_1 \leq 1/\beta$ such that

$$\frac{\|G\mathbb{1}_{S^*}\|^2}{\|\mathbb{1}_{S^*}\|_1} \geq \frac{\mathbb{E}_S\|G\mathbb{1}_S\|^2\mathbb{1}_{\{\beta\leq\mu(S)/\|f\|_1\leq 1/\beta\}}}{\mathbb{E}_S\|\mathbb{1}_{S^*}\|_1} \geq \frac{\|Gf\|^2 - 2\beta\|f\|_1 - \beta/(n\|f\|_1)}{\|f\|_1}.$$

The quantity $1 - \frac{\|G\mathbb{1}_{S^*}\|^2}{\|\mathbb{1}_{S^*}\|_1}$ is the expansion of S^* in G^2 so we are done. □

Combining the previous two lemmas, we can now show that a good partial assignment for the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ obtained by applying Reduction 3.1 to G implies that G^2 contains a set with low expansion and volume roughly $1/R$. We want to use this to find a set with low expansion in G. The following two lemmas describe how to reduce the problem of finding a non-expanding set in G to a non-expanding set in G^2.

We first consider the case where the fraction of constraints satisfied by the partial assignment is bounded away from 0. Here, we show how to construct a set of with volume $1/R$ and expansion bounded away from 1 in G.

Lemma 3.9 (Soundness close to 0). *Suppose there exists a partial assignment F for the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(G)$ (as defined in Reduction 3.1) with $\mathcal{U}(F) \geq \eta\alpha$, where α is the fraction of vertices of $\mathcal{U}$ labeled by F. Then, there exists a set $S^* \subset V$ with $\mu(S^*) = \alpha/R$ and $\Phi(S^*) \leq 1 - O(\epsilon\eta^4)$. Here, we make the (mild) assumptions that $1/\alpha\epsilon\eta \ll R \ll \alpha\eta n$.*

Proof. Let $\beta, \gamma > 0$. (We determine the best choice for these parameters later in the proof). By Lemma 3.4, there exists $U \in V^{R-1}$ such that the function f_U satisfies the following condition

$$\|Gf_U\|^2 \geq (\eta - \beta - \frac{1}{\alpha\epsilon R})\|f_U\|_1 \qquad \text{and} \qquad \frac{\alpha\beta}{R'} \leq \|f_U\|_1 \leq \frac{1}{\epsilon R}.$$

Using Lemma 3.5 we can round f_U to a vertex set $S \subset V$ with the following properties:

$$\frac{\alpha\beta^2}{R'} \leq \mu(S) \leq \frac{1}{\beta\epsilon R'} \qquad \text{and} \qquad \Phi_{G^2}(S) \leq 1 - \eta + \beta + \frac{1}{\alpha\epsilon R} + 2\beta + \frac{R'}{n\alpha}.$$

Choosing $\beta \ll \eta$ (say $\beta = \eta/10$) and using our assumptions on R, the expansion of S in G^2 is at most $1 - \eta/2$.

We now construct a set with expansion in G bounded away from 1 and roughly the same volume as S. Define

$$S' := \{x \in V \mid \mathbb{P}_{y \sim G(x)[y \in S] \geq \gamma}\}.$$

We see that $\mu(S') \leq \mu(S)/\gamma$, so the volume of S' is not much larger than that of S. On the other hand, we can relate the fraction of edges between S and S' to the expansion of S in G^2 as follows:

$$\begin{aligned}
(1 - \Phi_{G^2}(S))\mu(S) &= \mathbb{P}_{x \sim V, y, y' \sim G(x)}[y \in S, y' \in S] \\
&= \mu(S')\mathbb{P}_{x \sim V, y, y' \sim G(x)}[y \in S, y' \in S \mid x \in S'] \\
&\quad + \mu(V - S)\mathbb{P}_{x \sim V, y, y' \sim G(x)}[y \in S, y' \in S \mid x \notin S'] \\
&\leq \mu(S')\mathbb{P}_{x \sim V, y \sim G(x)}[y \in S \mid x \in S'] \\
&\quad + \gamma\mu(V - S)\mathbb{P}_{x \sim V, y \sim G(x)}[y \in S \mid x \notin S'] \\
&= (1 - \gamma)\mu(S')\mathbb{P}_{x \sim V, y \sim G(x)}[y \in S \mid x \in S'] + \gamma\mu(S).
\end{aligned}$$

It follows that $E(S, S') \geq (1 - \Phi_{G^2}(S) - \gamma)\mu(S)$ so for $S'' := S \cup S'$ we have

$$1 - \phi_G(S'') \geq \frac{E(S, S')}{\mu(S \cup S')} \geq (\frac{\eta}{2} - \gamma)\frac{\mu(S)}{\mu(S \cup S')} \geq (\frac{\eta}{2} - \gamma)\frac{\gamma}{2}.$$

Setting $\gamma := \eta/4$, we obtain $\Phi_G(S'') \leq 1 - \eta^2/32$. On the other hand, we observe that the volume of S'' satisfies

$$\Omega\left(\frac{\alpha\eta^2}{R}\right) \leq \mu(S'') \leq O\left(\frac{1}{\eta^2\epsilon R}\right).$$

To obtain a set S^* with the desired volume α/R, we can either pad S'' with the desired number of vertices (in the case $\mu(S^*) < \alpha/R$) or simply take a random subset of S'' with the desired cardinality (in the case $\mu(S^*) \geq \alpha/R$). Performing either operation obtains a set S^* with $\mu(S^*) = \alpha/R$ and $\Phi_G(S^*) \le 1 - O(\epsilon\eta^4)$ as desired. □

The goal in the following lemma will be to find vertex sets in G with certain volume and expansion close to 0. It turns out that in this case it is convenient to apply Reduction 3.1 to the graph $\frac{1}{2}I + \frac{1}{2}G$, i.e. the graph obtained from G by adding a self-loop at every vertex with weight $\frac{1}{2}$). The following lemma shows that if the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(\frac{1}{2}I + \frac{1}{2}G)$ (obtained by applying Reduction 3.1 to $\frac{1}{2}I + \frac{1}{2}G$) has a partial assignment that satisfies almost as many constraints as possible, then we can find a vertex set in G with volume roughly $1/R$ and expansion close to 0. The proof is similar to the previous lemma.

Lemma 3.10 (Soundness close to 1)**.** *Suppose there exists a partial assignment x for the unique game $\mathcal{U} = \mathcal{U}_{R,\epsilon}(\frac{1}{2}I + \frac{1}{2}G)$ (as defined in Reduction 3.1) with $\mathcal{U}(x) \ge (1-\eta)\alpha$, where α is the fraction of vertices of $\mathcal{U}$ labeled by x. Then there exists a set $S \subset V$ with $\Omega(\alpha\eta^2/R) \le \mu(S) \le O(1/\epsilon\eta R)$ and $\Phi(S) \le 4\eta$. Here, we make the (mild) assumptions that $1/\alpha\epsilon\eta \ll R \ll \alpha\eta n$.*

Proof. Let $\beta > 0$. (We again choose the best value for β later in the proof). Let $G_{\circlearrowleft} := \frac{1}{2}I + \frac{1}{2}G$. Combining Lemma 3.4 and Lemma 3.5, we obtain a set S with the following properties:

$$\frac{\alpha\beta^2}{R'} \le \mu(S) \le \frac{1}{\beta\epsilon R'}$$

$$\Phi_{G^2_{\circlearrowleft}}(S) \le \eta + \beta + \frac{1}{\alpha\epsilon R} + 2\beta + \frac{R'}{n\alpha}.$$

As before, choosing $\beta \ll \eta$ and using the assumptions on R obtains $\Phi_{G^2_{\circlearrowleft}}(S) \le 2\eta$. We can directly compare the quantities $\Phi_{G^2_{\circlearrowleft}}(S)$ and $\Phi_G(S)$ as follows:

$$\langle \mathbb{1}_S, G^2_{\circlearrowleft}\mathbb{1}_S\rangle = \frac{1}{4}\langle \mathbb{1}_S, I\mathbb{1}_S\rangle + \frac{1}{2}\langle \mathbb{1}_S, G\mathbb{1}_S\rangle + \frac{1}{4}\langle \mathbb{1}_S, G^2\mathbb{1}_S\rangle \le \frac{1}{2}\|\mathbb{1}_S\|^2 + \frac{1}{2}\langle \mathbb{1}_S, G\mathbb{1}_S\rangle.$$

It follows that $\Phi_{G^2_{\circlearrowleft}}(S) \ge \Phi_G(S)/2$ so that $\Phi_G(S) \le 4\eta$ as desired. □

3.2. Partial Unique Games to Unique Games.

Reduction 3.2 (PARTIAL UNIQUE GAMES to UNIQUE GAMES)**.**

- **Input:** A unique game $\mathcal{U}$ with vertex set V and alphabet Σ, and a parameter $c \in \mathbb{N}$.
- **Output:** A unique game $\mathcal{U}' = \Psi_c(\mathcal{U})$ with vertex set $V' = V^c$ and alphabet $\Sigma' = [c] \times \Sigma$.

The unique game $\mathcal{U}'$ corresponds to the following probabilistic verifier for an assignment $F : V' \to \Sigma'$:

1. Sample c random vertices $u_1, \dots, u_c \sim V$,
2. Sample two random constraints $(u_r, v_r, \pi_r), (u_r, v'_r, \pi'_r) \sim \mathcal{U} \mid u_r$ for every $r \in [c]$ (the notation $\mathcal{U} \mid u_r$ denotes the uniform distribution over constraints of $\mathcal{U}$ that contain vertex u_r). Let $(r, j) := F(v_1, \dots, v_c)$ and $(r', j') := F(v'_1, \dots, v'_c)$,
3. Verify that $r = r'$ and that $j = \pi_r(i)$ and $j' = \pi_{r'}(i)$ for some label $i \in \Sigma$. (Note that at most one label satisfies this condition.)

Reduction 3.2 has the following approximation guarantees:

Theorem 3.11. *Given a parameter $c \in \mathbb{N}$ and a unique game $\mathcal{U}$ with n vertices, Reduction 3.2 computes in time poly(n^c) a unique game $\mathcal{U}' = \Psi_c(\mathcal{U})$ such that the following assertions hold (for all $\alpha, \eta, \eta', \zeta > 0$):*

Completeness: *If $\mathrm{opt}_\alpha(\mathcal{U}) \geq 1 - \eta$, then opt($\mathcal{U}'$) $\geq 1 - 4\eta - 2e^{-\alpha c}$.*

Soundness I: *If $opt_{1/2c}(\mathcal{U}) < \zeta$, then opt($\mathcal{U}'$) $< 8\zeta$.*

Soundness II: *If $opt_{1/2c}(\mathcal{U}) < 1 - \eta'$ and half of the constraints of every vertex are trivial identity constraints, then opt($\mathcal{U}'$) $< 1 - \eta'/32$.*

3.2.1. *Completeness.* Let $\mathcal{U}$ be a unique game with vertex set $V = [n]$ and alphabet $\Sigma = [k]$. Recall that $\mathrm{opt}_\alpha(\mathcal{U})$ is the optimal value of an α-partial assignment for $\mathcal{U}$.

The following lemma shows that Reduction 3.2 is complete, i.e., given a unique game $\mathcal{U}$ with $\mathrm{opt}_\alpha(\mathcal{U}) \geq 1-\eta$ for some constant α, for $c = O(\log(1/\eta))$, the unique game $\mathcal{U}' = \Psi_c(\mathcal{U})$ obtained by Reduction 3.2 has value $\mathrm{opt}(\mathcal{U}') \geq 1 - O(\eta)$.

Lemma 3.12. *If $\mathrm{opt}_\alpha(\mathcal{U}) \geq 1 - \eta$, then opt($\mathcal{U}'$) $\geq 1 - 4\eta - 2e^{-\alpha c}$.*

Proof. Let $f : V \to \Sigma \cup \{\perp\}$ be an optimal α-partial assignment for $\mathcal{U}$. We may assume without loss of generality that $\mathbb{P}_{u\sim V}[f(u) \neq \perp] = \alpha$ and thus $\mathcal{U}(f) \geq (1-\eta)\alpha$. To lower bound the value of the unique game $\mathcal{U}'$, we consider the following assignment $F : V^c \to [c] \times \Sigma$ given by

$$(u_1, \dots, u_c) \longmapsto \begin{cases} (r, i) & \text{if } (f(u_r) = i \in \Sigma) \wedge (\forall j < r (f(u_j) = \perp)) \\ (1, 1) & \text{if } f(u_1) = \dots = f(u_c) = \perp . \end{cases}$$

In words, the prover returns the first answer in the list $f(u_1), \dots, f(u_c)$ (ignoring $\perp$). If the partial strategy (assignment) refuses to answer on all inputs, then the prover returns an arbitrary answer. Note that returning an arbitrary answer can only increase the number of satisfied constraints.

We claim that this assignment F satisfies at least $1-4\eta-2e^{-\alpha c}$ of the constraints of $\mathcal{U}'$, which proves the lemma. To establish this claim, it suffices to show that

$$\mathbb{P}_{(u_1,v_1,\pi_1)\sim\mathcal{U},\dots,(u_c,v_c,\pi_c)\sim\mathcal{U}}[\exists r \in [c], i \in \Sigma : F(u_1, \dots, u_c) = (r, i) \wedge f(v_1, \dots, v_c) = (r, \pi_r(i))] \geq 1 - 2\eta - e^{-\alpha c}$$

Observe that this probability simplifies to the product of two probabilities p_1, p_2 (which will be easier to lower bound):

$$\begin{aligned} &\mathbb{P}_{(u_1,v_1,\pi_1)\sim\mathcal{U},\dots,(u_c,v_c,\pi_c)\sim\mathcal{U}}[\exists r \in [c], i \in \Sigma : F(u_1, \dots, u_c) = (r, i) \wedge f(v_1, \dots, v_c) = (r, \pi_r(i))] \\ &\quad = \mathbb{P}_{(u_1,v_1,\pi_1)\sim\mathcal{U},\dots,(u_c,v_c,\pi_c)\sim\mathcal{U}}[\exists r \in [c] : f(u_r) \neq \perp \vee f(v_r) \neq \perp] \\ &\quad \cdot \mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[f(v) = \pi(f(u)) \mid f(u) \neq \perp \vee f(v) \neq \perp] \\ &\quad := p_1 \cdot p_2. \end{aligned}$$

Since f is α-partial, the event $f(u_1) = \dots = f(u_c) = \perp$ has probability $(1-\alpha)^c \leq e^{-\alpha c}$. Thus with probability at least $1 - e^{-\alpha c}$, one of the vertices $u_1, \dots, u_c$ is labeled, so $p_1 \geq 1 - e^{-\alpha c}$.

We lower bound p_2 as follows:

$$\begin{aligned} p_2 &= \mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[f(v) = \pi(f(u)) \mid f(u) \neq \perp \vee f(v) \neq \perp] \\ &= \frac{\mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[f(v) = \pi(f(u))]}{\mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[f(u) \neq \perp \vee f(v) \neq \perp]} \\ &= \frac{\mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[f(v) = \pi(f(u))]}{2\mathbb{P}_{u\sim V}[f(u) \neq \perp] - \mathbb{P}_{(u,v,\pi)\sim\mathcal{U}}[f(u) \in \Sigma, f(v) \in \Sigma]} \\ &\geq \frac{(1-\eta)\alpha}{2\alpha - (1-\eta)\alpha} \\ &\geq 1 - 2\eta. \end{aligned}$$

So we can conclude that $p_1 \cdot p_2 \geq 1 - 2\eta - e^{-\alpha c}$ as claimed. □

3.2.2. *Soundness.* In the following, we shows that Reduction 3.2 is sound, i.e., that given a good assignment for the unique game $\mathcal{U}' = \Psi_c(\mathcal{U})$ obtained via applying Reduction 3.2 to $\mathcal{U}$, one can construct a good partial assignment for the original unique game $\mathcal{U}$.

It turns out that it is easier to bound the values of such a partial assignment in the unique game $\mathcal{U}^2$: the square of a unique game $\mathcal{U}$ on vertex set V with alphabet Σ is a unique game on the same vertex set and alphabet corresponding to the following probabilistic verifier:

(1) Sample a random vertex $u \sim V$ according to the marginal distribution of the first vertex in a random constraint of $\mathcal{U}$,
(2) Sample two random constraints incident to u: $(u, v, \pi), (u, v', \pi') \sim \mathcal{U} \mid u$. (Here $\mathcal{U} \mid u$ denotes the constraint distribution of $\mathcal{U}$ conditioned on the first vertex being u.)
(3) Verify that $\pi^{-1}(x_v) = (\pi')^{-1}(x_{v'}$.

Let $\mathcal{U}_{\circlearrowleft}$ be the unique game obtained from sampling with probability $1/2$ a random constraint from $\mathcal{U}$ and sampling with the remaining constraint a trivial constraint (u, u, id). Observe that $\mathcal{U}_{\circlearrowleft}(x) = \frac{1}{2}(1 + \mathcal{U}(x))$ so we have $\mathcal{U}^2_{\circlearrowleft}(x) = 1/4 + 1/2\mathcal{U}(x) + 1/4\mathcal{U}^2(x)$ for all $x \in (\Sigma \cup \{\perp\})^V$.

The following properties are immediate from unwrapping definitions, and will be used in Lemma 3.15 and Lemma 3.16:

Proposition 3.13.

(1) If $\mathcal{U}(x) \geq 1 - \eta$, then $\mathcal{U}^2(x) \geq 1 - 2\eta$.
(2) If $\mathcal{U}^2(x) \geq \zeta$, then one can construct an assignment y such that $\mathcal{U}(y) \geq \zeta/4$.
(3) If $\mathcal{U}(x) \geq 1 - \eta$, then $\mathcal{U}^2_{\circlearrowleft}(x) \geq 1 - \eta$.
(4) If $\mathcal{U}^2_{\circlearrowleft}(x) \geq 1 - \eta'$, then $\mathcal{U}(x) \geq 1 - 2\eta'$.

Let $F : V^c \to [c] \times \Sigma$ be an assignment for the unique game $\mathcal{U}' = \Psi_c(\mathcal{U})$ (recall that $\mathcal{U}$ is on vertex set $V = [n]$ and alphabet $\Sigma = [k]$). Based on F, we construct a collection $\{f_{U,r}\}$ of partial assignments for $\mathcal{U}$. For $U \in V^{c-1}$ and $r \in [c]$, define $f_{U,r} : V \to \Sigma \cup \{\perp\}$ by

$$x \mapsto \begin{cases} i & \text{if } F(U +_i x) = (r, i), \\ \perp & \text{otherwise.} \end{cases}$$

We first show that one of the partial assignments $f_{U,r}$ has good value in $\mathcal{U}^2$.

Lemma 3.14. *Let $\beta > 0$. Then, there exists a partial assignment $f_{U,r}$ that labels at least β/c of the vertices and has value at least*

$$\mathcal{U}^2(f_{U,r}) \geq \frac{\mathcal{U}'(F) - \beta}{1-\beta} \mathbb{P}_{x\sim V}[f_{U,r}(x) \neq \perp].$$

In particular, $\mathrm{opt}_{\beta/c}(\mathcal{U}^2) \geq (\mathcal{U}'(F) - \beta)/(1-\beta)$.

Proof. We first compute the expectation of $f_{U,r}$ for $\mathcal{U}^2$:

$$\begin{aligned}
&\mathbb{E}_{U\sim V^{c-1}, r\in[c]} \mathcal{U}^2(f_{U,r}) \\
&= \mathbb{P}_{U\sim V^{c-1}, r\in[c], u\sim V, (u,v,\pi),(u,v',\pi')\sim \mathcal{U}|u}[\pi^{-1}(f_{U,r}(v)) = (\pi')^{-1}(f_{U,r}(v'))] \\
&= \frac{\mathcal{U}'(F)}{c}.
\end{aligned}$$

For fixed $U \in V^{c-1}$ and $r \in [c]$, define

$$\alpha_{U,r} \stackrel{\text{def}}{=} \mathbb{P}_{x\sim V}[f_{U,r}(x) \neq \perp].$$

The expected value of $\alpha_{U,r}$ is $1/c$ since

$$\mathbb{E}_{U\sim V^{c-1}, r\in[c]} \alpha_{U,r} = \mathbb{P}_{U\sim V^{c-1}, r\in[c], x\sim V}[f_{U,r}(x) \neq \perp] = \frac{1}{c}.$$

So $\mathcal{U}^2(f_{U,r}) \leq \alpha_{U,r}$, which we can can use to bound the contribution of assignments $f_{U,r}$ with $\alpha_{U,r} \leq \beta/c$ to the expected value of $\mathcal{U}(f_{U,r})$:

$$\mathbb{E}_{U\in V^{c-1}, r\in[c]} \mathcal{U}^2(f_{U,r}) \cdot \mathbb{1}_{\alpha_{U,r} \leq \beta/c} \leq \beta/c.$$

Therefore we can lower bound the β/c-partial value of the unique game $\mathcal{U}^2$ via:

$$\begin{aligned}
\mathrm{opt}_{\beta/c}(\mathcal{U}) &\geq \frac{\mathbb{E}_{U\sim V^{c-1}, r\in[c]} \mathcal{U}^2(f_{U,r}) \cdot \mathbb{1}_{\alpha_{U,r} \geq \beta/c}}{\mathbb{E}_{U\sim V^{c-1}, r\in[c]} \alpha_{U,r} \cdot \mathbb{1}_{\alpha_{U,r} \geq \beta/c}} \\
&\geq \frac{\mathcal{U}'(F)/c - \beta/c}{1/c - \beta/c} \\
&= \frac{\mathcal{U}'(F) - \beta}{1-\beta}.
\end{aligned}$$

□

Using Lemma 3.14 and basic relations between the optimal value of a unique game and its square, we can show the following lemmas:

Lemma 3.15 (Soundness close to 0)**.** *Let $\mathcal{U}' = \Psi_c(\mathcal{U})$ be the unique game obtained by applying Reduction 3.2 with parameter $c \in \mathbb{N}$ to the unique game $\mathcal{U}$. If* $\mathrm{opt}(\mathcal{U}') \geq \zeta$*, then* $\mathrm{opt}_{1/2c}(\mathcal{U}) \geq \zeta/8$.

Let $\mathcal{U}_{\circlearrowleft}$ be the unique game obtained by sampling with probability $1/2$ a random constraint from $\mathcal{U}$ and sampling with the remaining probability a trivial constraint (u, u, id). Note that $\mathcal{U}(f) = 1/2(\alpha + \mathcal{U}(f))$ for every assignment $f : V \to \Sigma \cup \{\perp\}$that labels an α fraction of the vertices.

Lemma 3.16 (Soundness close to 1)**.** *Let $\mathcal{U}' = \Psi_c(\mathcal{U}_{\circlearrowleft})$ be the unique game obtained by applying Reduction 3.2 with parameter $c \in \mathbb{N}$ to the unique game $\mathcal{U}_{\circlearrowleft}$. If* $\mathrm{opt}(\mathcal{U}') \geq 1 - \eta$*, then* $\mathrm{opt}_{1/2c}(\mathcal{U}) \geq \eta/16$.

4. The Unique Games Conjecture on Small-Set Expanders implies the Small-Set Expansion Hypothesis

In this section, we introduce a variant of the Unique Games Conjecture wherein one additionally assumes mild expansion of small sets on the constraint graph of the instance. We then describe the reduction from the such instances of Unique games, which we call Unique Games on Small-Set Expanders, to the Small-Set Expansion problem. This establishes an equivalence between the Unique Games Conjecture on Small-Set Expanders and the Small-Set Expansion Hypothesis. We omit all proofs and refer the interested reader to [RST10a].

Conjecture 4.1 (Unique Games Conjecture on Small-Set Expanders)**.** *For every $\eta > 0$, there exists $\delta = \delta(\eta) > 0$ such that for every $\zeta > 0$, the following promise problem is **NP**-hard for some $R = R(\eta, \zeta)$: given a unique game $\mathcal{U}$ with alphabet $[R]$ and constraint graph G, distinguish whether*

(1) $\mathrm{opt}(\mathcal{U}) \geq 1 - \eta$*, or*
(2) $\mathrm{opt}(\mathcal{U}) \leq \zeta$ *and* $\Phi_G(S) \geq 1 - \eta$ *for all sets* S *with* $\mu(S) = \delta$*.*

Reduction 4.1 (From Unique Games to Small-Set Expansion)**.**

Input: A unique game $\mathcal{U}$ with vertex set V and alphabet $[R]$, and parameters $q \in \mathbb{N}$ and $\epsilon > 0$.

Output: A graph $H = H_{q,\epsilon}(\mathcal{U})$ on the vertex set $V \times [q]^R$ with edge distribution given by the following sampling protocol:

(1) Sample a random vertex $u \sim V$,
(2) Sample a random point $x \in [q]^R$,
(3) Sample two random constraints $(u, v, \pi), (u, v', \pi) \sim \mathcal{U} \mid u$,
(4) Sample two random points $y, y' \sim T_{1-\epsilon}(x)$ (Here, $T_{1-\epsilon}(x)$ is the usual noise graph on $[q]^R$ with noise parameter $1 - \epsilon$; see [RST10a] for a detailed definition),
(5) Output an edge between (v, y) and (v', y').

The Reduction 4.1 plus the appropriate analysis yields the following theorem:

Theorem 4.2. *Conjecture 4.1 (Unique Games Conjecture on Small-Set Expanders) implies the Small-Set Expansion Hypothesis.*

5. If the Unique Games Conjecture is false

Raghavendra and Steurer observe in [RS10] that an immediate corollary of Theorem 3.1 is that the following hypothesis would emerge out of an algorithmic refutation of the Unique Games Conjecture:

Conjecture 5.1 (Unique Games is Easy)**.** *There exists a constant $\epsilon > 0$ and a function $f : \mathbb{N} \to \mathbb{N}$ such that given a UNIQUE GAMES instance $\mathcal{U}$ with n vertices and k labels, it is possible to distinguish between the cases* $\mathrm{opt}(\mathcal{U}) \geq 1 - \epsilon$ *and* $\mathrm{opt}(\mathcal{U}) \leq \epsilon$ *in time $n^{f(k)}$.*

Theorem 5.2. *Suppose Conjecture 5.1 is true for some constant ϵ_0 and function f. Then there exists a function $g : [0, 1] \to \mathbb{N}$ such that given any graph G on n vertices and $\delta \in [0, 1]$, it is possible to distinguish whether $\Phi_G(\delta) \leq \epsilon_1$ or $\Phi_G(\delta) > 1 - \epsilon_1$ for some absolute constant ϵ_1 in time $n^{g(\delta)}$.*

Using the parallel repetition theorem of Rao [Rao08], we can actually say something stronger. The following corollary is an immediate consequence of our reduction from Small-Set Expansion to Unique Games and the aforementioned theorem:

Corollary 5.3. *Suppose Conjecture 5.1 is true for some constant* ϵ_0 *and function* f*. Then, given a graph* G *with vertex set* $V = [n]$ *and parameters* ϵ, δ *such that* $\epsilon < \epsilon_1$ *for some absolute constant* ϵ_1*, we can distinguish between the following cases in time* $n^{g(\epsilon,\delta)}$ *for some* $g : [0,1]^2 \to \mathbb{N}$

(1) *There exists* $S \subset V$ *with* $\mu(S) = \delta$ *and* $\Phi_G(S) \leq \epsilon$*,*
(2) *Every set* $S \subset V$ *with* $\mu(S) \leq 1500\delta/\epsilon$ *satisfies* $\Phi_G(S) \geq 1500\sqrt{\epsilon}$*.*

6. Conclusion

The recent connections between the Small-Set Expansion Hypothesis and the Unique Games Conjecture put the veracity of the latter conjecture on potentially more stable ground: graph expansion is well-studied, so it is natural to ask about the veracity of the Small-Set Expansion Hypothesis. On one hand, the recent resolution of the $2-2$ Games Conjecture [KMS18] (a lesser variant where there are two possible values each constraint is allowed to satisfy; the unique in Unique Games is for similar reasons) gives credence to the truth of the Unique Games Conjecture and thus by extension to the Small-Set Expansion Hypothesis. On the other hand, the Small-Set Expansion Hypothesis implies the following (improbable) combinatorial conjecture:

Conjecture 6.1. *For all* $\epsilon > 0$*, there exists* $\delta > 0$ *and a family of* n*-vertex graphs* $\{G_n\}_{n\to\infty}$ *such that (1) every subset of size between* $\frac{\delta}{2}n$ *and* δn *in* G_n *has expansion at least (say)* $1/2$ *and (2) the number of eigenvalues of the normalized adjacency matrix of* G_n *that exceed* $1-\epsilon$ *is at least* $n\delta$*.*

We refer the interested reader to [Kho14].

The Unique Games Conjecture has led to the creation of many interesting algorithms. As it relates to this paper, there are two algorithms ([Aretal08], [MM10]) which efficiently solve Unique Games instances where the underlying constraint graph has good expansion (these results both clearly generalize to the case of small-set expansion). In particular, since a random graph has good expansion by a straightforward combinatorial argument, these results show that UGC is easy on random instances. If one believes that Unique Games is indeed hard in the worst case (as the conjecture contends), then these results perhaps indicate that (small-set) expansion in the constraint graph is the wrong place to look for hard instances for Unique Games. On the other hand, if one believes that the Unique Games Conjecture is false, then these results indicate that one need only find a way to unite the expanding and non-expanding cases to disprove the conjecture. Proving or disproving the $2-2$ conjecture for graphs with small-set expansion properties would help to resolve this confusion, as a proof in either direction would help determine whether expansion is intrinsically tied-up with the veracity of the Unique Games Conjecture.

7. Acknowledgements

I would like to thank my mentor Antares Chen for bearing with me as I slowly read every paper and watched every lecture necessary to write this exposition. In addition, I want to thank Professor Peter May for organizing the excellent REU here at Chicago. Most importantly, I want to thank my friends and family for supporting me in this life. I love you Charlie.

Finally, I would like to remind Zach Cogan that he owes me one (1) dollar.

References

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. In Proceedings of the thirty-fourth annual ACM symposium on Theory of computing (STOC '02), 2002.

[Kho14] S. Khot. Hardness of approximation. In Proceedings of the International Congress of Mathematicians, 2014.

[KKMO07] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? SIAM J. Comput., 37(1):319–357, 2007.

[KR08] S. Khot, O. Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. Journal of Computer and System Sciences Volume 74, Issue 3, 2008.

[KV05] S. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ^1. In FOCS, pages 53–62, 2005.

[KMS18] S. Khot, D Minzer, and M. Safra. Pseudorandom sets in Grassmann graph have near-perfect expansion. In 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS, 2018.

[Rag08] P. Raghavendra. Optimal algorithms and inapproximability results for every csp? In Proc. 40th STOC, pages 245–254. ACM, 2008.

[Ste10a] D. Steurer. On the complexity of Unique Games and graph expansion. dissertation, 2011.

[Ste10b] D. Steurer. Fast SDP algorithms for constraint satisfaction problems. In SODA, 2010.

[RS10] P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture, In STOC, 2010.

[RST10a] P. Raghavendra, D. Steurer, and M. Tulsiani. Reductions between graph expansion problems. manuscript, 2010.

[RST10b] P. Raghavendra, D. Steurer, and P. Tetali. Approximations for the isoperimetric and spectral profile of graphs and related parameters. In Proc. STOC. ACM, 2010.

[ABS10] S. Arora, B. Barak, and D. Steurer Subexponential Algorithms for Unique Games and Related problems In FOCS, 2010.

[Rao08] A. Rao. Parallel repetition in projection games and a concentration bound. In STOC, pages 1–10. ACM, 2008.

[MM10] K. Makarychev, Y. Makarychev. How to Play Unique Games on Expanders. In WAOA, 2010.

[Aretal08] S. Arora, D. Steurer, S. Khot, M. Tulsiani, A. Kolla, and N. Vishnoi. Unique games on expanding constraint graphs are easy In STOC, 2008.

Appendix A. Facts About Functions Over Graphs

We now prove Lemma 3.5

Proof of Lemma 3.5. We proceed by proving each item.

(1) We evaluate $\mathbb{E}_{U\sim V^{R-1}}\|f_U\|_1$ as follows:

$$\begin{aligned}\mathbb{E}_{U\sim V^{R-1}}\|f_U\|_1 &= \mathbb{P}_{\substack{U\sim V^R, W\sim G^{\otimes(R-1)}(U),\\ Z\sim V^{\epsilon R}, x\sim V, \pi\in\mathcal{S}_{R'}, i\in[R']}}\left[F\left(\pi.\left((W,Z)+_i x\right)\right)=\pi(i)\right]\\ &= \mathbb{P}_{x\sim V^{R'}, \pi\in\mathcal{S})R', i\in[R']}[F(\pi.X)=\pi(i)]\\ &= \frac{1}{R'}\mathbb{P}_{X\in V^{R'}}[F(X)\neq\perp] = \frac{\alpha}{R'}.\end{aligned}$$

Here the second line uses that the joint distribution $(W,Z)+_i x$ and i is the same as the joint distribution of X and i and the third line uses that the distribution of $\pi(i)$ is uniformly random in $[R']$ even when condition on X and π (and thus on $F(\pi.X)$).

(2) For fixed $U\in V^{R-1}$, the L^1 norm of f_U evaluates to

$$\begin{aligned}\|f_U\|_1 &= \mathbb{P}_{\substack{W\sim G^{\otimes(R-1)}(U)\\ Z\sim V^{\epsilon R}, x\sim V, \pi\in\mathcal{S}_{R'},\\ i\in\{R,R+1,\dots,R'\}}}\left[F\left(\pi.\left((W,Z)+_i x\right)\right)=\pi(i)\right]\\ &= \mathbb{P}_{\substack{W\sim G^{\otimes(R-1)}(U)\\ Z'\sim V^{\epsilon R+1}, x\sim V, \pi\in\mathcal{S}_{R'},\\ i\in\{R,R+1,\dots,R'\}}}\left[F\left(\pi.(W,Z')\right)=\pi(i)\right]\\ &= \frac{1}{\epsilon R+1}\mathbb{P}_{\substack{W\sim G^{\otimes(R-1)}(U)\\ Z'\sim V^{\epsilon R+1}, \pi\in\mathcal{S}_{R'}}}\left[\pi^{-1}\left(F\left(\pi.(W,Z')\right)\right)\in\{R,R+1,\dots,R'\}\right]\\ &\leq \frac{1}{\epsilon R+1}\end{aligned}$$

where the second set uses that $(i,(U,Z)+_i x)$ has the same distribution as $(i,(U,Z'))$. As opposed to the proof of (1), we insert x in a random coordinate among $\{R,\dots,R'\}$ (as opposed to totally random in $[R']$). The experiment as a whole doesn't change since π is a random permutation, as remarked previously.

(3) Sample tuples A,B,B' as specified in Reduction 3.1). Note that B and B' are distributed independent and identically conditioned on A. We denote this conditional distribution $B\mid A$. We can compute $\mathcal{U}(F)$ as

$$\begin{aligned}\mathcal{U}(F) &= \sum_{r=1}^{R'}\mathbb{P}_{A,B,B',\pi,\pi'\in\mathcal{S}_{R'}}\left[F(\pi.B)=\pi(r)\wedge F(\pi'.B')=\pi'(r)\right]\\ &= \sum_{r=1}^{R'}\mathbb{E}_{A\sim V^R}\left(\mathbb{P}_{\substack{B\mid A\\ \pi\in\mathcal{S}_{R'}}}[F(\pi.B)=\pi(r)]\right)^2.\end{aligned}$$

Fix $A\in V^R$. Observe that for any $r,r'\in\{R+1,\dots,R'\}$, it holds that

$$\mathbb{P}_{\substack{B\mid A\\ \pi\in\mathcal{S}_{R'}}}[F(\pi.B)=\pi(r)] = \mathbb{P}_{\substack{B\mid A\\ \pi\in\mathcal{S}_{R'}}}[F(\pi.B)=\pi(r')]$$

since all coordinates b_r of B with $r\in\{R+1,\dots,R'\}$ are distributed identically (even conditioned on A). Thus, for every $r\in\{R+1,\dots,R'\}$,

$$\mathbb{P}_{\substack{B\mid A\\ \pi\in\mathcal{S}_{R'}}}[F(\pi.B)=\pi(r)]\leq\frac{1}{\epsilon R}.$$

Now consider $r\in[R]$. For $B=(b_1,\dots,b_{R'})\in V^{R'}$, let $B^{-r}\in V^{R'-1}$ denote the tuple obtained from B by removing the r-th coordinate b_r. Then

$$\mathbb{P}_{\substack{B\mid A\\ \pi\in\mathcal{S}_{R'}}}[F(\pi.B)=\pi(r)] = \mathbb{P}_{\substack{B\mid A\\ \pi\in\mathcal{S}_{R'}}}[F(\pi.(B^{-r}+_r b_r))=\pi(r)] = \mathbb{E}_{B\mid A}f(B^{-r},b_r).$$

Since b_r is distributed as a random neighbor of a_r (where $A=(a_1,\dots,a_{R'})$), it follows that

$$\mathbb{E}_{B\mid A}f(B^{-r},b_r) = \mathbb{E}_{b_r\sim G(a_r)}f_{A^{-r}}(b_r) = Gf_{A^{-r}}(a_r).$$

Recall that we identify the graph G with its stochastic adjacency matrix so that Gf_{A-r} denotes the function on V obtained by applying the linear operator G to the function f_{A-r}. Combining the two previous bounds, we get

$$\mathcal{U}(F) \leq \sum_{r=1}^{R} \mathbb{E}_{A\sim V^R} \left(Gf_{A-r}(a_r)\right)^2 + \epsilon R \left(\frac{1}{\epsilon R}\right)^2 = R\mathbb{E}_{U\sim V^{R-1}} \|Gf_U\}^2 + \frac{1}{\epsilon R}$$

which implies the desired bound. $\square$