

Graph Expansion and the Unique Games Conjecture

Prasad Raghavendra
MSR New England
Cambridge, MA

David Steurer^{*}
Princeton University
Princeton, NJ

ABSTRACT

The edge expansion of a subset of vertices $S \subseteq V$ in a graph G measures the fraction of edges that leave S . In a d -regular graph, the edge expansion/conductance $\Phi(S)$ of a subset $S \subseteq V$ is defined as $\Phi(S) = \frac{|E(S, V \setminus S)|}{d|S|}$. Approximating the conductance of small linear sized sets (size δn) is a natural optimization question that is a variant of the well-studied SPARSEST CUT problem. However, there are no known algorithms to even distinguish between almost complete edge expansion ($\Phi(S) = 1 - \varepsilon$), and close to 0 expansion.

In this work, we investigate the connection between Graph Expansion and the Unique Games Conjecture. Specifically, we show the following:

–We show that a simple decision version of the problem of approximating small set expansion reduces to UNIQUE GAMES. Thus if approximating edge expansion of small sets is hard, then UNIQUE GAMES is hard. Alternatively, a refutation of the UGC will yield better algorithms to approximate edge expansion in graphs.

This is the first non-trivial “reverse” reduction from a natural optimization problem to UNIQUE GAMES.

–Under a slightly stronger UGC that assumes mild expansion of small sets, we show that it is UG-hard to approximate small set expansion.

–On instances with sufficiently good expansion of small sets, we show that UNIQUE GAMES is easy by extending the techniques of [4].

Categories and Subject Descriptors

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

^{*}Supported by NSF grants 0830673, 0832797, 528414. Part of this work done while at Microsoft Research New England.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'10, June 5–8, 2010, Cambridge, Massachusetts, USA.
Copyright 2010 ACM 978-1-4503-0050-6/10/06 ...\$10.00.

General Terms

Algorithms, Theory.

Keywords

Graph Expansion, Unique Games Conjecture, Hardness of Approximation, Spectral profile.

1. INTRODUCTION

The phenomenon of vertex and edge expansion in graphs has been a subject of intense study with applications pervading almost all branches of theoretical computer science. From an algorithmic standpoint, approximating expansion or lack thereof (finding good cuts or separators) is a fundamental optimization problem with numerous applications. Yet, the computational complexity of detecting and approximating expansion in graphs is not very well understood.

Among the two notions of expansion, this work will concern mostly with edge expansion. For simplicity, let us first consider the case of a d -regular graph $G = (V, E)$. The edge expansion of a subset of vertices $S \subseteq V$ measures the fraction of edges that leave S . Formally, the edge expansion $\Phi(S)$ of a subset $S \subseteq V$ can be defined as,

$$\Phi_G(S) = \frac{|E(S, V \setminus S)|}{d|S|},$$

where $E(S, V \setminus S)$ denotes the set of edges with one endpoint in S and the other endpoint in $V \setminus S$. The conductance or the Cheeger’s constant associated with the graph G is the minimum of $\Phi(S)$ over all sets S with at most half the vertices, i.e.,

$$\Phi_G = \min_{|S| \leq n/2} \Phi_G(S).$$

The definitions of conductance of sets $\Phi_G(S)$ and the graph Φ_G can be extended naturally to non-regular graphs, and finally to arbitrary weighted graphs. For the sake of simplicity, we restrict our attention to regular graphs here, and defer the discussion of general weighted graphs to the full version. Henceforth, we will use the notation $\mu(S)$ to denote the normalized set size $\mu(S) = |S|/n$.

The problem of approximating Φ_G also referred to as the the uniform SPARSEST CUT (equivalent within a factor of 2), is among the fundamental problems in approximation algorithms. Efforts towards approximating Φ_G have led to a rich body of work with strong connections to spectral techniques and metric embeddings.

The first approximation for conductance was obtained by discrete analogues of the Cheeger inequality [9] shown by

Alon-Milman [2] and Alon [1]. Specifically, they show the following:

THEOREM 1.1 (CHEEGER’S INEQUALITY). *If λ_2 denotes the second largest eigenvalue of the suitably normalized adjacency matrix of a graph G then,*

$$\frac{1 - \lambda_2}{2} \leq \Phi_G \leq \sqrt{2(1 - \lambda_2)}.$$

Since the second eigenvalue λ_2 can be efficiently computed, Cheeger’s inequality yields an approximation algorithm for Φ_G , indeed one that is used heavily in practice for graph partitioning. However, the approximation for Φ_G obtained via Cheeger’s inequality is poor in terms of a approximation ratio, especially when the value of Φ_G is small (λ_2 is close to 1).

An $O(\log n)$ approximation algorithm for Φ_G was obtained by Leighton and Rao [17]. Later work by Linial et al. [18] and Aumann and Rabani [6] established a strong connection between the SPARSEST CUT problem and the theory of metric spaces, in turn spurring a large and rich body of literature. More recently, in a breakthrough result Arora et al. [5] obtained an $O(\sqrt{\log n})$ approximation for the problem using semidefinite programming techniques.

Small Set Expansion.

Note that the Φ_G is a fairly coarse measure of edge expansion, in that it is the worst case edge expansion over sets S of all sizes. In a typical graph (say a random d -regular graph), smaller sets of vertices expand to a larger extent than sets with half the vertices. For instance, all sets S of δn vertices in a random d -regular graph have $\Phi(S) \approx 1 - \frac{2}{d}$ with very high probability, while the conductance Φ_G of the entire graph is roughly $\frac{1}{2}$. Moreover, the stronger expansion exhibited by small sets has numerous applications in graph theory.

A natural finer measure of the edge expansion of a graph is its expansion profile. Specifically, for a regular graph G the expansion profile is given by the curve ¹

$$\Phi_G(\delta) = \min_{\mu(S)=\delta} \Phi(S) \quad \forall \delta \in [0, 1/2].$$

The problem of approximating the expansion profile is seemingly far-less tractable than approximating Φ_G itself. For instance, there is no known algorithm for the following easily stated decision problem concerning the expansion profile:

Problem 1. GAP-SMALL-SET EXPANSION (η, δ) Given a graph G and constants $\eta, \delta > 0$, distinguish whether

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

Spectral techniques fail in approximating the expansion of small sets in graphs. On one hand, even with the largest possible spectral gap, the Cheeger’s inequality cannot yield a lower bound greater than $1/2$ for the conductance $\Phi_G(\delta)$. More importantly, there exists graphs such as hypercube where there are sets S of half the vertices with small conductance ($\Phi(S) < \eta$), yet every sufficiently small set S satisfies $\Phi(S) \geq 1 - \eta$. This implies that Φ_G (and the second eigenvalue λ_2) do not yield any information about $\Phi_G(\delta)$ for small δ .

¹ In irregular graphs, it is more convenient to permit sets S within a range of sizes say $[\delta, 10\delta]$, since in an arbitrary (non-regular) graph there could be no sets satisfying $\mu(S) = \delta$.

Unique Games Conjecture.

The Unique Games Conjecture (UGC) of Khot [13] is among the central open problems in hardness of approximation, and has fueled many developments in the area in recent years. The UGC is shown to imply optimal inapproximability results for classic problems like MAX CUT [14], VERTEX COVER [15] and SPARSEST CUT [16] and constraint satisfaction problems [19].

For the sake of concreteness, we formally state the unique games conjecture here. The UNIQUE GAMES problem is defined as follows:

Definition 1. An instance of UNIQUE GAMES represented as $\Upsilon = (\mathcal{V}, \mathcal{E}, \Pi, [R])$ consists of a graph over vertex set \mathcal{V} with the edges \mathcal{E} between them. Also part of the instance is a set of labels $[R] = \{1, \dots, R\}$, and a set of permutations $\pi_{v \leftarrow w} : [R] \rightarrow [R]$ for each edge $e = (w, v) \in \mathcal{E}$. An assignment \mathcal{A} of labels to vertices is said to satisfy an edge $e = (w, v)$, if $\pi_{v \leftarrow w}(\mathcal{A}(w)) = \mathcal{A}(v)$. The objective is to find an assignment \mathcal{A} of labels that satisfies the maximum number of edges.

As is customary in hardness of approximation, one defines a gap-version of the UNIQUE GAMES problem as follows:

Problem 2. (UNIQUE GAMES $(R, 1 - \varepsilon, \eta)$) Given a UNIQUE GAMES instance $\Upsilon = (\mathcal{V}, \mathcal{E}, \Pi = \{\pi_{v \leftarrow w} : [R] \rightarrow [R] \mid e = (w, v) \in \mathcal{E}\}, [R])$ with number of labels R , distinguish between the following two cases:

– $(1 - \varepsilon)$ -satisfiable instances: There exists an assignment \mathcal{A} of labels that satisfies a $1 - \varepsilon$ fraction of edges.

–Instances that are not η -satisfiable: No assignment satisfies more than a η -fraction of the edges \mathcal{E} .

The Unique Games Conjecture asserts that the above decision problem is NP-hard when the number of labels is large enough. Formally,

CONJECTURE 1.2 (UNIQUE GAMES CONJECTURE [13]). *For all constants $\varepsilon, \eta > 0$, there exists large enough constant R such that UNIQUE GAMES $(R, 1 - \varepsilon, \eta)$ is NP-hard.*

While the implications of the conjecture are well-understood, there has been much slower progress towards its resolution. Results supporting the truth of UGC have been especially difficult to show. In particular, it is unknown whether many of the implications of UGC are equivalent to the conjecture. In other words, it is entirely consistent with existing literature that all the implications of UGC on problems like MAX CUT and VERTEX COVER hold, but the conjecture itself is false. More precisely, although the UNIQUE GAMES problem is known to efficiently reduce to classic problems like MAX CUT and VERTEX COVER, there are no known “reverse” reductions from these problems back to UNIQUE GAMES.

The only *reverse* reduction towards which there is any literature is the reduction from MAX CUT to UNIQUE GAMES. Note that the MAX CUT problem is a special case of UNIQUE GAMES over the binary alphabet $\{0, 1\}$. Hence a MAX CUT instance is readily reduced to a UNIQUE GAMES instance by using parallel repetition. With a sufficiently strong parallel repetition theorem (conjectured by Feige et al. [11]), this would yield a *reverse* reduction from MAX CUT to UNIQUE GAMES. Unfortunately, a strong parallel repetition theorem of this nature was shown not to hold by Raz [24]. Subsequent work by Barak et al. [7] almost entirely ruled out this approach to reduce MAX CUT to UNIQUE GAMES.

Expansion and Unique Games.

Vertex and edge expansion in graphs appear to be closely tied to the hard instances for linear and semidefinite programming relaxations. Many integrality gap instances have been constructed for problems like VERTEX COVER or MAX CUT against the linear programming hierarchies such as Lovász–Schröder and Sherali–Adams hierarchies (see [8, 25] and the references therein). Not only do most of these instances consist of expanding graphs but the arguments rely crucially on either vertex or edge expansion.

The situation is little bit more subtle in case of semidefinite programming. Semidefinite programs can approximate MAX CUT well on instances that have very good conductance (spectral gap). Hence, the SDP integrality gap instances known are graphs where small sets expand well, while the larger sets do not. Indeed, SDP integrality gap constructions for MAX CUT [10, 16, 21], VERTEX COVER [12], UNIQUE GAMES [16, 20] and SPARSEST CUT [16] all have near-perfect edge expansion for small sets.

In case of UNIQUE GAMES, not only do all known integrality gap instances have near-perfect edge expansion of small sets, even the analysis relies directly on this property. Furthermore, it is known that the best possible soundness for UNIQUE GAMES with label size R and completeness $1 - \varepsilon$ is a constant $\eta(R, \varepsilon)$ which is roughly $\frac{1}{R^{\varepsilon/2}}$. The constant $1 - \eta$ arises directly from the expansion of sets of size $\frac{1}{R}$ in a certain graph defined over the Gaussian space.

While this suggests that UNIQUE GAMES is closely tied to expansion of small sets in graphs, somewhat contrastingly, Arora et al. [4] show that UNIQUE GAMES is easy when the constraint graph involved is a good spectral expander, i.e., has a non-trivial spectral gap for the Laplacian.

Motivated by the above reasons, we investigate the connection between graph expansion and UNIQUE GAMES in this work.

1.1 Results

The main result of this work is a reduction from the problem of approximating expansion of small sets to the UNIQUE GAMES problem. The main implication of the result can be succinctly stated as follows.

Let us consider the following hardness assumption about the complexity of the GAP-SMALL-SET EXPANSION problem.

CONJECTURE 1.3. (*Gap-Small-Set Expansion Conjecture*) *For every $\eta > 0$, there exists δ such that the problem GAP-SMALL-SET EXPANSION (η, δ) is NP-hard.*

Then, an immediate consequence of the reduction presented in this work is,

THEOREM 1.4. *The GAP-SMALL-SET EXPANSION conjecture implies the Unique Games Conjecture.*

To the best of our knowledge, this is the first non-trivial “reverse” reduction from a natural combinatorial optimization problem to UNIQUE GAMES. On one hand, it connects the somewhat non-standard problem of UNIQUE GAMES to the much well-studied problem of approximating graph expansion. Furthermore, the result makes concrete the conspicuous presence of small set expansion in SDP integrality gaps for UNIQUE GAMES and related problems.

While a confirmation of UGC was known to imply optimal inapproximability results for fundamental problems, a

refutation of UGC was not known to formally imply any new algorithmic technique. An important implication of the above result is that a refutation of the UGC would yield an algorithm for approximating edge expansion (only in a certain regime) in graphs – a basic optimization problem. Now we state the main result of the paper from this algorithmic standpoint. To this end, we formally state a hypothesis that would emerge from a refutation of UGC.

HYPOTHESIS 1.5 (UNIQUE GAMES IS EASY). *There exists a constant $\varepsilon > 0$ and a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that given a UNIQUE GAMES instance Υ with n vertices and k labels, it is possible to distinguish between the cases $\text{opt}(\Upsilon) \geq 1 - \varepsilon$ and $\text{opt}(\Upsilon) \leq \varepsilon$ in time $n^{f(k)}$.*

THEOREM 1.6. *Suppose the above hypothesis (UNIQUE GAMES is easy) holds for a constant ε_0 and a function f , then there exists a function $g: [0, 1] \rightarrow \mathbb{N}$ such that given any graph G with n vertices, and it is possible to distinguish whether $\Phi_G(\delta) \leq \varepsilon_1$ or $\Phi_G(\delta) > 1 - \varepsilon_1$ for some absolute constant ε_1 in time $n^{g(\delta)}$.*

A somewhat stronger consequence of the UNIQUE GAMES is easy hypothesis follows using the parallel repetition theorem. The following corollary is an immediate consequence of the parallel repetition theorem of Rao [23] and our reduction from GAP-SMALL-SET EXPANSION to UNIQUE GAMES.

COROLLARY 1.7. *Suppose the hypothesis – UNIQUE GAMES instance is easy) holds for a constant ε_0 and a function f . Then, given a graph G with n vertices and parameters ε, δ such that $\varepsilon < \varepsilon_1$ for some absolute constant ε_1 , we can distinguish the following cases in time $n^{g(\varepsilon, \delta)}$ for some $g: [0, 1]^2 \rightarrow \mathbb{N}$.*

1. There exists $S \subseteq V$ with $\mu(S) = \delta$ and $\Phi(S) \leq \varepsilon$
2. Every set $S \subseteq V$ with $\mu(S) \leq 1500\delta/\varepsilon$ satisfies $\Phi(S) \geq 1500\sqrt{\varepsilon}$.

1.1.1 Towards an Equivalence

A natural question that arises from Theorem 1.4 is whether the Unique Games Conjecture is equivalent to the GAP-SMALL-SET EXPANSION conjecture. More specifically, could it be true that Unique Games Conjecture implies the GAP-SMALL-SET EXPANSION conjecture.

Showing a result of this nature amounts to obtaining a reduction from UNIQUE GAMES to GAP-SMALL-SET EXPANSION. Despite the large number of UG reductions and a fairly thorough understanding of how to construct reductions from UNIQUE GAMES, obtaining reductions to graph expansion problems is often problematic.

The main issue is that hardness reductions via local gadgets do not alter the global structure of the graph. For example, if the UNIQUE GAMES instance is disconnected, the resulting graph produced by a gadget reduction is also disconnected. Hence, to show a UG-hardness result for a global property such as expansion, it often seems necessary to assume the corresponding global property on the constraint graph of the UNIQUE GAMES. More specifically, in our case we require that the UNIQUE GAMES instance has good local expansion, in that sufficiently small sets have conductance close to 1. The formal statement of the modified Unique Games Conjecture with mild expansion on small sets is given below:

CONJECTURE 1.8. (UNIQUE GAMES WITH SMALL SET EXPANSION Conjecture) For every $\varepsilon > 0$, there exists δ such that for all $\eta > 0$ the following problem is NP-hard with a sufficiently large $R = R(\varepsilon, \eta)$:

Given an instance $\Upsilon = (\mathcal{V}, \mathcal{E}, \Pi = \{\pi_{v \leftarrow w} : [R] \rightarrow [R]\}, [R])$ of UNIQUE GAMES with alphabet size $R \geq R_0$ distinguish between the following cases:

–*Completeness* There exists an assignment $A : \mathcal{V} \rightarrow [R]$ that satisfies at least $1 - \varepsilon$ fraction of edges.

–*Soundness* Every assignment $A : \mathcal{V} \rightarrow [R]$ satisfies at most η fraction of edges and the constraint graph $(\mathcal{V}, \mathcal{E})$ satisfies the following expansion property:

–(Local Expansion Property) For every set S with $\mu(S) \leq \delta$, $\Phi(S) \geq 1 - \varepsilon$.

The above conjecture assumes fairly mild expansion in that sufficiently small sets have conductance close to 1. In fact, existing SDP integrality gap instances for UNIQUE GAMES (see [16, 20]) satisfy the above local expansion property. Under this stronger Unique Games Conjecture stated above, we show the following hardness result for GAP-SMALL-SET EXPANSION.

THEOREM 1.9. The Unique Games with Small Set Expansion conjecture implies the GAP-SMALL-SET EXPANSION conjecture.

As an immediate consequence of the UG hardness reduction we also obtain super-constant hardness of approximation for the k -DENSEST SUBGRAPH PROBLEM.

THEOREM 1.10. Assuming the Unique Games with Small Set Expansion conjecture, it is NP-hard to approximate the k -DENSEST SUBGRAPH PROBLEM to any constant factor.

Using standard techniques, the UG-hardness reductions can be used to construct SDP integrality gaps for the GAP-SMALL-SET EXPANSION and the k -DENSEST SUBGRAPH PROBLEM. Due to space constraints, we defer the proofs of Theorem 1.9 and Theorem 1.10 to the full version.

1.2 Unique Games with Small Set Expansion are Easy

Further exploring the connection between UNIQUE GAMES and small set expansion, we show that the UNIQUE GAMES problem is easy when there is sufficiently high expansion of small sets. This result generalizes the work of [4] which showed that UNIQUE GAMES is easy when the associated constraint graph is a good spectral expander.

We will use an analytic notion of a graph being a small set expander for this purpose. This is akin to spectral expansion of a graph instead of combinatorial expansion.

Definition 2. For a graph G and $\delta > 0$, define $\tilde{\Lambda}_G(\delta)$ as the minimum of

$$\mathbb{E}_{(u,v) \sim G} \|\mathbf{X}_u - \mathbf{X}_v\|^2$$

over all collection of vectors $\{\mathbf{X}_v\}_{v \in V}$ that satisfy the constraints

$$\mathbb{E}_{u,v \sim V} \langle \mathbf{X}_u, \mathbf{X}_v \rangle \leq \delta, \quad (1.1)$$

$$\mathbb{E}_{u \sim V} \|\mathbf{X}_u\|^2 = 1. \quad (1.2)$$

Unlike the combinatorial small set expansion $\Phi_G(\delta)$, the quantity $\tilde{\Lambda}_G(\delta)$ is efficiently computable, thus making it easy to recognize graphs that satisfy the required expansion property. Raghavendra et al. [22] show that this parameter approximates the combinatorial expansion profile $\Phi_G(\delta)$ as follows:

$$\tilde{\Lambda}_G(\delta) \leq \Phi_G(\delta) \leq O\left(\tilde{\Lambda}_G(\delta/2) \log(1/\delta)\right)^{1/2}.$$

In this work, we show that UNIQUE GAMES is easy if the constraint graph G has sufficiently large $\tilde{\Lambda}_G(\delta)$. The formal statement of the result is as follows:

THEOREM 1.11. Let Υ be unique game with SDP value $1 - \varepsilon$ and constraint graph G . Suppose there exists $\delta > 0$ such that

$$\tilde{\Lambda}_G(\delta) > 10^6 \varepsilon \log(1/\varepsilon \delta),$$

then the integral value of Υ is at least $\delta/10$.

COROLLARY 1.12. If the constraint graph G associated with a unique games instance Υ satisfies $\tilde{\Lambda}_G(\delta) > 10^6 \varepsilon \log(1/\varepsilon \delta)$, then it is easy to distinguish whether instance has value $1 - \varepsilon$ or less than δ .

Independent of this work, Arora et al. [3] also show that UNIQUE GAMES is easy if the underlying constraint graph has sufficient local expansion using different techniques. Due to space constraints, we defer the proof of Theorem 1.11 to the full version.

Relation with Conjecture 1.8: For $\delta < 2^{-1/\varepsilon}$ the quantity $10^6 \varepsilon \log(1/\varepsilon \delta) > 1$, while $\tilde{\Lambda}_G(\delta) \leq 1$ for all graphs G and all $\delta > 0$. Hence, Theorem 1.11 is not applicable on graphs with expansion of sets of size $\delta < 2^{-1/\varepsilon}$. In Conjecture 1.8, there is no lower bound on the size δ of sets with local expansion. Consequently, Theorem 1.11 only yields an upper bound on the choice of the set size δ in Conjecture 1.8.

2. PRELIMINARIES

Notation: Unless otherwise stated, all graphs considered henceforth can be assumed to be regular unweighted graphs. The proofs and results generalize to arbitrary weighted graphs, but we defer its discussion to the full version.

Given a graph $G = (V, E)$, we will use $v \sim V$ to denote a vertex sampled from the probability distribution given by degrees $\deg(v)$ (in this case, the uniform distribution). Furthermore, for a set $S \subseteq V$, $\mu(S) = \mathbb{P}_{v \sim V} \{v \in S\}$. Similarly, the notation $e \sim E$ will denote an edge sampled from a distribution proportional to the weights (the uniform distribution in unweighted graphs). We denote by $E(A, B)$ the set of edges with one endpoint in A and the other endpoint in B . For a subset of edges $E' \subseteq E$, $|E'|$ will denote the sum of the edge weights within E' . Therefore for a subset of edges E' , $\mathbb{P}_{e \sim E} [e \in E'] = |E'|/|E| = 2|E'|$.

Definition 3. For a vertex subset $S \subseteq V$, we define

$$\partial(S) \stackrel{\text{def}}{=} |E(S, V \setminus S)| \quad \text{and} \quad \mu(S) \stackrel{\text{def}}{=} \sum_{s \in S} \deg(s).$$

Definition 4. (Conductance) The conductance/Cheeger's constant associated with a subset $S \subseteq V$ is given by $\Phi(S) = \frac{\partial(S)}{\mu(S)}$. The conductance/Cheeger's constant for the graph G is $\Phi_G = \min_{\mu(S) \leq \frac{1}{2}} \Phi(S)$.

Due to space constraints we omit the proofs of the following simple facts.

FACT 2.1. For all $a, b \in [0, 1]$, $|a^2 - b^2| \leq 2a + 2b - 4ab$.

LEMMA 2.2 (GLORIFIED MARKOV INEQUALITY). Let Ω be a probability space and let $X, Y: \Omega \rightarrow \mathbb{R}_+$ be two jointly distributed non-negative random variables over Ω . Suppose $\mathbb{E}X \leq \gamma \mathbb{E}Y$. Then, there exists $\omega \in \Omega$ such that $X(\omega) \leq 2\gamma Y(\omega)$ and $Y(\omega) \geq \mathbb{E}Y/2$.

Partial 2-prover games: A general 2-prover game Γ is specified by a vertex set V , an edge set E , an alphabet Σ , and a collection of predicates $\{\Pi_{u,v}\}$ indexed by vertex pairs $u, v \in V$. The value of the game Γ is defined as the maximum success probability

$$\mathbb{P}_{uv \in E} \left\{ (F(u), F(v)) \in \Pi_{u,v} \right\},$$

over all strategies $F: V \rightarrow \Sigma$.

Given a 2-prover game Γ , a *partial game* is one in which the provers are permitted to *refuse* to answer the question. The provers win only if both of them refuse to answer the question or else they both answer correctly as per the game Γ . To preclude the trivial strategy that always refuses to answer questions, we require that the provers answer at least an α -fraction of the questions for some constant $\alpha > 0$. Formally, we define the α -*partial value* of a 2-prover game Γ as the maximum success probability

$$\mathbb{P}_{uv \in E} \left\{ (F(u), F(v)) \in \Pi_{u,v} \right\} / \mathbb{P}_{u \sim V} \left\{ F(u) \neq \perp \right\},$$

over all α -partial strategies $F: V \rightarrow \Sigma \cup \{\perp\}$. Here, \perp is a designated symbol (not a member of Σ) and an α -*partial strategy* is an assignment $F: V \rightarrow \Sigma \cup \{\perp\}$ such that $\mathbb{P}_{u \sim V} \{F(u) \neq \perp\} \geq \alpha$. As usual, the notation $v \sim V$ means that v is sampled with probability proportional to its degree.

Organization: In the remainder of the paper, we present the reduction from GAP-SMALL-SET EXPANSION to UNIQUE GAMES and its analysis. For conceptual clarity, we subdivide the reduction in to two parts: In the first part (Section 3), we reduce GAP-SMALL-SET EXPANSION to a *partial* Unique games, and then in Section 4 exhibit a generic reduction from *partial* 2-prover games to a corresponding traditional 2-prover game. Finally, in Section 5 we wrap up the proof of the main result of this work - Theorem 1.4.

3. FROM GAP-SMALL-SET-EXPANSION TO PARTIAL UNIQUE GAMES

In this section, we outline the key ideas of the reduction from GAP-SMALL-SET EXPANSION to Partial UNIQUE GAMES.

First Attempt.

Let $G = (V, E)$ be an instance of the GAP-SMALL-SET EXPANSION (η, δ) problem. For the sake of simplicity, let G be a d -regular unweighted graph. We define the following unique game:

- Fix $R = \lceil \frac{1}{\delta} \rceil$. The referee/verifier picks R edges $M = \{(u_1, v_1), \dots, (u_R, v_R)\}$ uniformly at random from G .
- The referee sends a random permutation of the tuple $U = (u_1, \dots, u_R)$ to the first prover, and a random permutation of the tuple $V = (v_1, \dots, v_R)$ to second prover.

-The provers are expected to pick one of the vertices of the tuple they receive. Specifically, the answer set of the first prover is $\{u_1, \dots, u_R\}$, while the second prover's answer set is $\{v_1, \dots, v_R\}$.

-The provers win if they pick two vertices (u_i, v_i) corresponding to some edge in the set M .

Completeness: Let us suppose there exists a set S of δn vertices such that $\Phi(S) \leq \varepsilon$. In this case, the provers can use the following simple strategy: If exactly one vertex from S appears in the tuple, return that vertex, else refuse to answer the question.

The set S is of size δn and a question $U \in V^R$ has $\lceil \frac{1}{\delta} \rceil$ vertices. Therefore, with constant probability (at least 0.01) exactly one of the vertices from S appears in the question U , and the provers answer the question. Consequently, the above strategy is a valid partial strategy. Observe that if the set S has small conductance $\Phi(S) \leq \varepsilon$, a random edge incident on S is with high probability completely contained within S . In other words, if $u_i \in S$ then with very high probability its neighbour v_i also belongs to S . This implies that whenever the first prover decides to answer the vertex u_i , the second prover also answers v_i with very high probability. Hence a small non-expanding set S translates directly in to a *good* partial strategy for the unique game.

Soundness: Let $F: V^R \rightarrow [R]$ be a strategy for the two provers that succeeds with probability at least $\frac{1}{2}$. For a vertex sequence $U \in V^{R-1}$, let $U +_i x \in V^R$ denote the vertex sequence obtained by inserting x at index i .

For every vertex sequence $U \in V^{R-1}$ and an index $i \in [R]$, let us define a $\{0, 1\}$ -valued function $F_U^{(i)}$ as follows:

$$F_U^{(i)}(x) = 1 \text{ if } F(U +_i x) = i \text{ else it is } 0$$

Specifically, $F_U^{(i)}$ is the indicator function of the set of vertices x such that the strategy decides to pick up the vertex x , when it is inserted at the i^{th} location in U .

Notice that, for the strategy proposed in the completeness case, for every setting of U , the set $F_U^{(i)}$ is either the non-expanding set S or the empty set. Extrapolating from here, it is natural to look for non-expanding sets by considering the set $F_U^{(i)}$ for a random choice of $U \in V^{R-1}$ and i . This is the basic intuition behind the soundness analysis presented in Section 3.1.

Over a random choice of $U \in V^{R-1}$ and i , the expected size of $F_U^{(i)}$ is indeed roughly $\Theta(\frac{1}{R}) = \Theta(\delta)$. However, it could be possible that with very high probability over the choice of U and i , $F_U^{(i)}$ is either too large or too small a set. To rule out this possibility, we modify the above unique game by including random noise in to the questions. More specifically, the referee changes each vertex of the questions U, V independently with probability ε , before sending it to the provers. In the modified game, we bound the second moment of the random variable - the size of $F_U^{(i)}$ over a random choice of U and i .

Formal Reduction.

Here we present the details of the reduction from GAP-SMALL-SET EXPANSION problem to Partial UNIQUE GAMES.

For a vertex $v \in V$ and $\varepsilon > 0$, we define a distribution $\mathcal{N}_\varepsilon(v)$ on V as follows: with probability $1 - \varepsilon$, we output $u := v$ and with probability ε , we output $u \sim V$. For a vertex

sequence v_1, \dots, v_R , the notation $\mathcal{N}_\varepsilon(v_1, \dots, v_R)$ refers to the product distribution $\mathcal{N}_\varepsilon(v_1) \otimes \dots \otimes \mathcal{N}_\varepsilon(v_R)$ on V^R .

For an edge $e \in E$, we define $\mathcal{D}(e)$ to be the distribution that picks a random endpoint of e (uniformly). For a edge sequence e_1, \dots, e_R , the distribution $\mathcal{D}(e_1, \dots, e_R)$ on V^R is the product of the distributions $\mathcal{D}(e_1), \dots, \mathcal{D}(e_R)$. For a permutation $\pi: [R] \rightarrow [R]$ and a sequence $A \in V^R$, we write $A' = \pi.A$ to denote the permutation of A according to π , i.e., $A'_{\pi(i)} = A_i$ for all $i \in [R]$.

Let $R \in \mathbb{N}$. We define a UNIQUE GAMES instance $\Upsilon = \Upsilon_{R,\varepsilon}(G)$ with vertex set V^R and alphabet $\Sigma = [R]$. The edge constraints in Υ are given by the tests performed by the following verifier.

<p>SSE-TO-UG Reduction</p> <p>Let $F: V^R \rightarrow \Sigma$ be an assignment to the instance Υ. The value of the assignment is the success probability of the following test:</p> <ol style="list-style-type: none"> 1. Sample R random edges e_1, \dots, e_R from E and let $M := (e_1, \dots, e_R)$. 2. Sample $A \sim \mathcal{D}(M)$ and $B \sim \mathcal{D}(M)$. 3. Sample $\tilde{A} \sim \mathcal{N}_\varepsilon(A)$ and $\tilde{B} \sim \mathcal{N}_\varepsilon(B)$. 4. Sample two permutations $\pi_1, \pi_2: [R] \rightarrow [R]$. 5. Output Success if $\pi_1^{-1} \left(F(\pi_1.\tilde{A}) \right) = \pi_2^{-1} \left(F(\pi_2.\tilde{B}) \right). \quad (3.1)$
--

In the remainder of the section, we will show the following theorem.

THEOREM 3.1. *For $\eta < \frac{1}{24}$, given a graph $G = (V, E)$, the 3 reduction produces an instance Υ of UNIQUE GAMES such that:*

–**Completeness:** *If $\Phi(S) \leq \eta$ for some S with $\mu(S) \in [\frac{1}{20R}, \frac{1}{R}]$ then there exists a $1/10\varepsilon$ -partial strategy of value at least $1 - 40\varepsilon\eta - 4\varepsilon$.*

–**Soundness:** *If there exists a α -partial strategy with value $1 - \eta$ for the UNIQUE GAMES instance Υ , then there exists a set $S \subseteq V$ with $\Phi(S) \leq 96\eta$ and $\mu(S) \in [\frac{\alpha}{4R}, \frac{6}{\varepsilon\eta R}]$.*

Due to space limitations, we defer the somewhat straightforward proof of the completeness case to the full version. The soundness claim of the above theorem is proven in the next section (Lemma 3.3).

3.1 Soundness

Let $F: V^R \rightarrow [R] \cup \{\perp\}$ be a partial assignment for the unique game Υ . For $U \in V^{R-1}$ and $x \in V$, we let $f(U, x) \in [0, 1]$ denote the probability that F selects the coordinate of x after we place it at a random position of U and permute the sequence randomly, i.e.,

$$f(U, x) \stackrel{\text{def}}{=} \mathbb{P}_{i \in [R]} \{F(\pi.(U +_i x)) = \pi(i)\}. \quad (3.2)$$

Here, $U +_i x$ denotes the vertex sequence in V^R obtained from U by inserting x as the i -th coordinate (the original coordinates $i, \dots, R - 1$ of U are moved by one to the right).

For $M \in E^{R-1}$, we define a function $f_M: V \rightarrow [0, 1]$ as follows

$$f_M(x) \stackrel{\text{def}}{=} \mathbb{E}_{U \sim \mathcal{D}(M)} \mathbb{E}_{(\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)} f(\tilde{U}, \tilde{x}). \quad (3.3)$$

We establish the following three relatively straight-forward properties of the functions f_M .

PROPOSITION 3.2. *Suppose that F has value $1 - \eta$ for the game Υ and that F differs from \perp on exactly an α fraction of the inputs, i.e., $\alpha := \mathbb{P}_{X \sim V^R} \{F(X) \neq \perp\}$. Then the following hold,*

$$\mathbb{E}_{M \in E^{R-1}} \mathbb{E}_{x \sim V} f_M(x) = \frac{\alpha}{R}, \quad (3.4)$$

$$\mathbb{E}_{x \sim V} f_M(x) \leq \frac{2}{\varepsilon R} \quad \forall M \in E^{R-1} \quad (3.5)$$

$$\mathbb{E}_{M \in E^{R-1}} \mathbb{E}_{e \in E} \left(\mathbb{E}_{x \sim \mathcal{D}(e)} f_M(x) \right)^2 = (1 - \eta) \frac{\alpha}{R}. \quad (3.6)$$

Before proving the proposition, we will first show that the existence of such functions $\{f_M\}$ as outlined in the proposition, suffices to finish the soundness analysis.

LEMMA 3.3. *Suppose there exists a collection of functions $\{f_M\}$ satisfying the three conditions of 3.2 for some $\eta \leq \frac{1}{24}$. Then there exists a set $S \subseteq V$ with $\mu(S) \in [\frac{\alpha}{4R}, \frac{6}{\varepsilon R}]$ and $\Phi(S) \leq 96\eta$*

PROOF. Let X and Y be the following non-negative random variables over the probability space E^{R-1} ,

$$X(M) := \mathbb{E}_{xy \in E} f_M(x)f_M(y) \quad Y(M) := \mathbb{E}_{x \sim V} f_M(x).$$

Since $X \leq Y$ (pointwise), $X' = Y - X$ is also a non-negative random variable over E^{R-1} . The conditions (3.6) and (3.4) imply that $\mathbb{E} X' \leq 2\eta \mathbb{E} Y$ and $\mathbb{E} Y = \alpha/R$. Hence, Lemma 2.2 (Glorified Markov Inequality) asserts that there exists an edge sequence $M^* \in E^{R-1}$ such that $X'(M^*) \leq 4\eta Y(M^*)$ and $Y(M^*) \geq \mathbb{E} Y/2$. In other words,

$$\mathbb{E}_{xy \in E} f_{M^*}(x)f_{M^*}(y) \geq (1 - 4\eta) \mathbb{E}_{x \sim V} f_{M^*}(x), \quad (3.7)$$

$$\mathbb{E}_{x \sim V} f_{M^*}(x) \geq \frac{\alpha}{2R}. \quad (3.8)$$

By condition (3.5), we can also upper bound the expectation of f_{M^*} ,

$$\mathbb{E}_{x \sim V} f_{M^*}(x) \leq \frac{2}{\varepsilon R}. \quad (3.9)$$

Lemma 3.4 shows that we can round the function f_{M^*} to a cut S^* that satisfies

$$\Phi(S^*) \leq 96\eta, \quad \text{and} \quad \mu(S^*) \in \left[\frac{\alpha}{4R}, \frac{6}{\varepsilon R} \right].$$

□

LEMMA 3.4. *Let G be a graph with vertex set V and edge set E . Suppose $f: V \rightarrow [0, 1]$ is a bounded function on V such that*

$$\mathbb{E}_{xy \in E} f(x)f(y) \geq (1 - \eta) \mathbb{E}_{x \sim V} f(x), \quad (3.10)$$

for $\eta < \frac{1}{6}$. Let $\beta := \mathbb{E}_{x \sim V} f(x)$. Then, there exists a vertex set $S^* \subseteq V$ such that $\mu(S^*) \in [\frac{\beta}{2}, 3\beta]$ and $\Phi(S^*) \leq 24\eta$.

PROOF. Define a distribution over sets S given by the following sampling procedure:

1. sample a threshold $t \in [1/9, 4/9]$ uniformly at random,
2. output the set $S := \{x \mid f^2(x) > t\}$.

Since $\mathbb{E} f = \beta$ the set S always satisfies

$$\mu(S) \leq \frac{\mathbb{E}[f]}{1/3} = 3\beta. \quad (3.11)$$

Note that condition (3.10) implies

$$(1-\eta)\beta \leq \mathbb{E}_{xy \sim E} f(x)f(y) \leq \mathbb{E}_{xy \sim E} \frac{1}{2}(f^2(x)+f^2(y)) = \mathbb{E}_{x \sim V} f^2.$$

Therefore we have $\mathbb{E}_{x \sim V} f(1-f) \leq \eta\beta$. Rewriting the previous inequality,

$$\begin{aligned} \eta\beta &\geq \mathbb{P}[f < 2/3] \mathbb{E}_{x \sim V} [f(1-f) \mid f < 2/3] \\ &\geq \frac{1}{3} \mathbb{P}[f < 2/3] \mathbb{E}_{x \sim V} [f \mid f < 2/3]. \end{aligned}$$

This implies a lower bound on the size of the set S as follows,

$$\begin{aligned} \mu(S) &\geq \mathbb{P}[f \geq 2/3] \geq \mathbb{P}[f \geq 2/3] \mathbb{E}_{x \sim V} [f \mid f \geq 2/3] \\ &= \mathbb{E}[f] - \mathbb{P}[f < 2/3] \mathbb{E}_{x \sim V} [f \mid f < 2/3] \geq \beta(1-3\eta). \end{aligned} \quad (3.12)$$

Finally, we estimate the boundary of the set S .

$$\begin{aligned} \partial(S) &= \mathbb{E}_{xy \sim E} [|\mathbf{1}_S(x) - \mathbf{1}_S(y)|] \leq \mathbb{E}_{xy \sim E} [3|f^2(x) - f^2(y)|] \\ &\leq 12 \mathbb{E}_x f - 12 \mathbb{E}_{xy \sim E} f(x)f(y) \quad (\because 0 \leq f \leq 1 \text{ and Fact 2.1}) \\ &\leq 12\eta\beta \end{aligned} \quad (3.13)$$

The claim follows from (3.11), (3.12) and (3.13). \square

3.2 Proof of Proposition 3.2

Let $F: V^R \rightarrow [R] \cup \{\perp\}$ be a partial strategy of value $1-\eta$ for the game Υ . Suppose F differs from \perp on exactly an α fraction of inputs, i.e., $\mathbb{P}_{X \sim V} \{F(X) \in [R]\} = \alpha$. We prove the three conditions of 3.2 in the following lemmas.

LEMMA 3.5. $\mathbb{E}_{M \in E^{R-1}} \mathbb{E}_{x \sim V} f_M(x) = \frac{\alpha}{R}$.

PROOF. We unroll the definition of $f_M(x)$,

$$\begin{aligned} &\mathbb{E}_{M \in E^{R-1}} \mathbb{E}_{x \sim V} f_M(x) \\ &= \mathbb{E}_{M \in E^{R-1}} \mathbb{E}_{x \sim V, (\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)} \mathbb{P}_\pi \left\{ F(\pi.(\tilde{U} + \tilde{x})) = \pi(i) \right\} \\ &= \mathbb{P}_{X \sim V^R, r \in [R]} \{F(X) = r\} = \frac{\alpha}{R}. \end{aligned}$$

In the last equality, we use that the joint distribution of $\pi.(\tilde{U} + \tilde{x})$ and $\pi(i)$ is the same as the joint distribution of X and r . \square

LEMMA 3.6. $\mathbb{E}_{M \in E^{R-1}} \mathbb{E}_{e \in E} \left(\mathbb{E}_{x \sim \mathcal{D}(e)} f_M(x) \right)^2 = (1-\eta) \frac{\alpha}{R}$.

PROOF. For an edge sequence $M \in E^R$ and $r \in [R]$, the notation M^{-r} refers to the edge sequence in E^{R-1} obtained from M by removing the r -th coordinate. Similarly, M_r denotes the r -th coordinate of M . Recall that F is a partial

strategy with value $(1-\eta) \cdot \alpha$ on the game Υ . Expressing this fact we get,

$$\begin{aligned} (1-\eta) \frac{\alpha}{R} &= \mathbb{E}_{\substack{M \in E^R, A, B \sim \mathcal{D}(M) \\ \tilde{A} \sim \mathcal{N}_\varepsilon(A), \tilde{B} \sim \mathcal{N}_\varepsilon(B) \\ r \in [R]}} \mathbb{P}_{\pi_1, \pi_2} \left\{ F(\pi_1.\tilde{A}) = \pi_1(r) \wedge F(\pi_2.\tilde{B}) = \pi_2(r) \right\} \\ &= \mathbb{E}_{\substack{M \in E^R, \\ r \in [R]}} \left(\mathbb{E}_{A \sim \mathcal{D}(M)} \mathbb{E}_{\tilde{A} \sim \mathcal{N}_\varepsilon(A)} \mathbb{P}_\pi \left\{ F(\pi.\tilde{A}) = \pi(r) \right\} \right)^2 \\ &= \mathbb{E}_{\substack{M \in E^R, \\ r \in [R]}} \left(\mathbb{E}_{\substack{x \sim \mathcal{D}(M_r) \\ U \sim \mathcal{D}(M^{-r}) \\ (\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)}} \mathbb{P}_\pi \left\{ F(\pi.(\tilde{U} + r \tilde{x})) = \pi(r) \right\} \right)^2 \\ &= \mathbb{E}_{\substack{M \in E^R, \\ r \in [R]}} \left(\mathbb{E}_{x \sim \mathcal{D}(M_r)} \mathbb{E}_{U \sim \mathcal{D}(M^{-r})} \mathbb{E}_{(\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)} f(\tilde{U}, \tilde{x}) \right)^2 \\ &= \mathbb{E}_{\substack{M \in E^R, \\ r \in [R]}} \left(\mathbb{E}_{x \sim \mathcal{D}(M_r)} f_{M^{-r}}(x) \right)^2 \\ &= \mathbb{E}_{M' \in E^{R-1}} \mathbb{E}_{e \in E} \left(\mathbb{E}_{x \sim \mathcal{D}(e)} f_{M'}(x) \right)^2. \end{aligned}$$

In the second equality, we use the fact that the events $F(\pi_1.\tilde{A}) = \pi_1(r)$ and $F(\pi_2.\tilde{B}) = \pi_2(r)$ are independent given $M \in E^R$ and $r \in [R]$. In the third equality, e_r denotes the edge in the r -th coordinate of M . \square

LEMMA 3.7. Suppose εR is sufficiently large. Then for every $M \in E^{R-1}$,

$$\mathbb{E}_{x \sim V} f_M(x) \leq \frac{2}{\varepsilon R}.$$

PROOF. As before, we first unroll the definition of $f_M(x)$

$$\begin{aligned} \mathbb{E}_{x \sim V} f_M(x) &= \mathbb{E}_{U \sim \mathcal{D}(M)} \mathbb{E}_{x \sim V} \mathbb{E}_{(\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)} f(\tilde{U} + \tilde{x}, \tilde{x}) \\ &= \mathbb{E}_{U \sim \mathcal{D}(M)} \mathbb{E}_{x \sim V} \mathbb{E}_{(\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)} \mathbb{P}_\pi \left\{ F(\pi.(\tilde{U} + \tilde{x})) = \pi(i) \right\} \end{aligned} \quad (3.14)$$

To estimate this expectation, we generate random variables $\pi.(\tilde{U} + \tilde{x})$ and $\pi(i)$ in a different fashion. Let \circledast denote a placeholder symbol (not a member of V). Consider the following randomized procedure generating $A \in V^R$, $r \in [R]$:

1. Sample $U \sim \mathcal{D}(M)$, a permutation $\pi: [R] \rightarrow [R]$ and an index $i \in [R]$. Set $B := \pi.(U + i \circledast)$.
2. Generate $B' \in (V \cup \{\circledast\})^R$ by changing each coordinate of B randomly to \circledast with probability ε . Let $\mathcal{I} \subseteq [R]$ denote the set of coordinates of B' that contain \circledast .
3. Generate $A \in V^R$ by replacing all the placeholders in B' by vertices sampled from V .
4. Sample $r \in \mathcal{I}$ uniformly at random.

We claim that the joint distribution of A and r is the same as the joint distribution of $\pi.(\tilde{U} + \tilde{x})$ and $\pi(i)$ in (3.14). By construction, the joint distribution of A and $\pi(i)$ is the same as the joint distribution of $\pi.(\tilde{U} + \tilde{x})$ and $\pi(i)$ in (3.14) (with an obvious coupling). From the description of the construction, it is also clear that the r -th coordinate of A has

the same distribution as the $\pi(i)$ -th coordinate of A (both are sampled from V).

Since the pair (A, r) is identical in distribution to the pair $(\pi.(\tilde{U} +_i x), \pi(i))$, we can finish the proof of the lemma as follows

$$\mathbb{E}_{x \sim V} f_M(x) = \mathbb{E}_{\substack{U \sim \mathcal{D}(M), x \sim V \\ (\tilde{U}, \tilde{x}) \sim \mathcal{N}_\varepsilon(U, x)}} \mathbb{P}_{i \in [R]} \left\{ F(\pi.(\tilde{U} +_i \tilde{x})) = \pi(i) \right\}$$

(by (3.14))

$$\begin{aligned} &= \mathbb{P}_{A, r} \left\{ F(A) = r \right\} \quad (A \text{ and } r \text{ generated as above}) \\ &= \mathbb{E}_{A, \mathcal{I}} \mathbb{P}_{r \in \mathcal{I}} \left\{ F(A) = r \right\} \quad (A \text{ and } \mathcal{I} \text{ generated as above}) \\ &= \mathbb{E}_{A, \mathcal{I}} \frac{1}{|\mathcal{I}|} \mathbb{1}_{\{F(A) \in \mathcal{I}\}} \\ &\leq \mathbb{E}_{t \sim \text{Binom}(R-1, \varepsilon)} \frac{1}{t+1} \quad (\because |\mathcal{I}| \sim 1 + \text{Binom}(R-1, \varepsilon)) \\ &\leq \frac{1.5}{\varepsilon R} + \mathbb{P}_{t \sim \text{Binom}(R-1, \varepsilon)} \{t+1 \leq \varepsilon R/1.5\} \\ &\leq \frac{1.5}{\varepsilon R} + 2^{-\Omega(\varepsilon R)} \quad (\text{using Chernoff bound}) \\ &\leq \frac{2}{\varepsilon R} \quad (\text{for sufficiently large } \varepsilon R) \end{aligned}$$

□

4. FROM PARTIAL 2-PROVER GAMES TO TOTAL 2-PROVER GAMES

We show a general reduction from a partial 2-prover game Γ to a corresponding 2-prover game Γ' .

Recall that in the partial game Γ , the provers are allowed to refuse to answer questions. However, in the game Γ' no such choice should be available. To achieve this, the referee gives multiple questions from Γ simultaneously and the provers have to answer any question of their choice.

Specifically, the questions in Γ' consists of a sequence of c questions from the game Γ . The provers are required to choose one of the c questions to answer, and return the answer. The provers win only if they pick the corresponding pair of questions (say the i^{th} question), and also answer the question correctly as per the game Γ . Formally, we will show the following theorem in the remainder of the section.

THEOREM 4.1. *For all positive integers c , given a 2-prover game Γ with n vertices, there is a reduction to another 2-prover game Γ' running in time $n^{O(c)}$ such that:*

- Completeness** *If the α -partial value of Γ is at least $1 - \varepsilon$, then the value of Γ' is at least $1 - \varepsilon - e^{-\alpha \cdot c}$.*
- Soundness** *If the value of Γ' is at least $1 - \eta$, then the $1/3c$ -partial value of Γ is at least $1 - 12\eta$.*

Furthermore, the reduction preserves the uniqueness property of the games.

We will assume that the edge set E of the game Γ does not contain any self-loops. We denote by E_\circlearrowleft the following distribution over vertex pairs $(u, v) \in V^2$: sample a random edge $e \in E$ and let u and v be (independent) random endpoints of e (in particular, $u = v$ with probability $1/2$). We add predicates $\Pi_{u, u} = \{(a, a) \mid a \in \Sigma\}$ for all self-loops $(u, u) \in V^2$ to the collection $\{\Pi_{u, v}\}$. For a parameter $c \in \mathbb{N}$, let Γ' be the following 2-prover game:

Reduction from a partial 2-prover game Γ to a total 2-prover game Γ'

1. Sample c vertex pairs $(u_1, v_1), \dots, (u_c, v_c) \sim E_\circlearrowleft$.
2. Send the vertex sequence u_1, \dots, u_c to the first prover and the sequence v_1, \dots, v_c to the second prover.
3. Let (i, a) and (j, b) be their answers.
4. The provers win if $i = j$ and $(a, b) \in \Pi_{u_i, v_j}$.

We observe that the reduction preserves the uniqueness property.

OBSERVATION 4.2. *If Γ is a unique game, then Γ' is a unique game as well.*

We will show [Theorem 4.1](#) in two parts: [Lemma 4.3](#) and [Lemma 4.4](#).

LEMMA 4.3 (COMPLETENESS). *If the α -partial value of Γ is at least $1 - \varepsilon$, then the value of Γ' is at least $1 - \varepsilon - e^{-\alpha \cdot c}$.*

PROOF. Let $F: V \rightarrow \Sigma \cup \{\perp\}$ be an α -partial strategy of value $1 - \varepsilon$. We consider the following strategy F' for Γ' .

$$F'(u_1, \dots, u_c) := \begin{cases} (i, a) & \text{if } (F(u_i) \in \Sigma) \wedge (\forall j < i, F(u_j) = \perp), \\ (1, 1) & \text{if } F(u_1) = \dots = F(u_c) = \perp. \end{cases}$$

(In words, the prover answers with the first answer in the list $F(u_1), \dots, F(u_c)$ (ignoring \perp). If the partial strategy refuses to answer on all inputs u_1, \dots, u_c , then the prover returns an arbitrary answer.)

Let $u_1 v_1, \dots, u_c v_c$ be a sequence of c vertex pairs, independently drawn from E_\circlearrowleft . The probability of the event $F(u_1) = \dots = F(u_c)$ is at most $(1 - \alpha)^c \leq e^{-\alpha c}$. Let us condition on the complementary event, i.e., the event that $F(u_i) \neq \perp$ for at least one coordinate $i \in [c]$. Let i_0 be the first coordinate such that $F(u_{i_0}) \neq \perp$ or $F(v_{i_0}) \neq \perp$. The winning probability of the provers (conditioned on the event that $F(u_i) \neq \perp$ for at least one coordinate $i \in [c]$) is equal to

$$\begin{aligned} &\mathbb{P} \left\{ (F(u_{i_0}), F(v_{i_0})) \in \Pi_{u_{i_0}, v_{i_0}} \right\} \\ &= \mathbb{P}_{uv \sim E_\circlearrowleft} \left\{ (F(u), F(v)) \in \Pi_{u, v} \mid F(u) \in \Sigma \vee F(v) \in \Sigma \right\} \\ &= \frac{\mathbb{P}_{uv \sim E_\circlearrowleft} \left\{ (F(u), F(v)) \in \Pi_{u, v} \right\}}{\mathbb{P}_{uv \sim E_\circlearrowleft} \left\{ F(u) \in \Sigma \vee F(v) \in \Sigma \right\}} \\ &= \frac{\mathbb{P}_{uv \sim E_\circlearrowleft} \left\{ (F(u), F(v)) \in \Pi_{u, v} \right\}}{2 \mathbb{P}_{u \sim V} \left\{ F(u) \in \Sigma \right\} - \mathbb{P}_{uv \sim E_\circlearrowleft} \left\{ F(u), F(v) \in \Sigma \right\}}. \end{aligned}$$

Without loss of generality, we may assume $\mathbb{P}_{u \sim V} \{F(u) \in \Sigma\} = \alpha$. Since the partial strategy F' has value $1 - \varepsilon$, we have $\mathbb{P}_{uv \in E_\circlearrowleft} \{(F(u), F(v)) \in \Pi_{u, v}\} \geq 1/2 \cdot (1 - \varepsilon)\alpha + 1/2 \cdot \alpha$. Therefore, also $\mathbb{P}_{uv \in E_\circlearrowleft} \{F(u), F(v) \in \Sigma\} \geq (1 - \varepsilon/2)\alpha$. We can conclude

$$\begin{aligned} &\mathbb{P}_{uv \sim E_\circlearrowleft} \left\{ (F(u), F(v)) \in \Pi_{u, v} \mid F(u) \in \Sigma \vee F(v) \in \Sigma \right\} \\ &\geq \frac{(1 - \varepsilon/2)\alpha}{(1 + \varepsilon/2)\alpha} \geq 1 - \varepsilon. \end{aligned}$$

It follows that the value of the strategy F' for the game Γ' is at least $1 - \varepsilon - e^{-\alpha c}$. □

LEMMA 4.4 (SOUNDNESS). *If the value of Γ' is at least $1 - \eta$, then the $1/3c$ -partial value of Γ is at least $1 - 12\eta$.*

PROOF. Let $F' : V^c \rightarrow [c] \times \Sigma$ be a strategy for Γ' of value $1 - \eta$. We first construct a partial strategy for Γ that uses shared randomness. For $i \in [c]$ and $\underline{u} \in V^{c-1}$, we define a partial strategy $F_{i,\underline{u}} : V \rightarrow \Sigma \cup \{\perp\}$ as

$$F_{i,\underline{u}}(u) := \begin{cases} a & \text{if } F'(\underline{u} +_i u) = (i, a), \\ \perp & \text{otherwise.} \end{cases}$$

Here, $\underline{u} +_i u$ denotes the vertex sequence $\underline{u}' \in V^c$ obtained from \underline{u} by inserting u as the i -th coordinate (the original coordinates $i, \dots, c-1$ of \underline{u} are moved by one to the right).

It is clear that

$$\mathbb{P}_{i \in [c], \underline{u} \in V^{c-1}, u \sim V} \left\{ F_{i,\underline{u}}(u) \in \Sigma \right\} = 1/c. \quad (4.1)$$

Since the value of F' is $1 - \eta$, we have

$$\begin{aligned} (1 - \eta)/c &= \mathbb{P}_{i \in [c], \underline{u}, v \in E_{\circ}^{c-1}, u, v \sim E_{\circ}} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\} \\ &= \frac{1}{2} \mathbb{P}_{i \in [c], \underline{u}, v \in E_{\circ}^{c-1}, u, v \in E} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\} \\ &\quad + \frac{1}{2} \mathbb{P}_{i \in [c], \underline{u}, v \in E_{\circ}^{c-1}, u \sim V} \left\{ F_{i,\underline{u}}(u) = F_{i,\underline{v}}(v) \in \Sigma \right\} \end{aligned}$$

From (4.1), it follows that both probabilities in the previous equation are at most $1/c$. Hence, for their average to be $(1 - \eta)/c$, both of them have to be at least $(1 - 2\eta)/c$,

$$\mathbb{P}_{i \in [c], \underline{u}, v \in E_{\circ}^{c-1}, u, v \in E} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\} \geq (1 - 2\eta)/c, \quad (4.2)$$

$$\mathbb{P}_{i \in [c], \underline{u}, v \in E_{\circ}^{c-1}, u \sim V} \left\{ F_{i,\underline{u}}(u) = F_{i,\underline{v}}(v) \in \Sigma \right\} \geq (1 - 2\eta)/c. \quad (4.3)$$

To further analyze the partial strategies, we define two random variables

$$\begin{aligned} \text{Vol}(i, \underline{u}, \underline{v}) &:= \mathbb{P}_{u \sim V} \left\{ F_{i,\underline{u}}(u) \in \Sigma \vee F_{i,\underline{v}}(u) \in \Sigma \right\}, \\ \text{Val}(i, \underline{u}, \underline{v}) &:= \mathbb{P}_{uv \sim E_{\circ}} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\}. \end{aligned}$$

The measure on $(i, \underline{u}, \underline{v})$ is as follows: We choose $i \in [c]$ uniformly at random, and sample $\underline{u}, \underline{v}$ from E_{\circ}^{c-1} . It is clear that $\text{Val} \leq \text{Vol}$ (pointwise) and

$$\begin{aligned} \mathbb{E} \text{Vol} &= \mathbb{E}_{i \in [c], \underline{u}, \underline{v} \sim E_{\circ}^{c-1}} \left[\mathbb{P}_{u \sim V} \left\{ F_{i,\underline{u}}(u) \in \Sigma \right\} \right. \\ &\quad \left. + \mathbb{P}_{u \sim V} \left\{ F_{i,\underline{v}}(u) \in \Sigma \right\} - \mathbb{P}_{u \sim V} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(u)) \in \Sigma \right\} \right] \\ &= 2/c - \mathbb{P}_{i \in [c], \underline{u}, \underline{v} \sim E_{\circ}^{c-1}, u \sim V} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(u)) \in \Sigma \right\} \\ &\leq (1 + 2\eta)/c \quad (\text{using (4.3)}). \end{aligned} \quad (4.4)$$

On the other hand, since the value of F' is $1 - \eta$, $\mathbb{E} \text{Val}$ is given by,

$$\mathbb{P}_{i \in [c], \underline{u}, \underline{v} \sim E_{\circ}^{c-1}, u, v \sim E_{\circ}} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\} = (1 - \eta)/c. \quad (4.6)$$

At this point, we could use [Lemma 2.2](#) (Glorified Markov Inequality) to argue that there exists a triple $(i, \underline{u}, \underline{v})$ such that $\text{Val}(i, \underline{u}, \underline{v}) \geq (1 - O(\eta))\text{Vol}(i, \underline{u}, \underline{v})$ and $\text{Vol}(i, \underline{u}, \underline{v}) \geq \mathbb{E} \text{Vol}/2 \geq 1/2c$.

For the reader's convenience, we repeat the (short) argument: Let $\eta' \leq 3\eta$ be such that $1 - \eta' = (1 - \eta)/(1 + 2\eta)$. Let Good be the event $\text{Val} \geq (1 - 2\eta')\text{Vol}$ and let Bad be the complementary event. We denote by $\mathbb{1}_{\text{Good}}$ and $\mathbb{1}_{\text{Bad}}$ the indicator variables of these events. We can relate $\mathbb{E} \mathbb{1}_{\text{Good}} \text{Vol}$ and $\mathbb{E} \text{Vol}$ as follows

$$\begin{aligned} 0 &\leq \mathbb{E} \text{Val} - (1 - \eta') \mathbb{E} \text{Vol} \quad (\text{using (4.5) and (4.6)}) \\ &\leq \mathbb{E} \mathbb{1}_{\text{Good}} \text{Vol} + (1 - 2\eta') \mathbb{E} \mathbb{1}_{\text{Bad}} \text{Vol} - (1 - \eta') \mathbb{E} \text{Vol} \\ &= 2\eta' \mathbb{E} \mathbb{1}_{\text{Good}} \text{Vol} - \eta' \mathbb{E} \text{Vol}. \end{aligned}$$

Hence, we can find i, \underline{u} , and \underline{v} such that $\mathbb{1}_{\text{Good}}(i, \underline{u}, \underline{v}) = 1$ and $\text{Vol}(i, \underline{u}, \underline{v}) \geq \mathbb{E} \text{Vol}/2$, i.e.,

$$\mathbb{P}_{uv \sim E_{\circ}} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\} \geq (1 - 2\eta') \cdot \mathbb{P}_{u \sim V} \left\{ F_{i,\underline{u}}(u) \in \Sigma \vee F_{i,\underline{v}}(u) \in \Sigma \right\}, \quad (4.7)$$

$$\mathbb{P}_{u \sim V} \left\{ F_{i,\underline{u}}(u) \in \Sigma \vee F_{i,\underline{v}}(u) \in \Sigma \right\} \geq (1 - 2\eta)/2c. \quad (4.8)$$

We claim that the two conditions (4.7) and (4.8) allows us to construct an $O(1/c)$ -partial strategy of value $1 - O(\eta)$. More concretely, we can modify $F_{i,\underline{u}}$ and $F_{i,\underline{v}}$ in such a way that $F_{i,\underline{u}}(u) = F_{i,\underline{v}}(u)$ for all $u \in V$ while maintaining the conditions (4.7) and (4.8). (Here, we use the fact that with probability $1/2$, we test $F_{i,\underline{u}}(u) = F_{i,\underline{v}}(u) \in \Sigma$.) In this way, we get an assignment $F^* : V \rightarrow \Sigma \cup \{\perp\}$ such that $F^*(u) = \perp$ if $F_{i,\underline{u}}(u) = F_{i,\underline{v}}(u) = \perp$ and $F^*(u) \in \{F_{i,\underline{u}}(u), F_{i,\underline{v}}(u)\} \cap \Sigma$ otherwise. Furthermore, the assignment F^* satisfies

$$\begin{aligned} \mathbb{P}_{u \sim V} \left\{ F^*(u) \in \Sigma \right\} &= \mathbb{P}_{u \sim V} \left\{ F_{i,\underline{u}}(u) \in \Sigma \vee F_{i,\underline{v}}(u) \in \Sigma \right\}, \\ \mathbb{P}_{uv \sim E_{\circ}} \left\{ (F^*(u), F^*(v)) \in \Pi_{u,v} \right\} &\geq \mathbb{P}_{uv \sim E_{\circ}} \left\{ (F_{i,\underline{u}}(u), F_{i,\underline{v}}(v)) \in \Pi_{u,v} \right\}. \end{aligned}$$

Let $\alpha = \mathbb{P}_{u \sim V} \{F^*(u) \in \Sigma\}$. We have

$$\begin{aligned} (1 - 2\eta') \cdot \alpha &\leq \mathbb{P}_{uv \sim E_{\circ}} \left\{ (F^*(u), F^*(v)) \in \Pi_{u,v} \right\} \\ &= \frac{1}{2} \mathbb{P}_{uv \in E} \left\{ (F^*(u), F^*(v)) \in \Pi_{u,v} \right\} + \frac{1}{2} \alpha \end{aligned}$$

It follows that $\mathbb{P}_{uv \in E} \{ (F^*(u), F^*(v)) \in \Pi_{u,v} \} \geq (1 - 4\eta') \cdot \alpha$. We can conclude that F^* is an $(1 - 2\eta)/2c$ -partial strategy of value at least $1 - 4\eta' \geq 1 - 12\eta$ for Γ . \square

5. WRAPPING UP

The main result of this work namely [Theorem 1.4](#) is an immediate consequence of the following theorem.

THEOREM 5.1. *Given a regular graph G of size n and parameters $\varepsilon, \delta > 0$ such that $\delta = m/n$ for some integer m , we can compute in polynomial time a UNIQUE GAMES instance Υ with $(n/\varepsilon)^{O(1/\delta)}$ variables and with $\text{poly}(1/\varepsilon\delta)$ labels such that the following holds with $C = 1500$:*

–**Completeness:** *For all $\eta > 0$ if*

$$\Phi_G(\delta) \leq \eta \implies \text{opt}(\Upsilon) \geq 1 - C(\eta + \varepsilon)$$

–**Soundness:**

$$\Phi_G(\delta) \geq 1 - \varepsilon^2/2C^2 \log^2(1/\varepsilon) \implies \text{opt}(\Upsilon) \leq 1 - \frac{1}{2C}$$

PROOF. Execute the SSE-TO-UG reduction with parameters $R = \lfloor \frac{1}{\delta} \rfloor$ and ε on the input graph $G = (V, E)$. This yields a UNIQUE GAMES instance Υ^* on alphabet size R . Apply the reduction from partial UNIQUE GAMES to UNIQUE GAMES presented in Section 4 with parameter $c = 40e \log(1/\varepsilon)$, to get the UNIQUE GAMES instance Υ .

Completeness If $\Phi_G(\delta) \leq \eta$ then there exists a subset S such that $\mu(S) \in [\delta/10, \delta]$ and $\Phi(S) \leq \varepsilon$. By choice of R , we have $\mu(S) \in [\frac{1}{20R}, \frac{1}{R}]$. Hence by Theorem 3.1, there exists a $1/20e$ -partial strategy with value $(1 - 40e\eta - 4\varepsilon)$ for Υ^* . Finally, by Theorem 4.1 there exists a strategy for the game Υ succeeding with probability at least $1 - 40e\eta - 4\varepsilon - \varepsilon^2$.

Soundness Let $\beta = \frac{1}{2C}$. Suppose not, let us say $\text{opt}(\Upsilon) \geq 1 - \beta$. By Theorem 4.1, if $\text{opt}(\Upsilon) \geq 1 - \beta$ then there exists a $\frac{1}{3c}$ -partial strategy for Υ^* with value at least $1 - 12\beta$. Applying Theorem 3.1, we get that there exists a set S with $\Phi(S) \leq 1152\beta$ and

$$\mu(S) \in \left[\frac{\delta}{480e \log(1/\varepsilon)}, \frac{6\delta}{\varepsilon} \right].$$

If $\mu(S) < \delta$ then pad it with arbitrary vertices to construct a set S' such that $\mu(S') = \delta$. On the other hand, if $\mu(S) > \delta$, randomly subsample a subset $S' \subseteq S$ of volume δ . In either case, it is easy to check that $\Phi(S') \leq 1 - (1 - 1152\beta)\varepsilon^2/C^2 \log^2(1/\varepsilon) \leq 1 - \varepsilon^2/2C^2 \log^2(1/\varepsilon)$. \square

Acknowledgments: We thank Boaz Barak for suggesting the notion of partial games, which helped to improve the presentation of the results. We also thank Sanjeev Arora, Venkatesan Guruswami, Oded Regev and Prasad Tetali for insightful discussions.

6. REFERENCES

- [1] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- [2] N. Alon and V. D. Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory. Series B*, 38:73–88, 1985.
- [3] S. Arora, R. Impagliazzo, W. Matthews, and D. Steurer. Improved algorithms for unique games via divide and conquer. In *Electronic Colloquium on Computational Complexity ECCCTR: TR10-041*, 2010.
- [4] S. Arora, S. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. K. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *STOC*, pages 21–28. ACM, 2008.
- [5] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proceedings of the thirty-sixth annual ACM Symposium on Theory of Computing (STOC-04)*, pages 222–231, June 13–15 2004.
- [6] Y. Aumann and Y. Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, Feb. 1998.
- [7] B. Barak, M. Hardt, I. Haviv, A. Rao, O. Regev, and D. Steurer. Rounding parallel repetitions of unique games. In *FOCS*, pages 374–383. IEEE Computer Society, 2008.
- [8] M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for sherali-adams relaxations. In *STOC*, pages 283–292. ACM, 2009.
- [9] J. Cheeger. A lower bound on smallest eigenvalue of a laplacian. *Problems in Analysis (Papers dedicated to Salomon Bochner)*, pages 195–199, 1970.
- [10] Feige and Schechtman. On the optimality of the random hyperplane rounding technique for MAX CUT. *RSA: Random Structures Algorithms*, 20, 2002.
- [11] U. Feige, G. Kindler, and R. O’Donnell. Understanding parallel repetition requires understanding foams. In *IEEE Conference on Computational Complexity*, pages 179–192, 2007.
- [12] K. Georgiou, A. Magen, T. Pitassi, and I. Tourlakis. Integrality gaps of 2 - o(1) for vertex cover SDPs in the lovász-schrijver hierarchy. In *FOCS*, pages 702–712. IEEE Computer Society, 2007.
- [13] S. Khot. On the power of unique 2-prover 1-round games. In *STOC*, pages 767–775. ACM, 2002.
- [14] 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.
- [15] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2-\varepsilon$. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.
- [16] 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. IEEE Computer Society, 2005.
- [17] F. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [18] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [19] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *STOC*, pages 245–254. ACM, 2008.
- [20] P. Raghavendra and D. Steurer. How to round any CSP. In *FOCS*, pages 586–594. IEEE Computer Society, 2009.
- [21] P. Raghavendra and D. Steurer. Integrality gaps for strong SDP relaxations of UNIQUE GAMES. In *FOCS*, pages 575–585. IEEE Computer Society, 2009.
- [22] P. Raghavendra, D. Steurer, and P. Tetali. Approximations for the isoperimetric and spectral profile of graphs and for restricted eigenvalues of diagonally-dominant matrices. In *STOC*. ACM, 2010. To Appear.
- [23] A. Rao. Parallel repetition in projection games and a concentration bound. In *STOC*, pages 1–10. ACM, 2008.
- [24] R. Raz. A counterexample to strong parallel repetition. In *FOCS*, pages 369–373. IEEE Computer Society, 2008.
- [25] G. Schoenebeck, L. Trevisan, and M. Tulsiani. A linear round lower bound for Lovász-Schrijver SDP relaxations of vertex cover. In *IEEE Conference on Computational Complexity*, pages 205–216, 2007.