2012 IEEE 27th Conference on Computational Complexity

Reductions Between Expansion Problems

Prasad Raghavendra
Georgia Institute of Technology,
Atlanta, USA.
Email: raghavendra@cc.gatech.edu

David Steurer
Microsoft Research New England,
Cambridge, USA.
Email: dsteurer@cs.princeton.edu

Madhur Tulsiani
Toyota Technological Institute,
Chicago, USA.
Email: madhurt@ttic.edu

***Abstract*—The *Small-Set Expansion Hypothesis* (Raghavendra, Steurer, STOC 2010) is a natural hardness assumption concerning the problem of approximating the edge expansion of small sets in graphs. This hardness assumption is closely connected to the *Unique Games Conjecture* (Khot, STOC 2002). In particular, the Small-Set Expansion Hypothesis implies the Unique Games Conjecture (Raghavendra, Steurer, STOC 2010).**

Our main result is that the Small-Set Expansion Hypothesis is in fact equivalent to a variant of the Unique Games Conjecture. More precisely, the hypothesis is equivalent to the Unique Games Conjecture restricted to instance with a fairly mild condition on the expansion of small sets. Alongside, we obtain the first strong hardness of approximation results for the Balanced Separator and Minimum Linear Arrangement problems. Before, no such hardness was known for these problems even assuming the Unique Games Conjecture.

These results not only establish the Small-Set Expansion Hypothesis as a natural unifying hypothesis that implies the Unique Games Conjecture, all its consequences and, in addition, hardness results for other problems like Balanced Separator and Minimum Linear Arrangement, but our results also show that the Small-Set Expansion Hypothesis problem lies at the combinatorial heart of the Unique Games Conjecture.

The key technical ingredient is a new way of exploiting the structure of the Unique Games instances obtained from the Small-Set Expansion Hypothesis via (Raghavendra, Steurer, 2010). This additional structure allows us to modify standard reductions in a way that essentially destroys their local-gadget nature. Using this modification, we can argue about the expansion in the graphs produced by the reduction without relying on expansion properties of the underlying Unique Games instance (which would be impossible for a local-gadget reduction).

***Index Terms*—Unique games conjecture, Small set expansion, Balanced separator, Minimum linear arrangement.**

I. Introduction

Finding small vertex or edge separators in a graph is a fundamental computational task. Even from a purely theoretical standpoint, the phenomenon of vertex and edge expansion – the lack of good vertex and edge separators, has had numerous implications in all branches of theoretical computer science. Yet, the computational complexity of detecting and approximating expansion, or finding good vertex and edge separators in graphs is not 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 (non-empty) subset $S \subseteq V$ is 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 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.,

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

These notions of conductance can be extended naturally to non-regular graphs, and finally to arbitrary weighted graphs (see Section II). Henceforth, in this section, for a subset of vertices S in a graph G we will use the notation $\mu(S)$ to denote the normalized set size, i.e., $\mu(S) = |S|/n$ in a n vertex graph.

The problem of approximating the quantity Φ_G for a graph 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 [1] shown by Alon-Milman [2] and Alon [3]. Specifically, Cheeger's inequality relates the conductance Φ_G to the second eigenvalue of the adjacency matrix of the graph – an efficiently computable quantity. This yields an approximation algorithm for Φ_G, one that is used heavily in practice for graph partitioning. However, the approximation for Φ_G obtained via Cheeger's inequality is

1093-0159/12 $26.00 © 2012 IEEE
DOI 10.1109/CCC.2012.43

poor in terms of a approximation ratio, especially when the value of Φ_G is small. An $O(\log n)$ approximation algorithm for Φ_G was obtained by Leighton and Rao [4]. Later work by Linial et al. [5] 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. [7] obtained an $O(\sqrt{\log n})$ approximation for the problem using semidefinite programming techniques.

Small Set Expansion. It is easy to see that 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 $n/1000$ vertices in a random d-regular graph have $\Phi(S) \geqslant 0.99$ with very high probability, while the conductance Φ_G of the entire graph is roughly $1/2$.

A more refined measure of the edge expansion of a graph is its expansion profile. Specifically, for a graph G the expansion profile is given by the curve

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

The problem of approximating the expansion profile has received much less attention, and is seemingly far less tractable. The second eigenvalue λ_2 fails to approximate 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 Φ_G is small (say ε), yet every sufficiently small set has near perfect expansion ($\Phi(S) \geqslant 1 - \varepsilon$). This implies that Φ_G (and the second eigenvalue λ_2) does not yield any information about expansion of small sets.

In a recent work, Raghavendra, Steurer, and Tetali [8] give a polynomial-time algorithm based on semidefinite programming for this problem. Roughly speaking, the approximation guarantee of their algorithm for $\Phi_G(\delta)$ is similar to the one given by Cheeger's inequality for Φ_G, except with the approximation degrading by a $\log 1/\delta$ factor. In particular, the approximation gets worse as the size of the sets considered gets smaller.

In the regime when $\Phi_G(\delta)$ tends to zero as a function of the instance size n, an $O(\log n)$-approximation follows from the framework of Räcke [9]. Very recently, this approximation has been improved to a $O(\sqrt{\log n \cdot \log(1/\delta)})$-approximation [10]. Our work focuses on the regime when $\Phi_G(\delta)$ is not a function of the instance size n. In this regime, the algorithm of [8] gives the best known approximation for the expansion profile $\Phi_g(\delta)$.

In summary, the current state-of-the-art algorithms for approximating the expansion profile of a graph are still far from satisfactory. Specifically, the following hypothesis is consistent with the known algorithms for approximating expansion profile.

Hypothesis (Small-Set Expansion Hypothesis, [11])**.** *For every constant $\eta > 0$, there exists sufficiently small $\delta > 0$ such that given a graph G it is* NP*-hard to distinguish the cases,*

YES: *there exists a vertex set S with volume $\mu(S) = \delta$ and expansion $\Phi(S) \leqslant \eta$,*

NO: *all vertex sets S with volume $\mu(S) = \delta$ have expansion $\Phi(S) \geqslant 1 - \eta$.*

For the sake of succinctness, we will refer to the above promise problem as SMALL-SET EXPANSION with parameters (η, δ). Apart from being a natural optimization problem, the SMALL-SET EXPANSION problem is closely tied to the Unique Games Conjecture, as discussed in the next paragraph.

Recently, Arora, Barak, and Steurer [12] showed that the problem SMALL-SET EXPANSION(η, δ) admits a subexponential algorithm, namely an algorithm that runs in time $\exp(n^{\eta}/\delta)$. However, such an algorithm does not refute the hypothesis that the problem SMALL-SET EXPANSION(η, δ) might be hard for every constant $\eta > 0$ and sufficiently small $\delta > 0$.

Unique Games Conjecture. Khot's Unique Games Conjecture [13] is among the central open problems in hardness of approximation. At the outset, the conjecture asserts that a certain constraint satisfaction problem called the Unique Games is hard to approximate in a strong sense.

An instance of UNIQUE GAMES consists of a graph with vertex set V, a finite set of labels $[R]$, and a permutation $\pi_{v\leftarrow w}$ of the label set for each edge (v, w) of the graph. A labeling $F : V \to [R]$ of the vertices of the graph is said to *satisfy* an edge (v, w), if $\pi_{v\leftarrow w}(F(w)) = F(v)$. The objective is to find a labeling that satisfies the maximum number of edges.

The Unique Games Conjecture asserts that if the label set is large enough then even though the input instance has a labeling satisfying almost all the edges, it is NP-hard to find a labeling satisfying any non-negligible fraction of edges.

In recent years, Unique Games Conjecture has been shown to imply optimal inapproximability results for classic problems like MAX CUT [14], VERTEX COVER [15] SPARSEST CUT [16] and all constraint satisfaction problems [17]. Unfortunately, it is not known if the converse of any of these implications holds. In other words, there are no known polynomial-time *reductions* from these classic optimization problems to UNIQUE GAMES, leaving the possibility that while the its implications are true the conjecture itself could be false.

Recent work by two of the authors established a *reverse* reduction from the SMALL-SET EXPANSION problem to Unique Games [11]. More precisely, their work showed that Small-Set Expansion Hypothesis implies the Unique Games Conjecture. This result suggests that the problem of approximating expansion of small sets lies at the combinatorial heart of the Unique Games problem. In fact, this connection proved useful in the development of subexponential time algorithms for Unique Games by Arora, Barak and Steurer [12]. It was also conjectured in [11] that Unique Games Conjecture is equivalent to the Small-Set Expansion Hypothesis.

A. Results (Informal Description)

In this work, we further investigate the connection between SMALL-SET EXPANSION and the Unique Games problem. The main result of this work is that the Small-Set Expansion Hypothesis is equivalent to a variant of the Unique Games Conjecture. More precisely, we show the following:

Theorem (Main Theorem, Informal)**.** *The Small-Set Expansion Hypothesis is equivalent to assuming that the Unique Games Conjecture holds even when the input instances are required to be small set expanders, i.e., sets of roughly δn vertices for some small constant δ have expansion close to* 1.

As a corollary, we show that Small-Set Expansion Hypothesis implies hardness of approximation results for BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT problems. The significance of these results stems from two main reasons.

First, the Unique Games Conjecture is not known to imply hardness results for problems closely tied to graph expansion such as BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT. The reason being that the hard instances of these problems are required to have certain global structure namely expansion. Gadget reductions from a unique games instance preserve the global properties of the unique games instance such as lack of expansion. Therefore, showing hardness for BALANCED SEPARATOR or MINIMUM LINEAR ARRANGEMENT problems often required a stronger version of the Unique Games Conjecture, where the instance is guaranteed to have good expansion. To this end, several such variants of the conjecture for expanding graphs have been defined in literature, some of which turned out to be false [18]. Our main result shows that the Small-Set Expansion Hypothesis serves as a natural unified assumption that yields all the implications of Unique Games Conjecture and, in addition, also hardness results for other fundamental problems such as BALANCED SEPARATOR.

Second, several results in literature point to the close connection between SMALL-SET EXPANSION problem and the Unique Games problem. One of the central implications of the Unique Games Conjecture is that certain semidefinite programs yield optimal approximation for various classes of problems. As it turns out, hard instances for semidefinite programs (SDP integrality gaps) for MAX CUT [19], [16], [20], [21], VERTEX COVER [22], UNIQUE GAMES [16], [21] and (uniform) SPARSEST CUT [23], [21] 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. All known integrality gap instances for semidefinite programming relaxations of Unique Games, can be translated in to gap instances for SMALL-SET EXPANSION problem, and are arguably more natural in the latter context. Furthermore, all the algorithmic results for SMALL-SET EXPANSION, including the latest work of Arora, Barak and Steurer [12] extend to Unique Games as well. This apparent connection was formalized in the result of Raghavendra et al. [11] which showed that Small-Set Expansion Hypothesis implies the Unique Games Conjecture. This work complements that of Raghavendra et al. [11] in exhibiting that the SMALL-SET EXPANSION problem lies at the combinatorial heart of the Unique Games problem.

We also show a "hardness amplification" result for SMALL-SET EXPANSION proving that if the Small-Set Expansion Hypothesis holds then the current best algorithm for SMALL-SET EXPANSION due to [8] is optimal within some fixed constant factor. One can view the reduction as a "scale change" operation for expansion problems, which starting from the qualitative hardness of a problem about expansion of sets with a sufficiently small measure δ, gives the optimal quantitative hardness results for problems about expansion of sets with any desired measure (larger than δ). This is analogous to (and based on) the results of [14] who gave a similar alphabet reduction for UNIQUE GAMES. An interesting feature of the reductions in the paper is that they produce instances whose expansion of small sets closely mimics a certain graph on the Gaussian space.

II. PRELIMINARIES

Unique Games. An instance of UNIQUE GAMES represented as $\mathcal{U} = (\mathcal{V}, \mathcal{E}, \Pi, [R])$ 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 $[R] = \{1, \ldots, R\}$, and a set of permutations $\Pi = \{\pi_{v \leftarrow w} : [R] \to [R]\}$, one permutation for each edge $e = (w, v) \in \mathcal{E}$. An assignment $F \colon \mathcal{V} \to [R]$ of labels to vertices is said to satisfy an edge $e = (w, v)$, if $\pi_{v \leftarrow w}(F) = F(v)$. The objective is to find an assignment F 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 II.1 (Unique Games $(R, 1 - \varepsilon, \eta)$)**.** Given a Unique Games instance $\mathcal{U} = (\mathcal{V}, \mathcal{E}, \Pi = \{\pi_{v \leftarrow w} : [R] \to [R] \mid e = (w, v) \in \mathcal{E}\}, [R])$ with number of labels R, distinguish between the following two cases:

- $(1 - \varepsilon)$- satisfiable instances: There exists an assignment F 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 II.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.*

Small-Set Expansion Hypothesis

Problem II.3 (Small-Set Expansion (η, δ))**.** Given a regular graph $G = (V, E)$, distinguish between the following two cases:

- Yes: There exists a non-expanding set $S \subseteq V$ with $\mu(S) = \delta$ and $\Phi_G(S) \leqslant \eta$.
- No: All sets $S \subseteq V$ with $\mu(S) = \delta$ are highly expanding having $\Phi_G(S) \geqslant 1 - \eta$.

Hypothesis II.4 (Hardness of approximating Small-Set Expansion)**.** *For all $\eta > 0$, there exists $\delta > 0$ such that the promise problem* Small-Set Expansion *(η, δ) is NP-hard.*

Definition II.5. Let $\mathcal{P}$ be a decision problem of distinguishing between two disjoint families (cases) of instances denoted by {Yes, No}. For a given instance $\mathcal{I}$ of $\mathcal{P}$, let $\mathsf{Case}(\mathcal{I})$ denote the family to which $\mathcal{I}$ belongs. We say that $\mathcal{P}$ is SSE-hard if for some $\eta > 0$ and all $\delta \in (0, \eta)$, there is a polynomial time reduction, which starting from an instance $G = (V, E)$ of Small-Set Expansion(η, δ), produces an instance $\mathcal{I}$ of $\mathcal{P}$ such that

- $\exists S \subseteq V$ with $\mu(S) = \delta$ and $\Phi_G(S) \leqslant \eta \implies \mathsf{Case}(\mathcal{I}) = \text{Yes}$.
- $\forall S \subseteq V$ with $\mu(S) = \delta$, $\Phi_G(S) \geqslant 1 - \eta \implies \mathsf{Case}(\mathcal{I}) = \text{No}$.

For the proofs, it shall be more convenient to use the following version of the Small-Set Expansion problem, in which high expansion is guaranteed not only for sets of measure δ, but also within an arbitrary multiplicative factor of δ.

Problem II.6 (Small-Set Expansion (η, δ, M))**.** Given a regular graph $G = (V, E)$, distinguish between the following two cases:

- Yes: There exists a non-expanding set $S \subseteq V$ with $\mu(S) = \delta$ and $\Phi_G(S) \leqslant \eta$.
- No: All sets $S \subseteq V$ with $\mu(S) \in \left(\frac{\delta}{M}, M\delta\right)$ have $\Phi_G(S) \geqslant 1 - \eta$.

The following proposition shows that for the purposes of showing that $\mathcal{P}$ is SSE-hard, it is sufficient to give a reduction from Small-Set Expansion (η, δ, M) for any chosen values of η, M and for all δ. We defer the proof to the full version.

Proposition II.7. *For all $\eta > 0, M \geqslant 1$ and all $\delta < 1/M$, there is polynomial time reduction from* Small-Set Expansion *$(\frac{\eta}{M}, \delta)$ to* Small-Set Expansion*(η, δ, M).*

III. Technical Preliminaries

Random walks on graphs. Consider the natural random walk on V defined by G. We write $j \sim G(i)$ to denote a random neighbor of vertex i in G (one step of the random walk started in i). The stationary measure for the random walk is given by the volume as defined earlier with $\mu(i) = G(\{i\}, V)$, where $G(\{i, \}, V)$ denotes the fraction of all the edges in the graph that are incident on i. If G is regular, then μ is the uniform distribution on V. In general, μ is proportional to the degrees of the vertices in G. We write $i \sim \mu$ to denote a vertex sampled according to the stationary measure. If G is clear from the context, we often write $i \sim V$ instead of $i \sim \mu$.

Spectral gap of graphs. We identify G with the stochastic matrix of the random walk on G. We equip the vector space $\{f \colon V \to \mathbb{R}\}$ with the inner product

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

We define $\|f\| = \langle f, f \rangle^{1/2}$. As usual, we refer to this (Hilbert) space as $L_2(V)$. Notice that G is self-adjoint with respect to this inner product, i.e., $\langle f, Gg \rangle = \langle Gf, g \rangle$ for all $f, g \in L_2(V)$. Let $\lambda_1 \geqslant \ldots \geqslant \lambda_n$ be the eigenvalues of G. The non-zero constants are eigenfunctions of G with eigenvalue $\lambda_1 = 1$.

For a vertex set $S \subseteq V$, let $\mathbb{1}_S$ be the $\{0, 1\}$-indicator function of S. We denote by $G(S, T) = \langle \mathbb{1}_S, G\mathbb{1}_T \rangle$ the total weight of all the edges in G that go between S and T.

Fact III.1. *Suppose the second largest eigenvalue of G is λ. Then, for every function $f \in L_2(V)$,*

$$\langle f, Gf \rangle \leqslant (\mathbb{E} f)^2 + \lambda \cdot \left(\|f\|^2 - (\mathbb{E} f)^2 \right) .$$

In particular, $\Phi_G(\delta) \geqslant 1 - \delta - \lambda$ for every $\delta > 0$.

Gaussian Graphs. For a constant $\rho \in (-1, 1)$, let $\mathcal{G}(\rho)$ denote the infinite graph over $\mathbb{R}$ where the weight of an edge (x, y) is the probability density that two standard Gaussian random variables X, Y with correlation ρ equal x and y respectively. The expansion profile of Gaussian graphs is given by $\Phi_{\mathcal{G}(\rho)}(\mu) = 1 - \Gamma_\rho(\mu)/\mu$ where the quantity $\Gamma_\rho(\mu)$ defined as

$$\Gamma_\rho(\mu) := \mathop{\mathbb{P}}_{(x,y) \sim \mathcal{G}_\rho} \{x \geqslant t, y \geqslant t\} ,$$

where $\mathcal{G}_\rho$ is the 2-dimensional Gaussian distribution with covariance matrix

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

and $t \geqslant 0$ is such that $\mathbb{P}_{(x,y) \sim \mathcal{G}_\rho} \{x \geqslant t\} = \mu$. A theorem of Borell [24] shows that for any set S of measure μ, $(\mathcal{G}(\rho))(S, S) \leqslant \Gamma_\rho(\mu)$. This expansion profile will be frequently used in the paper to state the results succinctly.

Noise graphs. For a finite probability space (Ω, ν) and $\rho \in [0, 1]$, we define $T = T_{\rho,\Omega}$ to be the following linear operator on $L_2(\Omega)$,

$$Tf(x) = \rho f(x) + (1 - \rho) \mathop{\mathbb{E}}_{x' \sim \nu} f(x') .$$

The eigenvalues of T are 1 (with multiplicity 1) and ρ (with multiplicity $|\Omega| - 1$). The operator T corresponds to the following natural (reversible) random walk on Ω: with probability ρ stay at the current position, with probability $(1 - \rho)$ move to a random position sampled according to the measure ν.

Product graphs. If G and G' are two graphs with vertex sets V and V', we let $H = G \otimes G'$ be the *tensor product* of G and G'. The vertex set of G is $V \times V'$. For $i \in V$ and $i' \in V'$, the distribution $H(i, i')$ is the product of the distributions $G(i)$ and $G'(i')$. For $R \in \mathbb{N}$, we let $G^{\otimes R}$ denote the *R-fold tensor product* of G. Sometimes the power R of the graph is clear from the context. In this case, we might drop the superscript for the tensor graph.

IV. Results

Towards stating the results succinctly, we introduce the notion of a decision problem being SSE-hard. It is the natural notion wherein a decision problem is SSE-hard if the Small-Set Expansion (η, δ) reduces to it by a polynomial time reduction for some constant η and all $\delta > 0$ (See Definition II.5).

A. Relation to the Unique Games Conjecture

We show that the Small-Set Expansion Hypothesis is equivalent to a certain variant of the Unique Games Conjecture with expansion. Specifically, consider the following version of the conjecture with near-perfect expansion of sufficiently small sets. The hypothesis is as follows: [1]

Hypothesis IV.1 (Unique Games with Small-Set Expansion)**.** *For every $\varepsilon, \eta > 0$ and $M \in \mathbb{N}$, there exists $\delta = \delta(\varepsilon, M) > 0$ and $q = q(\varepsilon, \eta, M) \in \mathbb{N}$ such that it is NP-hard to distinguish for a given* Unique Games *instance $\mathcal{U}$ with alphabet size q whether*

- Yes: *The* Unique Games *instance $\mathcal{U}$ is almost satisfiable,* $\mathrm{opt}(\mathcal{U}) > 1 - \varepsilon$.
- No: *The* Unique Games *instance $\mathcal{U}$ satisfies* $\mathrm{opt}(\mathcal{U}) < \eta$ *and its constraint graph G satisfies $\Phi(S) > 1 - \varepsilon$ for every vertex set with $\delta \leqslant \mu(S) \leqslant M\delta$.*

The main result of the paper is the following reduction from Small-Set Expansion to Unique Games on instances with small-set expansion.

Theorem IV.2. *For every $q \in \mathbb{N}$ and every $\varepsilon, \gamma > 0$, it is SSE-hard to distinguish between the following cases for a given* Unique Games *instance $\mathcal{U}$ with alphabet size q:*

- Yes: *The* Unique Games *instance $\mathcal{U}$ is almost satisfiable,* $\mathrm{opt}(\mathcal{U}) > 1 - 2\varepsilon - o(\varepsilon)$
- No: *The optimum of the* Unique Games *instance $\mathcal{U}$ is negligible, and the expansion profile of the instance resembles the Gaussian graph $\mathcal{G}(1 - \varepsilon)$. More precisely, the* Unique Games *instance $\mathcal{U}$ satisfies* $\mathrm{opt}(\mathcal{U}) < O\left(q^{-\varepsilon/(2-\varepsilon)}\right) + \gamma$ *and in addition, the constraint graph G of $\mathcal{U}$ satisfies*

$$\forall S \subseteq V(G). \qquad \Phi_G(S) \geqslant \Phi_{\mathcal{G}(1-\varepsilon)}(\mu(S)) - \gamma/\mu(S) .$$

The proof of the above theorem is deferred to the full version [25]. Together with Theorem 1.9 from [11], Theorem IV.2 implies the following equivalence:

Corollary IV.3. *The Small-Set Expansion Hypothesis is equivalent to Hypothesis IV.1 (Unique Games with Small-Set Expansion).*

Remark IV.4. If we choose $\gamma \ll \varepsilon$, then the constraint graph G in the No case satisfies $\Phi(S) \geqslant \Omega(\sqrt{\varepsilon})$ for every vertex set S with $\mu(S) \in (b, 1/2)$ for an arbitrarily small constant $b > 0$. In other words, the best balanced separator in G has cost $\Omega(\sqrt{\varepsilon})$. A hardness of Unique Games on

[1] The hypothesis in [11] is not quite the same. However, the reduction and its analysis in [11] also work for this hypothesis.

graphs of this nature was previously conjectured in [18], towards obtaining a hardness for BALANCED SEPARATOR.

As already mentioned, for several problems such as MAX CUT, the the hard instances for the semidefinite programs have high expansion of small sets. For instance, hard instances for semidefinite programs (SDP integrality gaps) for MAX CUT [19], [16], [20], [21], VERTEX COVER [22], UNIQUE GAMES [16], [21] and SPARSEST CUT [16], [20], [21] all have near-perfect edge expansion for small sets. In fact, in many of the cases, the edge expansion in the graph closely mimics the expansion of sets in some corresponding Gaussian graph. Confirming this observation, our techniques imply an optimal hardness result for MAX CUT on instances that are small-set expanders. More precisely, the Small-Set Expansion Hypothesis implies that the Goemans-Williamson algorithm is optimal even on graphs that are guaranteed to have good expansion of small sets, in fact an expansion profile that resembles the Gaussian graph. For the sake of succinctness, we omit the formal statement of the result.

B. Hardness Amplification for Graph Expansion

Observe that the Small-Set Expansion Hypothesis is a purely qualitative assumption on the approximability of expansion. Specifically, for every constant η the hypothesis asserts that there exists some δ such that approximating expansion of sets of size δ is NP-hard. The hypothesis does not assert any quantitative dependence on the set size and approximability. Surprisingly, we show that this qualitative hardness assumption is sufficient to imply precise quantitative bounds on approximability of graph expansion.

Theorem IV.5. *For all $q \in \mathbb{N}$ and $\varepsilon, \gamma > 0$, it is SSE-hard to distinguish between the following two cases for a given graph $H = (V_H, E_H)$*

YES: *There exist q disjoint sets $S_1, \ldots, S_q \subseteq V_H$ satisfying for all $l \in [q]$,*

$$\mu(S_l) = \tfrac{1}{q} \qquad \text{and} \qquad \Phi_H(S_l) \leqslant \varepsilon + o(\varepsilon).$$

NO: *For all sets $S \subseteq V_H$,*

$$\Phi_H(S) \;\geqslant\; \Phi_{\mathcal{G}(1-\varepsilon/2)}(\mu(S)) - \gamma/\mu(S)$$

where $\Phi_{\mathcal{G}(1-\varepsilon/2)}(\mu(S))$ is the expansion of sets of volume $\mu(S)$ in the infinite Gaussian graph $\mathcal{G}(1-\varepsilon/2)$.

The above hardness result matches (up to an absolute constant factor), the recent algorithmic result (Theorem 1.2) of [8] approximating the graph expansion. Furthermore, both the YES and the NO cases of the above theorem are even qualitatively stronger than in the Small-Set Expansion Hypothesis. In the YES case, not only does the graph have one non-expanding set, but it can be partitioned into small sets, *all* of which are non-expanding. This partition property is useful in some applications such as hardness reduction to MINIMUM LINEAR ARRANGEMENT. In the NO case, the expansion of all sets can be characterized only by their size $\mu(S)$. Specifically, the expansion of every set S of vertices with $\mu(S) \gg \gamma$, is at least the expansion of a set of similar size in the Gaussian graph $\mathcal{G}(1-\varepsilon/2)$.

Here we wish to draw an analogy to the Unique Games Conjecture. The Unique Games Conjecture is qualitative in that it does not prescribe a relation between its soundness and alphabet size. However, the work of Khot et al. [14] showed that the Unique Games Conjecture implies a quantitative form of itself with a precise relation between the alphabet size and soundness. Theorem IV.5 could be thought of as an analogue of this phenomena for the Small-Set Expansion problem.

As an immediate consequence of Theorem IV.5, we obtain the following hardness of the BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT problems (details deferred to full version [25]). Subsequent work [26] builds on Theorem IV.5 to show that achieving a constant-factor approximation for treewidth is SSE-hard.

Corollary IV.6 (Hardness of BALANCED SEPARATOR and MIN BISECTION)**.** *There is a constant c such that for arbitrarily small $\varepsilon > 0$, it is SSE-hard to distinguish the following two cases for a given graph $G = (V, E)$:*

YES: *There exists a cut $(S, V \setminus S)$ in G such that $\mu(S) = \frac{1}{2}$ and $\Phi_G(S) \leqslant \varepsilon + o(\varepsilon)$.*

NO: *Every cut $(S, V \setminus S)$ in G, with $\mu(S) \in \left(\frac{1}{10}, \frac{1}{2}\right)$ satisfies $\Phi_G(S) \geqslant c\sqrt{\varepsilon}$.*

Corollary IV.7 (Hardness of MINIMUM LINEAR ARRANGEMENT)**.** *It is SSE-hard to approximate* MINIMUM LINEAR ARRANGEMENT *to any fixed constant factor.*

V. WARM-UP: HARDNESS FOR BALANCED SEPARATOR

In this section we present a simplified version of our reduction from SMALL-SET EXPANSION to BALANCED SEPARATOR. Though it gives sub-optimal parameters, it illustrates the key ideas used in the general reduction.

A. Candidate Reduction from Unique Games

A natural approach for reducing UNIQUE GAMES to BALANCED SEPARATOR is to consider variants of the reduction from UNIQUE GAMES to MAX CUT in [14] (similarly, one could consider variants of the reduction from UNIQUE GAMES to the *generalized* SPARSEST CUT problem [16]).

Let $\mathcal{U}$ be a unique game with alphabet size R and vertex set V. (We assume that every vertex of the unique game participates in the same number of constraints. This assumption is without loss of generality.) The candidate reduction has a parameter $\varepsilon > 0$. The graph $H = H_\varepsilon(\mathcal{U})$ obtained from this candidate reduction has vertex set $V \times \{0,1\}^R$ and its edge distribution is defined as follows:

1) Sample a random vertex $u \in V$.
2) Sample two random constraints $(u, v, \pi), (u, v', \pi')$ of $\mathcal{U}$ that contain the vertex u. (Henceforth, we will write $(u, v, \pi) \sim \mathcal{U} \mid u$ to denote a random constraint of $\mathcal{U}$ containing vertex u.)
3) Sample a random edge (y, y') of the boolean noise graph $T_{1-\varepsilon}$ with noise parameter ε.
4) Output an edge between $(v, \pi(y))$ and $(v', \pi'(y'))$. (Here, $\pi(y)$ denotes the vector obtained by permuting the coordinates of y according to the permutation π.)

Completeness. Suppose there is a good assignment $F\colon V \to [R]$ for the unique game $\mathcal{U}$. Then, if we sample a random vertex $u \in V$ and two random constraint $(u, v, \pi), (u, v', \pi') \sim \mathcal{U} \mid u$, with probability very close to 1 (much closer than ε), the labels assigned to v and v' satisfy $\pi^{-1}(F(v)) = (\pi')^{-1}(F(v))$. Consider the vertex set $S = \{(u, x) \mid x_{F(u)} = 1\}$. in the graph H. We have $\mu(S) = 1/2$. We claim that the expansion of this set is essentially $\varepsilon/2$ (up to a lower-order term depending on the fraction of constraint of $\mathcal{U}$ violated by F). Consider a random edge e with endpoints $(v, \pi(y))$ and $(v', \pi'(y'))$, where the vertices $v, v' \in V$ and the permutations π, π' are generated as specified above. Let $r = \pi^{-1}(F(v))$ and $r' = (\pi')^{-1}(F(v))$. The edge e crosses the cut S if and only if $y_r \neq y'_{r'}$. As argued before, with probability very close to 1, we have $r = r'$. Conditioned on this event, the probability that $y_r \neq y_{r'}$ is equal to $\varepsilon/2$. This shows that S has expansion $\varepsilon/2$.

Soundness. Suppose no assignment for the unique game $\mathcal{U}$ satisfies a significant fraction of constraints. Let S be a vertex set in the graph H. The goal is to lower bound the expansion of S (which is the same as upper bounding the fraction of edges with both endpoints in S). Let $f\colon V \times \{0,1\}^R \to \{0,1\}$ be the indicator function of S. Following the analysis of [14], we consider functions $g_u\colon \{0,1\}^R \to [0,1]$ for every vertex $u \in V$,

$$g_u(x) = \mathop{\mathbb{E}}_{(u,v,\pi)\sim\mathcal{U}|u,} \left[f(v, \pi(x)) \right] .$$

By construction, the fraction $H(S, S)$ of edges of H with both endpoints in S is exactly

$$H(S, S) = \mathop{\mathbb{E}}_{u\in V} \langle g_u, T_{1-\varepsilon} g_u \rangle .$$

Let μ_u be the expected value of g_u and $\Gamma_{1-\varepsilon}(\cdot)$ is the noise stability profile of the Gaussian noise graph with parameter ε. Since $\mathcal{U}$ does not have a good assignment, standard arguments (invariance principle and influence decoding, see [14]) imply the following upper bound on $H(S, S)$,

$$H(S, S) \leqslant \mathop{\mathbb{E}}_{u\in V} \Gamma_{1-\varepsilon}(\mu_u) + o(1) .$$

(The notation $o(1)$ hides a term depending on the maximum fraction of constraints of $\mathcal{U}$ that can be satisfied. For us, this term is not significant.) We would like to show that every set S that contains a μ fraction of the vertices of H satisfies $H(S, S) \leqslant \Gamma_{1-\varepsilon}(\mu) + o(1)$. However, the function $\Gamma_{1-\varepsilon}$ is, of course, not concave. Hence, this upper bound holds only if μ_u is close to μ for most vertices $u \in V$.

In fact, it is easy to construct examples that show that the candidate reduction is not sound. For example, consider a unique game $\mathcal{U}$ that consists of two disjoint parts of the same size (i.e., without any constraint between the two parts). The reduction preserves this global structure, in the sense that the graph H also consists of two disjoint parts of the same size (with no edge between the parts). Hence, this graph contains a vertex set with volume 1/2 and expansion 0 irrespective of the optimal value of the unique game $\mathcal{U}$. In fact, any cut in the underlying graph of $\mathcal{U}$ can be translated to a cut in H and the resulting function f may have the values μ_u as (very close to) 0 or 1.

This example shows that the above candidate reduction can only work if one makes assumptions about structure of the constraint graph of the underlying unique game $\mathcal{U}$. However, such an assumption raises the question if Unique Games could be hard to approximate even if the constraint graph is expanding. This issue turns out to be delicate as demonstrated by the algorithm for Unique Games with expanding constraint graphs [18]. This algorithm achieves a good approximation for Unique Games if the expansion of the constraint graph exceeds a certain threshold.

B. Structured Unique Games from Small-Set Expansion

In this work, we present a different approach for fixing the above candidate reduction. Instead of assuming expansion properties of the constraint graph, we assume that the underlying unique game is obtained by the reduction from Small-Set Expansion to Unique Games in [11] [2]. This specific form of the underlying unique game will allows us to modify the reduction such that the global structure of the constraint graph is no longer preserved

[2] We remark that unique games of this form do not necessarily have expanding constraint graphs. In fact, it is still possible that the constraint graph consists of two disconnected components.

in the graph obtained from the reduction. (In particular, our modified reduction will break with the paradigm of composing unique games with local gadgets.)

In the following, we describe the reduction from Small-Set Expansion to Unique Games. Let G be a regular graph with vertex set V. For technical reasons, we assume that G contains a copy of the complete graph of weight $\eta > 0$. (Since we will be able to work with very small η, this assumption is without loss of generality.) Given a parameter $R \in \mathbb{N}$ and the graph G, the reduction outputs a unique game $\mathcal{U} = \mathcal{U}_R(G)$ with vertex set V^R and alphabet $[R]$. The constraints of the unique game $\mathcal{U}$ correspond to the following probabilistic verifier for an assignment $F\colon V^R \to [R]$:

1) Sample a random vertex $A \in V^R$.
2) Sample two random neighbors $B, C \sim G^{\otimes R}(A)$ of the vertex A in the tensor-product graph $G^{\otimes R}$.
3) Sample two random permutations π_B, π_C of $[R]$.
4) Verify that $\pi_B^{-1}(F(\pi_B(B))) = (\pi_C)^{-1}(F(\pi_C(C)))$.

Raghavendra and Steurer [11] show that this reduction is complete and sound in the following sense:

Completeness: If the graph G contains a vertex set with volume $1/R$ and expansion close to 0, then the unique game $\mathcal{U} = \mathcal{U}_R(G)$ has a partial assignment that labels an $\alpha \geqslant 1/e$ fraction of the vertices and satisfies almost an α fraction of the constraints.

Soundness: If the graph G contains no set with volume $1/R$ and expansion bounded away from 1, then no assignment for the unique game $\mathcal{U} = \mathcal{U}_R(G)$ satisfies a significant fraction of the constraints.

Hence, if one assumes the Small-Set Expansion Hypothesis, then the unique games instances obtained via the above reduction are hard to approximate.

We remark that the completeness of the reduction seems weaker than usual, because we are only guaranteed a partial assignment for the unique game. However, it is easy to check that the KKMO [14] reduction presented in the previous section also works if there is only a partial assignment in the completeness case. The only difference is that one now gets a set S with $\mu(S) = \alpha/2$ and expansion roughly $\varepsilon/2$.

C. Reduction from Small-Set Expansion to Balanced Separator

We now show how the combination of the above two reductions can be modified to give a reduction from Small-Set Expansion to Balanced Separator. Let $\mathcal{U} = \mathcal{U}_R(G)$ be the unique game given by the reduction of Raghavendra and Steurer. If we consider the graph H given by the reduction in Section V-A, each vertex of H is now of the form (A, x), where $A \in V^R$ and $x \in \{0, 1\}^R$.

The intuition is that in this case, we can think of x as picking a *subset* of the vertices in A, and that just the knowledge of this subset (instead of the whole of A) is sufficient for the provers to provide a good answer to the corresponding unique game. In particular, let $A' = \{A_i \mid x_i = 1\}$ be the subset picked by x. Then the argument for the completeness case in [11] actually shows that one can still find a good labeling for an α fraction of the vertices A, where the label of A only depends on A' [3].

Formally, if we replace A with the tuple $A'(x)$ defined by taking $A'_i = A_i$ if $x_i = 1$ and $A'_i = \perp$ otherwise. This gives a graph H' with the vertex set being a subset of $(V \cup \{\perp\})^R \times \{0, 1\}^R$. The the argument in completeness case for showing that H has a balanced cut of expansion roughly $\varepsilon/2$ can in fact be extended to show that H' also has a balanced cut of expansion roughly $\varepsilon/2$.

The soundness analysis in the previous reduction did not always work because H had the same structure as $G^{\otimes R}$, since we essentially replaced every vertex of $G^{\otimes R}$ by a gadget $\{0, 1\}^R$ to obtain H. However, the structure of H' is very different from that of $G^{\otimes R}$.

For example, consider the vertices $A = (u_1, u_2, \ldots, u_R)$ and $B = (v_1, u_2, \ldots u_R)$ in V^R which only differ in the first coordinate (A, B are not necessarily adjacent). Let $x \in \{0, 1\}^R$ be such that $x_1 = 0$. Then, while (A, x) and (B, x) are different vertices in H, $(A'(x), x)$ and $(B'(x), x)$ are in fact the same vertex in H'! On the other hand, if $x_1 = 1$, then $(A'(x), x)$ and $(B'(x), x)$ would be two different vertices in H'. Hence, the gadget structure of H is no longer preserved in H' - it is very different from a "locally modified" copy of $G^{\otimes R}$.

Remark V.1. The graph H' produced by the above reduction is similar to the graph produced as follows: Starting from an instance G of Small-Set Expansion, perform the reduction of [11] to Unique Games (using parameter $R/2$) to get $\mathcal{U}_{R/2}(G)$, and output $\mathcal{U}_{R/2}(G)$.

Note that every vertex in H' consists of a random tuple with roughly $R/2$ vertices, while the same holds for $\mathcal{U}_{R/2}(G)$. It is easy to see that for an appropriate choice of the noise parameter η in the [11] reduction, the resulting graph closely resembles the graph H'.

While this would be a much simpler description of the reduction, we were unable to analyze the reduction without going through the long code reduction outlined above. Looking at the instance H' as a result of a modification to the instance H produced by a long code reduction, is critical for the soundness analysis that we present.

[3]Given a non-expanding small set S, if $A' \cap S$ contains a single element A'_j, then we assign the label j to A. If $A' \cap S$ is not a singleton, we do not label A.

For the purposes of analysis, it will be more convenient to think of A' being obtained by replacing A_i where $x_i = 0$, by a random vertex of G instead of the symbol $\perp$. Instead of identifying different vertices in H with the same vertex in H', this now has the effect of re-distributing the weight of an incident on (A, x), uniformly over all the vertices that (A', x) can map to. Let M_x denote a Markov operator which maps A to a random A' as above.

We now state the combined reduction. The weight of an edge in the final graph H' is the probability that it is produced by the following process:

1) Sample a random vertex $A \in V^R$.
2) Sample two random neighbors $B, C \sim G^{\otimes R}(A)$ of the vertex A in the tensor-product graph $G^{\otimes R}$.
3) Sample $x_B, x_C \sim \{0,1\}^R$.
4) Sample $B' \sim M_{x_B}(B)$ and $C' \sim M_{x_C}(C)$.
5) Sample two random permutations π_B, π_C of $[R]$.
6) Output an edge between the vertices $\pi_B(B', x_B)$ and $\pi_C(C', x_C)$ ($\pi(A, x)$ denotes the tuple $(\pi(A), \pi(x))$).

As before, let $f : V^R \times \{0,1\}^R$ denote the indicator function of a set in H', with (say) $\mathbb{E} f = \mu = 1/2$. We define the functions

$$\bar{f}(A, x) \stackrel{\text{def}}{=} \mathop{\mathbb{E}}_{\pi} f(\pi.A, \pi.x),$$
$$g_A(x) \stackrel{\text{def}}{=} \mathop{\mathbb{E}}_{B \sim G^{\otimes R}(A)} \mathop{\mathbb{E}}_{B' \sim M_x(B)} \bar{f}(B', x).$$

By construction, each vertex (A, x) of H' has exactly the same neighborhood structure as $(\pi.A', \pi.x)$ for all $\pi \in S_R$ and $A' \in M_x(A)$. Hence, the fraction of edges crossing the cut can also be written in terms of $\bar{f}$ as $\langle f, H' f\rangle = \langle \bar{f}, H' \bar{f}\rangle$.

We will show that $\bar{f}$ gives a cut (actually, a distribution over cuts) with the same expansion in the graph H, such that the functions g_A satisfy $\mathbb{P}_A\{\mathbb{E}\, g_A \in (1/10, 9/10)\} \geqslant 1/10$. Recall that showing this was exactly the problem in making the reduction in Section V-A work.

Since $\mathbb{E}_A \mathbb{E}_x\, g_A = \mu$, we have $\mathbb{E}_A (\mathbb{E}_x\, g_A)^2 \geqslant \mu^2$. The following claim also gives an upper bound.

Claim V.2. $\mathbb{E}_A (\mathbb{E}_x\, g_A)^2 \leqslant \mu^2/2 + \mu/2$

Proof: We have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{A\sim V^R} \Big(\mathop{\mathbb{E}}_{x} g_A\Big)^2 \\
&= \mathop{\mathbb{E}}_{A\sim V^R} \Big(\mathop{\mathbb{E}}_{B\sim G^{\otimes R}(A)} \mathop{\mathbb{E}}_{x} \mathop{\mathbb{E}}_{B'\sim M_x(B)} \bar{f}\Big)^2 \\
&\leqslant \mathop{\mathbb{E}}_{A\sim V^R} \mathop{\mathbb{E}}_{B\sim G^{\otimes R}} \Big(\mathop{\mathbb{E}}_{x} \mathop{\mathbb{E}}_{B'\sim M_x(B)} \bar{f}\Big)^2 \\
&= \mathop{\mathbb{E}}_{B\sim V^R} \Big(\mathop{\mathbb{E}}_{x} \mathop{\mathbb{E}}_{B'\sim M_x(B)} \bar{f}\Big)^2 \\
&= \mathop{\mathbb{E}}_{B\sim V^R} \Big[\Big(\mathop{\mathbb{E}}_{x_1} \mathop{\mathbb{E}}_{B'_1\sim M_{x_1}(B)} \bar{f}\Big)\Big(\mathop{\mathbb{E}}_{x_2} \mathop{\mathbb{E}}_{B'_2\sim M_{x_2}(B)} \bar{f}\Big)\Big] \\
&= \mathop{\mathbb{E}}_{x_1} \mathop{\mathbb{E}}_{B'_1\sim M_{x_1}(B)} \bar{f}(B'_1, x_1) \mathop{\mathbb{E}}_{(B'_2,x_2)\sim M(B'_1,x_1)} f(B'_2, x_2).
\end{aligned}$$

For the last equality above, we define M to be a Markov operator which samples (B'_2, x_2) from the correct distribution given (B'_1, x'_1). Since x_1, x_2 are independent, x_2 can just be sampled uniformly. The fact that B'_1 and B'_2 come from the same (random) B can be captured by sampling each coordinate of B'_2 as

$$(B'_2)_i = \begin{cases} (B'_1)_i & \text{if } (x_1)_i = (x_2)_i = 1 \\ \text{random vertex in } V & \text{otherwise} \end{cases}.$$

Abusing notation, we also use M to denote the operator on the space of the functions which averages the value of the function over random (B'_2, x_2) generated as above. Then, if λ is the second eigenvalue of M, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{A}\Big(\mathop{\mathbb{E}}_{x} g_A\Big)^2 &\leqslant \langle \bar{f}, M\bar{f}\rangle \\
&\leqslant 1\cdot(\mathbb{E}\,\bar{f})^2 + \lambda\cdot\Big(\|\bar{f}\|^2 - (\mathbb{E}\,\bar{f})^2\Big) \\
&\leqslant (1-\lambda)\cdot\mu^2 + \lambda\cdot\mu.
\end{aligned}$$

Finally, it can be checked that the second eigenvalue of M is 1/2 which proves the claim. ■

This gives that $\mathbb{E}_x\, g_A$ cannot be always very far from μ. Formally,

$$\mathop{\mathbb{P}}_{A}\{|\mathbb{E}\, g_A - \mu| \geqslant \gamma\} \leqslant \frac{\mathbb{E}_A(\mathbb{E}\, g_A - \mu)^2}{\gamma^2} \leqslant \frac{\mu(1-\mu)}{2\gamma^2}.$$

Hence, for $\gamma = 2/5$, the probability is at most $25/32 < 9/10$. This can now be combined with the bound from Section V-A that gives

$$H'(S,S) \leqslant \mathop{\mathbb{E}}_{A} \Gamma_{1-\varepsilon}(\mathbb{E}\, g_A) + o(1).$$

Since $\mathbb{E}\, g_A \geqslant 1/10$ with probability at least $1/10$ over A, these "nice" A's contribute a volume of at least $1/100$. Also, for a nice A, we have $\Gamma_{1-\varepsilon}(\mathbb{E}\, g_A) \leqslant (\mathbb{E}\, g_A)(1 - \Omega(\sqrt{\varepsilon}))$. Hence,

$$H'(S,S) \leqslant (\mu - 1/100) + 1/100\cdot(1 - \Omega(\sqrt{\varepsilon})) + o(1)$$

which shows that S has expansion $\Omega(\sqrt{\varepsilon})$.

VI. Acknowledgments

This research was carried out while the authors where visiting Princeton University and Institute of Advanced Study. The first author would like to acknowledge support by NSF Career Grant CCF-1149843. The second author would like to acknowledge support from NSF Grants CCF-0832797, 0830673, and 0528414. The third author was supported by NSF grant CCF-0832797 and IAS Sub-contract no. 00001583.

References

[1] J. Cheeger, "A lower bound for the smallest eigenvalue of the Laplacian," in *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*. Princeton, N. J.: Princeton Univ. Press, 1970, pp. 195–199.

[2] N. Alon and V. D. Milman, "λ_1, isoperimetric inequalities for graphs, and superconcentrators," *J. Comb. Theory, Ser. B*, vol. 38, no. 1, pp. 73–88, 1985.

[3] N. Alon, "Eigenvalues and expanders," *Combinatorica*, vol. 6, no. 2, pp. 83–96, 1986.

[4] F. T. Leighton and S. Rao, "Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms," *J. ACM*, vol. 46, no. 6, pp. 787–832, 1999.

[5] N. Linial, E. London, and Y. Rabinovich, "The geometry of graphs and some of its algorithmic applications," *Combinatorica*, vol. 15, no. 2, pp. 215–245, 1995.

[6] Y. Aumann and Y. Rabani, "An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm," *SIAM J. Comput.*, vol. 27, no. 1, pp. 291–301, 1998.

[7] S. Arora, S. Rao, and U. V. Vazirani, "Expander flows, geometric embeddings and graph partitioning," in *STOC*, 2004, pp. 222–231.

[8] P. Raghavendra, D. Steurer, and P. Tetali, "Approximations for the isoperimetric and spectral profile of graphs and related parameters," in *STOC*, 2010, pp. 631–640.

[9] H. Räcke, "Optimal hierarchical decompositions for congestion minimization in networks," in *STOC*, 2008, pp. 255–264.

[10] N. Bansal, R. Krauthgamer, V. Nagarajan, K. Makarychev, J. S. Naor, and R. Schwartz, "Min-max graph partitioning and small set expansion," 2010, manuscript.

[11] P. Raghavendra and D. Steurer, "Graph expansion and the unique games conjecture," in *STOC*, 2010, pp. 755–764.

[12] S. Arora, B. Barak, and D. Steurer, "Subexponential algorithms for unique games and related problems," in *FOCS*, 2010, pp. 563–572.

[13] S. Khot, "On the power of unique 2-prover 1-round games," in *STOC*, 2002, pp. 767–775.

[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.*, vol. 37, no. 1, pp. 319–357, 2007.

[15] S. Khot and O. Regev, "Vertex cover might be hard to approximate to within $2-\varepsilon$," *J. Comput. Syst. Sci.*, vol. 74, no. 3, pp. 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*, 2005, pp. 53–62.

[17] P. Raghavendra, "Optimal algorithms and inapproximability results for every CSP?" in *STOC*, 2008, pp. 245–254.

[18] S. Arora, S. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. K. Vishnoi, "Unique games on expanding constraint graphs are easy," in *STOC*, 2008, pp. 21–28.

[19] U. Feige and C. Scheideler, "Improved bounds for acyclic job shop scheduling," *Combinatorica*, vol. 22, no. 3, pp. 361–399, 2002.

[20] S. Khot and R. Saket, "SDP integrality gaps with local ℓ_1-embeddability," in *FOCS*, 2009.

[21] P. Raghavendra and D. Steurer, "Integrality gaps for strong SDP relaxations of unique games," in *FOCS*, 2009, pp. 575–585.

[22] 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*, 2007, pp. 702–712.

[23] N. R. Devanur, S. Khot, R. Saket, and N. K. Vishnoi, "Integrality gaps for sparsest cut and minimum linear arrangement problems," in *STOC*, 2006, pp. 537–546.

[24] C. Borell, "Geometric bounds on the Ornstein-Uhlenbeck velocity process," *Z. Wahrsch. Verw. Gebiete*, vol. 70, no. 1, pp. 1–13, 1985. [Online]. Available: http://dx.doi.org/10.1007/BF00532234

[25] P. Raghavendra, D. Steurer, and M. Tulsiani, "Reductions between expansion problems," arXiv:1011.2586, 2010.

[26] P. Austrin, T. Pitassi, and Y. Wu, "Inapproximability of treewidth, one-shot pebbling, and related layout problems," arXiv:1109.4910v1, 2011.