THEORY OF COMPUTING, Volume 11 (9), 2015, pp. 241–256
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2013

A New Regularity Lemma and Faster Approximation Algorithms for Low Threshold-Rank Graphs

Shayan Oveis Gharan* Luca Trevisan†

Received February 4, 2014; Revised March 27, 2015; Published June 10, 2015

Abstract: Kolla and Tulsiani (2007, 2011) and Arora, Barak and Steurer (2010) introduced the technique of *subspace enumeration*, which gives approximation algorithms for graph problems such as unique games and small set expansion; the running time of such algorithms is exponential in the *threshold rank* of the graph.

Guruswami and Sinop (2011, 2012) and Barak, Raghavendra, and Steurer (2011) developed an alternative approach to the design of approximation algorithms for graphs of bounded threshold rank based on semidefinite programming relaxations obtained by using sum-of-squares hierarchy (2000, 2001) and on novel rounding techniques. These algorithms are faster than the ones based on subspace enumeration and work on a broad class of problems.

An extended abstract of this paper appeared in the proceedings of the 16th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2013) [15].

*Department of Computer Science and Engineering, University of Washington.

†Department of Electrical Engineering and Computer Sciences U.C. Berkeley. This material is based upon work supported by the National Science Foundation under grant No. CCF 1017403.

ACM Classification: F.2.2, G.2.2, G.1.6

AMS Classification: 68W25, 68Q25, 68R10

Key words and phrases: regularity lemma, approximation algorithms, expander graphs, low threshold-rank graphs, maximum cut

 DOI: 10.4086/toc.2015.v011a009

In this paper we develop a third approach to the design of such algorithms. We show, constructively, that graphs of bounded threshold rank satisfy a *weak Szemerédi regularity lemma* analogous to the one proved by Frieze and Kannan (1999) for dense graphs. The existence of efficient approximation algorithms is then a consequence of the regularity lemma, as shown by Frieze and Kannan.

Applying our method to the Max Cut problem, we devise an algorithm that is slightly faster than all previous algorithms, and is easier to describe and analyze.

1 Introduction

Kolla and Tulsiani [13, 12] and Arora, Barak and Steurer [2] proved that the Unique Games problem can be approximated efficiently if the adjacency matrix of a graph associated with the problem has few large eigenvalues; they show that, for every optimal solution, its indicator vector is close to the subspace spanned by the eigenvectors of the large eigenvalues, and one can find a solution close to an optimal one by enumerating an ε-net for such a subspace. This algorithm, also known as the *subspace enumeration* algorithm, runs in time exponential in the dimension of the subspace, which is the number of large eigenvalues; the number of large eigenvalues is called the *threshold rank* of the graph. Arora, Barak and Steurer show that the subspace enumeration algorithm can approximate other graph problems, in regular graphs, in time exponential in the threshold rank, including the Uniform Sparsest Cut problem, the Small-Set Expansion problem and the Max Cut problem. We remark that the subspace enumeration algorithm does not improve the 0.878 approximation guarantee of Goemans and Williamson [8], but it finds a solution of approximation factor $1-O(\varepsilon)$ if the optimum cuts at least $1-\varepsilon$ fraction of edges.

Barak, Raghavendra and Steurer [3] and Guruswami and Sinop [9, 10, 11] developed an alternative approach to the design of approximation algorithms running in time exponential in the threshold rank. Their algorithms are based on solving semidefinite programming relaxations obtained by using the sum-of-squares hierarchy [16, 14] and then applying sophisticated rounding schemes. This approach has several advantages. It is applicable to a more general class of graph problems and constraint satisfaction problems, that the approximation guarantee has a tighter dependency on the threshold used in the definition of threshold rank and that, in some cases, the algorithms have a running time of $f(k,\varepsilon)\cdot n^{O(1)}$ where k is the threshold rank and $1\pm\varepsilon$ is the approximation guarantee, instead of the running time of $n^{O(k)}$ which follows from an application of the subspace enumeration algorithm for constant ε.

In this paper we introduce a third approach to designing algorithms for graphs of bounded threshold rank, which is based on proving a *weak Szemerédi regularity lemma* for such graphs.

The regularity lemma of Szemerédi [17] states that every dense graph can be well approximated by the union of a constant number of bipartite complete subgraphs; the constant, however, has a tower-of-exponentials dependency on the quality of approximation. Frieze and Kannan [6, 7] prove what they call a *weak regularity lemma*, showing that every dense graph can be approximated up to an error εn^2 in the cut norm by a linear combination of $O(1/\varepsilon^2)$ cut matrices (a cut matrix is a bipartite complete subgraph) with bounded coefficients. Frieze and Kannan also show that such an approximation can be constructed “implicitly” in time polynomial in $1/\varepsilon$ and that, for a weighted graph which is a linear combination of σ cut matrices, several graph problems can be approximated in time $\exp(\tilde{O}(\sigma))+\mathrm{poly}(n)$

time.[1] Combining the two facts one has a $\exp(\text{poly}(1/\varepsilon)) + \text{poly}(n)$ time approximation algorithm for many graph problems on dense graphs.

We prove that a weak regularity lemma holds for all graphs of bounded threshold rank. Our result is a proper generalization of the weak regularity lemma of Frieze and Kannan, because dense graphs are known to have bounded threshold rank,[2] but there are many graphs with bounded threshold rank that are not dense. For a (weighted) $G = (V,E)$ with adjacency matrix A, and diagonal matrix of vertex degrees D, $D^{-1/2}AD^{-1/2}$ is called the normalized adjacency matrix of G. If the sum of squares of the eigenvalues of the normalized adjacency matrix outside the range $[-\varepsilon/2, \varepsilon/2]$ is equal to k (in particular, if there are at most k such eigenvalues), then we show that there is a linear combination of $O(k/\varepsilon^2)$ cut matrices that approximate A up to $2\varepsilon|E|$ in cut norm; furthermore, such a decomposition can be found in $\text{poly}(n,k,1/\varepsilon)$ time. (See Theorem 2.3 below.) Our regularity lemma, combined with an improvement of the Frieze-Kannan approximation algorithm for graphs that are linear combination of cut matrices, gives us algorithms of running time $2^{\tilde{O}(k^{1.5}/\varepsilon^3)} + \text{poly}(n)$ for several graph problems on graphs of threshold rank k, providing an additive approximation of $2\varepsilon|E|$. In problems such as Max Cut in which the optimum is $\Omega(|E|)$, this additive approximation is equivalent to a multiplicative approximation.

We remark that there are several generalizations of the weak regularity lemma to the matrices that are not necessarily dense, e. g., [5, 4], but to the best of our knowledge, none of these generalizations include matrices of low threshold rank. Let us provide a detailed example. Coja-Oghlan, Cooper and Frieze [4] consider sparse matrices that have a suitable boundedness property. For $S,T \subseteq V$, let the density of a sub-matrix $A_{S,T}$ be defined as follows:

$$\text{density}(A_{S,T}) := \frac{\sum_{u\in S, v\in T} A_{u,v}}{|S|\cdot|T|}.$$

Coja-Oghlan et al. [4] generalized weak regularity lemma to matrices where the density of each sub-matrix $A_{S,T}$ is within a constant factor, C, of the density of A, for any S,T such that $|S|,|T| \geq \Omega(n/2^{C^2})$. It turns out that if A represents the adjacency matrix of a graph that is a union of a constant number of constant degree expanders, then A has bounded threshold rank, but it doesn't satisfy the boundedness property.

Reference	Running time	Parameter k
BaRaSt [3]	$2^{O(k/\varepsilon^4)} \cdot \text{poly}(n)$	# of eigenvalues not in range $[-c\cdot\varepsilon^2, c\cdot\varepsilon^2]$, $c > 0$
GuSi [9]	$n^{O(k/\varepsilon^2)}$	# of eigenvalues $\leq -\varepsilon/2$
GuSi [10]	$2^{k/\varepsilon^3} \cdot n^{O(1/\varepsilon)}$	# of eigenvalues $\leq -\varepsilon/2$
this paper	$2^{\tilde{O}(k^{1.5}/\varepsilon^3)} + \text{poly}(n)$	sum of squares of eigenvalues not in range $[-\varepsilon/8, \varepsilon/8]$

Table 1: A comparison between previous algorithms applied to Max Cut and our algorithm.

Table 1 gives a comparison between previous algorithms applied to Max Cut and our algorithm. Unlike the previous algorithms, our algorithm rounds the solution to a fixed size LP, as opposed to a SDP hierarchy. The advantages over previous algorithms, besides the simplicity of the algorithm, is a faster

[1] Throughout the paper we use $\tilde{O}(.)$ to denote that logarithmic terms are ignored.

[2] The normalization one needs for dense graphs is different from what we use in this paper. If G is a graph with average degree $c \cdot n$, and A is the adjacency matrix of G, then $A/(c\cdot n)$ has a low threshold rank.

running time and the dependency on a potentially smaller threshold-rank parameter, because the running time of our algorithm depends on the *sum of squares* of eigenvalues outside of a certain range, rather than the number of such eigenvalues. Recall that the eigenvalues of $D^{-1/2}AD^{-1/2}$ are in the range $[-1,1]$.

We now give a precise statement of our results, after introducing some notation.

2 Statement of results

2.1 Notation

Let $G=(V,E)$ be a (weighted) undirected graph with $n:=|V|$ vertices. Let A be the adjacency matrix of G. For any vertex $u\in V$, let

$$d(u) := \sum_v A_{u,v}$$

be the degree of u. For a set $S\subset V$, let the *volume* of S be the summation of vertex degrees in S, $d(S)=\sum_{v\in S} d(v)$, and let

$$m := d(V) = \sum_{v\in V} d(v)\,.$$

Let D be the diagonal matrix of degrees. For any matrix $M\in\mathbb{R}^{V\times V}$, we use

$$M_{\mathcal{D}} := D^{-1/2}MD^{-1/2}\,.$$

Observe that if G is a d-regular graph, then $M_{\mathcal{D}}=M/d$. We call $A_{\mathcal{D}}$ the *normalized adjacency matrix* of G. It is straightforward to see that all eigenvalues of $A_{\mathcal{D}}$ are contained in the interval $[-1,1]$.

For two functions $f,g\in V\to\mathbb{R}$, let $\langle f,g\rangle := \sum_{v\in V} f(v)g(v)$. Also, let $f\otimes g$ be the tensor product of f,g; i. e., the matrix in $\mathbb{R}^{V\times V}$ such that (u,v) entry is $f(u)\cdot g(v)$. For a function $f\in\mathbb{R}^V$, and $S\subseteq V$ let $f(S):=\sum_{v\in S} f(v)$.

For a set $S\subseteq V$, let $\mathbf{1}_S$ be the indicator function of S, and let

$$d_S(v) := \begin{cases} d(v) & v\in S,\\ 0 & \text{otherwise.}\end{cases}$$

For any two sets $S,T\subseteq V$, and $\alpha\in\mathbb{R}$, we use the notation

$$\mathrm{CUT}(S,T,\alpha) := \alpha\cdot(d_S\otimes d_T)$$

to denote the matrix corresponding to the cut (S,T), where (u,v) entry of the matrix is $\alpha\cdot d(u)\cdot d(v)$ if $u\in S, v\in T$ and zero otherwise. We remark that $\mathrm{CUT}(S,T,\alpha)$ is not necessarily a symmetric matrix.

Definition 2.1 (Matrix norms)**.** For a matrix $M\in\mathbb{R}^{V\times V}$, and $S,T\subseteq V$, let

$$M(S,T) := \sum_{u\in S, v\in T} M_{u,v}\,.$$

The Frobenius norm and the cut norm are defined as follows:

$$\begin{aligned} \|M\|_F &:= \sqrt{\sum_{u,v} M_{u,v}^2}\,,\\ \|M\|_C &:= \max_{S,T\subseteq V} |M(S,T)|\,. \end{aligned}$$

Definition 2.2 (Sum-of-squares threshold rank)**.** For any unweighted graph G, with normalized adjacency matrix $A_{\mathcal{D}}$, let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of $A_{\mathcal{D}}$ with the corresponding eigenfunctions $f_1, \ldots, f_n$. For $\delta > 0$, the δ-sum-of-squares threshold rank of A is defined as

$$t_\delta(A_{\mathcal{D}}) := \sum_{i:|\lambda_i|>\delta} \lambda_i^2 .$$

Also, the δ-threshold approximation of $A_{\mathcal{D}}$ is defined as

$$T_\delta(A_{\mathcal{D}}) := \sum_{i:|\lambda_i|>\delta} \lambda_i f_i \otimes f_i .$$

In words, $T_\delta(A_{\mathcal{D}})$ is an approximation of $A_{\mathcal{D}}$ obtained by removing the projection on eigenvectors corresponding to the small eigenvalues. It is well known that $T_\delta(A_{\mathcal{D}})$ is the best approximation of its rank to $A_{\mathcal{D}}$ in ℓ_2 and Frobenius norm. In Lemma 3.1, below, we show that $D^{1/2} T_\delta(A_{\mathcal{D}}) D^{1/2}$ is always a good approximation of A in *cut norm*.

2.2 Matrix decomposition theorem

The following matrix decomposition theorem is the main technical result of this paper.

Theorem 2.3. *For any graph G, and $\varepsilon > 0$, let $k := t_{\varepsilon/2}(A_{\mathcal{D}})$. There is an algorithm that writes A as a linear combination of cut matrices, $W^{(1)}, W^{(2)}, \ldots, W^{(\sigma)}$, such that $\sigma \leq 16k/\varepsilon^2$, and*

$$\left\| A - W^{(1)} - \cdots - W^{(\sigma)} \right\|_C \leq \varepsilon m ,$$

where each $W^{(i)}$ is a cut matrix $\mathrm{CUT}(S, T, \alpha)$, for some $S, T \subseteq V$, such that $|\alpha| \leq \sqrt{k}/m$ and m is the sum of the degrees of all vertices of G. The running time of the algorithm is polynomial in $n, k, 1/\varepsilon$.

2.3 Algorithmic applications

Our main algorithmic application of Theorem 2.3 is the following theorem that approximates any cut on low threshold-rank graphs with a running time $2^{\tilde{O}(k^{1.5}/\varepsilon^3)} + \mathrm{poly}(n)$.

Theorem 2.4. *Let $G = (V, E)$, and for a given $\varepsilon > 0$, let $k := t_{\varepsilon/8}(A_{\mathcal{D}})$. There is a randomized algorithm such that for either of maximum cut or minimum cut problems over sets of volume*

$$\Gamma - \varepsilon m/2 \leq |S| \leq \Gamma + \varepsilon m/2 ,$$

in time $2^{\tilde{O}(k^{1.5}/\varepsilon^3)} + \mathrm{poly}(n, k, 1/\varepsilon)$, with constant probability finds a set S such that $|d(S) - \Gamma| \leq \varepsilon m$ and

$$A(S, \overline{S}) \geq \max_{\Gamma - \varepsilon m/2 \leq |S^*| \leq \Gamma + \varepsilon m/2} A(S^*, \overline{S^*}) - \varepsilon m$$

if it is a maximization problem, and

$$A(S, \overline{S}) \leq \min_{\Gamma - \varepsilon m/2 \leq |S^*| \leq \Gamma + \varepsilon m/2} A(S^*, \overline{S^*}) + \varepsilon m$$

otherwise.

In the minimum bisection problem we want to find the smallest cut with equal volume in both sides of the cut,

$$\min_{S:d(S)=m/2} A(S,\overline{S}).$$

Similarly, in the maximum bisection problem we want to find the maximum cut with equal volume in both sides. Although in the literature a bisection is typically defined as the cut with equal number of *vertices* in the both sides, here we study cuts with (approximately) equal volume in both sides. This is because of a limitation of spectral algorithms (cf. Cheeger's inequality for finding the minimum bisection). Nonetheless, the applications are very similar (e. g., we can use above corollary in divide and conquer algorithms to partition a given graph into small pieces with few edges in between).

We use the above theorem to provide a PTAS for maximum cut, maximum bisection, and minimum bisection problems.

Corollary 2.5. *Let $G=(V,E)$, and for a given $\varepsilon>0$, let $k:=t_{\varepsilon/8}(A_{\mathcal{D}})$. There is a randomized algorithm that in time $2^{\tilde{O}(k^{1.5}/\varepsilon^3)}+\mathrm{poly}(n,k,1/\varepsilon)$ finds an εm additive approximation of the maximum cut.*

Proof. We can simply guess the size of the optimum within an $\varepsilon m/2$ additive error and then use Theorem 2.4. □

Corollary 2.6. *Let $G=(V,E)$, and for a given $\varepsilon>0$, let $k:=t_{\varepsilon/8}(A_{\mathcal{D}})$. For any of the maximum bisection and minimum bisection problems, there is a randomized algorithm that in time $2^{\tilde{O}(k^{1.5}/\varepsilon^3)}+\mathrm{poly}(n,k,1/\varepsilon)$ finds a cut $(S,\overline{S})$ such that $|d(S)-m/2|\leq\varepsilon m$ and that $A(S,\overline{S})$ provides an εm additive approximation of the optimum.*

Proof. For the maximum/minimum bisection the optimum must have size $m/2$. So we can simply use Theorem 2.4 with $\Gamma=m/2$. □

3 Regularity lemma for low threshold-rank graphs

In this section we prove Theorem 2.3. The first step is to approximate A by a low-rank matrix B. In the next lemma we construct B such that the value of any cut in A is approximated within a small additive error in B.

Lemma 3.1. *Let A be the adjacency matrix of G. For $0\leq\delta<1$, let*

$$B:=D^{1/2}T_{\delta}(A_{\mathcal{D}})D^{1/2}.$$

Then, $\|A-B\|_C\leq\delta m$.

Proof. Let $\lambda_1,\ldots,\lambda_n$ be the eigenvalues of $A_{\mathcal{D}}$, with the corresponding orthonormal eigenfunctions

$f_1, \ldots, f_n$. For any $S, T \subseteq V$, we have

$$\begin{aligned}
\langle \mathbf{1}_S, (A-B)\mathbf{1}_T \rangle &= \langle D^{1/2}\mathbf{1}_S, (A-B)_{\mathcal{D}} D^{1/2}\mathbf{1}_T \rangle \\
&= \langle \sqrt{d_S}, (A_{\mathcal{D}} - T_\delta(A_{\mathcal{D}}))\sqrt{d_T} \rangle \\
&\leq \delta \cdot \sum_{i:|\lambda_i|\leq\delta} \langle \sqrt{d_S}, f_i \rangle \langle \sqrt{d_T}, f_i \rangle \\
&\leq \delta \cdot \sqrt{\sum_{i:|\lambda_i|\leq\delta} \langle \sqrt{d_S}, f_i \rangle^2} \cdot \sqrt{\sum_{i:|\lambda_i|\leq\delta} \langle \sqrt{d_T}, f_i \rangle^2} \\
&\leq \delta \cdot \left\| \sqrt{d_S} \right\| \cdot \left\| \sqrt{d_T} \right\| \leq \delta \cdot \left\| \sqrt{d_V} \right\|^2 = \delta m \,,
\end{aligned}$$

where the second inequality follows by the Cauchy–Schwarz inequality. The lemma follows by noting the fact that $\|A-B\|_C$ is the maximum of the above expression for any $S, T \subseteq V$. □

By the above lemma if we approximate B by a linear combination of cut matrices, then it is also a good approximation of A. Moreover, since $t_\delta(A_{\mathcal{D}}) = t_\delta(B_{\mathcal{D}})$, B has a small sum-of-squares threshold rank iff A has a small sum-of-squares threshold rank.

Lemma 3.2. *For any graph G with adjacency matrix A, $\delta > 0$, and $B = D^{1/2} T_\delta(A_{\mathcal{D}}) D^{1/2}$,*

$$\|B_{\mathcal{D}}\|_F^2 = t_\delta(A_{\mathcal{D}}) \,.$$

Proof. The lemma follows from the fact that the square of the Frobenius norm of any matrix is equal to the summation of square of eigenvalues. If $\lambda_1, \ldots, \lambda_n$ are the eigenvalues of $A_{\mathcal{D}}$, then

$$\|B_{\mathcal{D}}\|_F^2 = \operatorname{trace}({B_{\mathcal{D}}}^2) = \sum_{|\lambda_i|>\delta} \lambda_i^2 = t_\delta(A_{\mathcal{D}}) \,.$$

Note that in the first equality we are using the fact that $B_{\mathcal{D}}$ is a symmetric matrix. □

The next proposition is the main technical part of the proof of Theorem 2.3. We show that we can write any (not necessarily symmetric) matrix B as a linear combination of $O(\|B\|_F^2 / \varepsilon^2)$ cut matrices such that the cut norm of B is preserved within an additive error of εm. The proof builds on the existential theorem of Frieze and Kannan [7, Theorem 7].

Proposition 3.3. *For any matrix $B \in \mathbb{R}^{V \times V}$, $k = \|B_{\mathcal{D}}\|_F^2$, and $\varepsilon > 0$, there exist cut matrices*

$$W^{(1)}, W^{(2)}, \ldots, W^{(\sigma)} \,,$$

such that $\sigma \leq 1/\varepsilon^2$, and for all $S, T \subseteq V$,

$$\left| \left(B - W^{(1)} - W^{(2)} - \cdots - W^{(\sigma)} \right)(S,T) \right| \leq \varepsilon \sqrt{k \cdot d(S) \cdot d(T)} \,,$$

where each $W^{(i)}$ is a cut matrix $\mathrm{CUT}(S, T, \alpha)$, for some $S, T \subseteq V$, and $\alpha \in \mathbb{R}$.

Proof. Let $R^{(0)} = B$. We use the potential function $h(R) := \|R_{\mathcal{D}}\|_F^2$. We show that as long as there are $S,T \subseteq V$ such that

$$|R(S,T)| > \varepsilon\sqrt{kd(S)d(T)}$$

we can add new cut matrices iteratively while maintaining the invariant that each time the value of the potential function decreases by at least $\varepsilon^2 h(B)$. Since $h(R^{(0)}) = h(B)$, after at most $1/\varepsilon^2$ we obtain a good approximation of B.

Assume that after $i < 1/\varepsilon^2$ iterations, $R^{(i)} = B - W^{(1)} - \cdots - W^{(i)}$. Suppose for some $S,T \subseteq V$,

$$\left|R^{(i)}(S,T)\right| > \varepsilon \cdot \sqrt{h(B)\cdot d(S)\cdot d(T)} = \varepsilon \cdot \sqrt{k \cdot d(S)\cdot d(T)}\,. \tag{3.1}$$

Choose $W^{(i+1)} = \mathrm{CUT}(S,T,\alpha)$, for

$$\alpha = \frac{R^{(i)}(S,T)}{d(S)\cdot d(T)}\,,$$

and let $R^{(i+1)} = R^{(i)} - W^{(i+1)}$. Note that for any pair of vertices $u,v \in V$ if $u \notin S$ or $v \notin T$, then $R^{(i+1)}_{u,v} = R^{(i)}_{u,v}$. So,

$$\begin{aligned} h(R^{(i+1)}) - h(R^{(i)}) &= \sum_{u\in S, v\in T} \frac{(R^{(i)}_{u,v} - \alpha d(u)d(v))^2 - {R^{(i)}_{u,v}}^2}{d(u)d(v)} \\ &= -2\alpha R^{(i)}(S,T) + \alpha^2 d(S)d(T) \\ &= \frac{-R^{(i)}(S,T)^2}{d(S)d(T)} \le -\varepsilon^2 \cdot h(B)\,. \end{aligned}$$

The last equality follows by the definition of α, and the last inequality follows from Equation (3.1). Therefore, after at most $\sigma \le 1/\varepsilon^2$ iterations, (3.1) cannot hold for all $S,T \subseteq V$. □

Although the previous proposition only proves the existence of a decomposition into cut matrices, we can construct such a decomposition efficiently using the following nice result of Alon and Naor [1] that gives a constant factor approximation algorithm for the cut norm of any matrix.

Theorem 3.4 (Alon and Naor [1])**.** *There is a polynomial time randomized algorithm such that for any given $A \in \mathbb{R}^{V\times V}$, with high probability, finds sets $S,T \subseteq V$, such that*

$$|A(S,T)| \ge 0.56\,\|A\|_C\,.$$

Now we are ready to prove Theorem 2.3.

Proof of Theorem 2.3. Let $\delta := \varepsilon/2$, and $B := D^{1/2}T_\delta(A_{\mathcal{D}})D^{1/2}$. By Theorem 3.1, we have that

$$\|A - B\|_C \le \delta m = \varepsilon m/2\,. \tag{3.2}$$

So we just need to approximate B by a set of cut matrices within an additive error of $\varepsilon m/2$. For a matrix R, let $h(R) := \|R_{\mathcal{D}}\|_F^2$. By Lemma 3.2 we have $h(B) = k$.

Let $\varepsilon' := \varepsilon/\sqrt{4k}$. We use the proof strategy of Proposition 3.3. Let $R^{(i)} = B - W^{(1)} - \cdots - W^{(i)}$. If $\left\|R^{(i)}\right\|_C \geq \varepsilon'\sqrt{k}m$, then by Theorem 3.4 in polynomial time we can find $S,T \subseteq V$ such that

$$\left|R^{(i)}(S,T)\right| \geq \varepsilon' \cdot \sqrt{k} \cdot m/2 \geq \varepsilon' \cdot \sqrt{h(B)} \cdot m/2\,. \tag{3.3}$$

Choose $W^{(i+1)} = \mathrm{CUT}(S,T,\alpha)$, for $\alpha = R^{(i)}(S,T)/m^2$, and let $R^{(i+1)} = R^{(i)} - W^{(i+1)}$. We get

$$h(R^{(i+1)}) - h(R^{(i)}) = -2\alpha R^{(i)}(S,T) + \alpha^2 d(S)d(T) \leq -\frac{R^{(i)}(S,T)^2}{m^2} \leq -\frac{\varepsilon'^2 \cdot h(B)}{4}\,.$$

Since $h(R^{(0)}) = h(B)$, after $\sigma \leq 4/\varepsilon'^2 = 16k/\varepsilon^2$, we have $\left\|R^{(\sigma)}\right\|_C \leq \varepsilon'\sqrt{k}m$. Using the triangle inequality on the cut norm we obtain

$$\begin{aligned}\left\|A - W^{(1)} - \cdots - W^{(\sigma)}\right\|_C &\leq \|A - B\|_C + \left\|B - W^{(1)} - \cdots - W^{(\sigma)}\right\|_C \\ &\leq \varepsilon m/2 + \varepsilon'\sqrt{k}m = \varepsilon m\,.\end{aligned}$$

where the second inequality uses (3.2).

This proves the correctness of the algorithm. It remains to upper bound α. For each cut matrix $W^{(i)} = \mathrm{CUT}(S,T,\alpha)$ constructed throughout the algorithm we have

$$\begin{aligned}|\alpha| = \frac{|R^{(i)}(S,T)|}{m^2} &= \frac{1}{m^2}\left|\sum_{u\in S, v\in T} R^{(i)}_{u,v}\frac{\sqrt{d(u)d(v)}}{\sqrt{d(u)d(v)}}\right| \\ &\leq \frac{1}{m^2}\sqrt{\sum_{u\in S, v\in T}\frac{{R^{(i)}_{u,v}}^2}{d(u)d(v)}}\sqrt{d(S)d(T)} \\ &\leq \frac{1}{m}\sqrt{\sum_{u\in S, v\in T}\frac{{R^{(i)}_{u,v}}^2}{d(u)d(v)}} \\ &= \frac{\sqrt{h(R^{(i)})}}{m} \leq \frac{\sqrt{h(B)}}{m} = \frac{\sqrt{k}}{m}\,.\end{aligned}$$

where the first inequality follows by the Cauchy–Schwarz inequality, the second inequality uses $d(S), d(T) \leq m$, and the last inequality follows by the fact that the potential function is decreasing throughout the algorithm. This completes the proof of theorem. □

4 Fast approximation algorithm for low threshold-rank graphs

In this section we prove Theorem 2.4. First, by Theorem 2.3 in time $\mathrm{poly}(n,k,1/\varepsilon)$ we can find cut matrices $W^{(1)},\ldots,W^{(\sigma)}$ for $\sigma = O(k/\varepsilon^2)$, such that for all $1 \leq i \leq t$, $W^{(i)} = \mathrm{CUT}(S_i,T_i,\alpha_i)$, $\alpha_i \leq \sqrt{k}/m$, and

$$\|A - W\|_C \leq \varepsilon m/4\,,$$

where $W := W^{(1)} + \cdots + W^{(\sigma)}$. It follows from the above equation that for any set $S \subseteq V$,

$$|A(S,\overline{S}) - W(S,\overline{S})| = \left| A(S,\overline{S}) - \sum_{i=1}^{\sigma} \alpha_i \cdot d(S \cap S_i) \cdot d(\overline{S} \cap T_i) \right| \leq \frac{\varepsilon m}{4}. \tag{4.1}$$

Fix $S^* \subseteq V$ of volume $\Gamma - \varepsilon/2 \leq d(S^*) \leq \Gamma + \varepsilon/2$ (think of $(S^*, \overline{S^*})$ as the optimum cut), and let $s_i^* := d(S_i \cap S^*)$, and $t_i^* := d(T_i \cap \overline{S^*})$. Observe that by Equation (4.1),

$$\left| A(S^*, \overline{S^*}) - \sum_{i=1}^{\sigma} \alpha_i s_i^* t_i^* \right| \leq \frac{\varepsilon m}{4}. \tag{4.2}$$

Let $\alpha_{\max} := \max_{1 \leq i \leq \sigma} |\alpha_i|$. Choose $\Delta = \Theta(\varepsilon^3 m / k^{1.5})$ such that

$$\Delta \leq \min\left\{ \frac{\varepsilon^3 \cdot m}{48}, \frac{\varepsilon}{48 \alpha_{\max} \cdot \sigma} \right\}. \tag{4.3}$$

Note that this is achievable since $k \geq 1$, $\alpha_{\max} \leq \sqrt{k}/m$ and $\sigma = O(k/\varepsilon^2)$.

We define an approximation of s_i^*, t_i^* by rounding them down to the nearest multiple of Δ, i. e.,

$$\tilde{s}_i^* := \Delta \cdot \lfloor s_i^* / \Delta \rfloor,$$
$$\tilde{t}_i^* := \Delta \cdot \lfloor t_i^* / \Delta \rfloor.$$

We use $\tilde{s}^*, \tilde{t}^*$ to denote the vectors of the approximate values. It follows that we can obtain a good approximation of the size of the cut $(S^*, \overline{S^*})$ just by guessing the vectors $\tilde{s}^*$ and $\tilde{t}^*$. Since $|s_i^* - \tilde{s}_i^*| \leq \Delta$ and $|t_i^* - \tilde{t}_i^*| \leq \Delta$, we get

$$\sum_{i=1}^{\sigma} |s_i^* t_i^* \alpha_i - \tilde{s}_i^* \tilde{t}_i^* \alpha_i| \leq \sigma \cdot \alpha_{\max} (2 \cdot \Delta \cdot m + \Delta^2) \leq 3 \alpha_{\max} \cdot \sigma \cdot \Delta \cdot m \leq \varepsilon \cdot m / 16, \tag{4.4}$$

where we used (4.3).

Observe that by Equations (4.1), (4.2), and (4.4), if we know the vectors $\tilde{s}^*, \tilde{t}^*$, then we can find $A(S^*, \overline{S^*})$ within an additive error of $\varepsilon m/2$. Since $\tilde{s}_i^*, \tilde{t}_i^* \leq m$, there are only $O(m/\Delta)$ possibilities for each $\tilde{s}_i^*$ and $\tilde{t}_i^*$. Therefore, we afford to enumerate all possible values of them in time $(m/\Delta)^{2\sigma}$, and choose the one that gives the largest cut. Unfortunately, for a given assignment of $\tilde{s}^*, \tilde{t}^*$ the corresponding cut $(S^*, \overline{S^*})$ may not exist. Next we give an algorithm that for a given assignment of $\tilde{s}^*, \tilde{t}^*$ finds a cut $(S, \overline{S})$ such that

$$A(S, \overline{S}) = \sum_i \tilde{s}_i^* \tilde{t}_i^* \alpha_i \pm \varepsilon m,$$

if one exists.

First we distinguish the large degree vertices of G and simply guess which side they are mapped to in the optimum cut. For the rest of the vertices we use the solution of LP(1). Let

$$U := \{v : d(v) \geq \Delta\}$$

be the set of large degree vertices. Observe that $|U| \leq m/\Delta$. Let $\mathcal{P}$ be the coarsest partition of the set $V \setminus U$ such that for any $1 \leq i \leq \sigma$, both $S_i \setminus U$ and $T_i \setminus U$ can be written as a union of sets in $\mathcal{P}$, and for each $P \in \mathcal{P}$, $d(P) \leq \Delta$. Observe that $|\mathcal{P}| \leq 2^{2\sigma} + m/\Delta$. For a given assignment of $\tilde{s}^*, \tilde{t}^*$, first we guess the set of vertices in U that are contained in S^*, $U_{S^*} := S^* \cap U$, and $U_{\overline{S^*}} := U \setminus U_{S^*}$. For the rest of the vertices we use the linear program LP(1) to find the unknown $d(S^* \cap P)$.

$$\begin{array}{lclclll}
\text{LP(1)} & & & & & & \\
0 & \leq & y_P & \leq & 1 & \forall P \in \mathcal{P} & \\
\Gamma - \varepsilon m/2 & \leq & \sum_P y_P d(P) + d(U_{S^*}) & \leq & \Gamma + \varepsilon m/2 & & (4.5) \\
\tilde{s}_i^* & \leq & \sum_{P \subseteq S_i} y_P d(P) + d(U_{S^*} \cap S_i) & \leq & \tilde{s}_i^* + \Delta & \forall 1 \leq i \leq \sigma & (4.6) \\
\tilde{t}_i^* & \leq & \sum_{P \subseteq T_i} (1 - y_P) d(P) + d(U_{\overline{S^*}} \cap T_i) & \leq & \tilde{t}_i^* + \Delta & \forall 1 \leq i \leq \sigma. & (4.7)
\end{array}$$

Observe that $y_P = d(S^* \cap P)/d(P)$ is a feasible solution to the linear program. In the next lemma which is the main technical part of the analysis we show how to construct a set based on a given solution of the LP.

Lemma 4.1. *There is a randomized algorithm such that for any $S^* \subset V$, given $\tilde{s}_i^*, \tilde{t}_i^*$ and U_{S^*} returns a random set S such that*

$$\mathbb{P}\left[W(S, \overline{S}) \geq A(S^*, \overline{S^*}) - \frac{3\varepsilon m}{4} \wedge |d(S) - \Gamma| \leq \varepsilon m\right] \quad \geq \quad \frac{\varepsilon}{4}, \tag{4.8}$$

$$\mathbb{P}\left[W(S, \overline{S}) \leq A(S^*, \overline{S^*}) + \frac{3\varepsilon m}{4} \wedge |d(S) - \Gamma| \leq \varepsilon m\right] \quad \geq \quad \frac{\varepsilon}{4}. \tag{4.9}$$

Proof. Let y be a feasible solution of LP(1). We use a simple independent rounding scheme to compute a *random* set S. We always include U_{S^*} in S. For each $P \in \mathcal{P}$, we include P in S, independently, with probability y_P. We prove that S satisfies the lemma's statements.

First of all, by linearity of expectation,

$$\mathbb{E}[d(S \cap S_i)] = d(U_{S^*}) + \sum_{P \subseteq S_i} y_P d(P), \quad \text{and}$$
$$\mathbb{E}\left[d(\overline{S} \cap T_i)\right] = d(U_{\overline{S^*}}) + \sum_{P \subseteq T_i} (1 - y_P) d(P).$$

Therefore, by (4.5),

$$|\mathbb{E}[d(S)] - \Gamma| \leq \varepsilon m/2. \tag{4.10}$$

Furthermore, by (4.6) and (4.7) for any $1 \leq i \leq \sigma$,

$$|\mathbb{E}[S \cap S_i] - \tilde{s}_i^*| \leq \Delta \quad \text{and} \quad |\mathbb{E}\left[\overline{S} \cap T_i\right] - \tilde{t}_i^*| \leq \Delta. \tag{4.11}$$

In the following claim we show that $d(S)$ is highly concentrated around its expectation. This shows that $d(S)$ is close to Γ.

Claim 4.2. $\mathbb{P}[|d(S) - \Gamma| \geq \varepsilon m] \leq \dfrac{\varepsilon}{12}$.

Proof. We use the theorem of Hoeffding to prove the claim:

Theorem 4.3 (Hoeffding Inequality)**.** *Let $X_1, \ldots, X_n$ be independent random variables such that for each $1 \leq i \leq n$, $X_i \in [0, a_i]$. Let $X := \sum_{i=1}^n X_i$. Then, for any $\varepsilon > 0$*

$$\mathbb{P}[|X - \mathbb{E}[X]\,| \geq \varepsilon] \leq 2\exp\left(-\frac{2\varepsilon^2}{\sum_{i=1}^n a_i^2}\right).$$

Now, by the independent rounding procedure, we obtain

$$\begin{aligned}\mathbb{P}[|d(S) - \mathbb{E}[d(S)]\,| \geq \varepsilon m/2] \leq 2\exp\left(-\frac{\varepsilon^2 m^2}{2\sum_P d(P)^2}\right) &\leq 2\exp\left(-\frac{\varepsilon^2 m^2}{2m\Delta}\right) \\ &\leq 2\exp(-24/\varepsilon) \leq \frac{\varepsilon}{12},\end{aligned}$$

where the second inequality follows by the fact that $d(P) \leq \Delta$ and $\sum_P d(P) \leq m$ and the third inequality follows by (4.3). The claim flows by (4.10). □

In the next claim we upper bound the expected value of $W(S, \overline{S}) - A(S^*, \overline{S^*})$.

Claim 4.4. $\left|\mathbb{E}\left[W(S,\overline{S})\right] - A(S^*, \overline{S^*})\right| \leq \dfrac{\varepsilon m}{2}$.

Proof. First, we calculate $\mathbb{E}\left[W(S,\overline{S})\right]$ in terms of $\mathbb{E}[d(S \cap S_i)]\,, \mathbb{E}\left[d(\overline{S} \cap T_i)\right]$.

$$\begin{aligned}\mathbb{E}\left[W(S,\overline{S})\right] &= \mathbb{E}\left[\sum_{i=1}^{\sigma} d(S\cap S_i) d(\overline{S}\cap T_i)\alpha_i\right] \\ &= \sum_{i=1}^{\sigma} \alpha_i \mathbb{E}\left[\left(\sum_{P\in\mathcal{P}:P\subseteq S_i} d(P)\,\mathbb{I}[P\subseteq S]\right)\left(\sum_{Q\in\mathcal{P}:Q\subseteq T_i} d(Q)\,\mathbb{I}\left[Q\subseteq \overline{S}\right]\right)\right] \\ &\quad + \sum_{i=1}^{\sigma} \alpha_i \Big(d(U_{S^*}\cap S_i)\,\mathbb{E}\left[d(\overline{S}\cap T_i)\right] + d(U_{\overline{S^*}}\cap T_i)\,\mathbb{E}[d(S\cap S_i)]\Big).\end{aligned} \tag{4.12}$$

Since the event that $P \subseteq S$ is independent of $Q \subseteq S$, iff $P \neq Q$ we get

$$\mathbb{E}\left[\mathbb{I}[P\subseteq S]\,\mathbb{I}\left[Q\subseteq\overline{S}\right]\right] = \begin{cases} y_P(1-y_Q) & \text{if } P\neq Q, \\ 0 & \text{otherwise.}\end{cases}$$

Let $s_i := \mathbb{E}[d(S\cap S_i)]$ and $t_i := \mathbb{E}\left[d(\overline{S}\cap T_i)\right]$. By (4.12) and the above equation,

$$\mathbb{E}\left[W(S,\overline{S})\right] \quad = \quad \sum_{i=1}^{\sigma}\alpha_i s_i t_i - \sum_{i=1}^{\sigma}\alpha_i \sum_{P\in\mathcal{P}} y_P(1-y_P)d(P)^2. \tag{4.13}$$

Using Equation (4.2) we write

$$\begin{aligned}\left|\mathbb{E}\left[W(S,\overline{S})\right]-A(S^*,\overline{S^*})\right| &\leq \left|\mathbb{E}\left[W(S,\overline{S})\right]-\sum_{i=1}^{\sigma}\alpha_i s_i^* t_i^*\right|+\frac{\varepsilon m}{4}\\ &= \left|\sum_{i=1}^{\sigma}(\alpha_i s_i t_i-\alpha_i s_i^* t_i^*)-\sum_{i=1}^{\sigma}\sum_{P\in\mathcal{P}}\alpha_i y_P(1-y_P)d(P)^2\right|+\frac{\varepsilon m}{4}\\ &\leq \sum_{i=1}^{\sigma}|\alpha_i s_i t_i-\alpha_i\tilde{s}_i^*\tilde{t}_i^*|+\sum_{i=1}^{\sigma}|\alpha_i\tilde{s}_i^*\tilde{t}_i^*-\alpha_i s_i^* t_i^*|+\sigma\alpha_{\max}m\Delta+\frac{\varepsilon m}{4}\\ &\leq \sum_{i=1}^{\sigma}|\alpha_i s_i t_i-\alpha_i\tilde{s}_i^*\tilde{t}_i^*|+\frac{\varepsilon m}{16}+\frac{\varepsilon m}{48}+\frac{\varepsilon m}{4}\leq\frac{\varepsilon m}{2},\end{aligned}$$

where the equality follows by (4.13), the second inequality follows by the fact that $d(P)\leq\Delta$ for all $P\in\mathcal{P}$ and $\sum_P d(P)\leq m$, the third inequality follows by (4.4) and (4.3), and the last equation can be proved similar to (4.4) using (4.11). This completes the proof of Claim 4.4. □

Now, we are ready finish the proof of Lemma 4.1. Here, we prove (4.8). Equation (4.9) can be proved similarly. By Claim 4.4,

$$\mathbb{E}\left[W(S,\overline{S})\right]\geq A(S^*,\overline{S^*})-\frac{\varepsilon m}{2}.$$

Since the size of any cut in G is at most $m/2$, by (4.1), for any $S\subseteq V$, $W(S,\overline{S})\leq\varepsilon m/4+m/2<3m/4$. Therefore, by Markov's inequality,

$$\mathbb{P}\left[W(S,\overline{S})\geq A(S^*,\overline{S^*})-\frac{3\varepsilon m}{4}\right]\geq\frac{\varepsilon}{3}.$$

By the union bound, Lemma 4.1 follows from the discussion above and Claim 4.2. □

Our rounding algorithm is described in Algorithm 1. First, we prove the correctness, then we calculate the running time of the algorithm. Let S be the output set of the algorithm. First, observe that the output always satisfies $|d(S)-\Gamma|\leq\varepsilon m$. Now let $A(S^*,\overline{S^*})$ be the maximum cut among all sets of size Γ (the minimization case can be proved similarly). In the iteration that the algorithm correctly guesses $\tilde{s}_i^*,\tilde{t}_i^*,U_{S^*}$, there exists a feasible solution y of LP(1). by Lemma 4.1, for all $1\leq i\leq 4/\varepsilon$,

$$\mathbb{P}\left[W(S_y(i),\overline{S_y(i)})\geq A(S^*,\overline{S^*})-\frac{3\varepsilon m}{4}\wedge|d(S_y(i))-\Gamma|\leq\varepsilon m\right]\geq\frac{\varepsilon}{4}.$$

Since we take the best of $4/\varepsilon$ samples, with probability $1/e$ the output set S satisfies

$$W(S,\overline{S})\geq A(S^*,\overline{S^*})-3\varepsilon m/4.$$

But by (4.1), we must have $A(S,\overline{S})\geq A(S^*,\overline{S^*})-\varepsilon m$. This proves the correctness of the algorithm.

It remains to upper-bound the running time of the algorithm. First observe that if $|U|=O(k/\varepsilon^2)$, the running time of the algorithm is dominated by the time it takes to compute a feasible solution of

Algorithm 1 Approximate Maximum Cut $(S,\overline{S})$ such that $d(S) = \Gamma \pm \varepsilon m$

for all possible values of $\tilde{s}_i^*, \tilde{t}_i^*$, and $U_{S^*} \subseteq U$ **do**
  if there is a feasible solution y of LP(1) **then**
    for $i = 1 \to 4/\varepsilon$ **do**
      $S_y(i) \leftarrow U_{S^*}$.
      For each $P \in \mathcal{P}$ include P in $S_y(i)$, independently, with probability y_P.
    end for
  end if
end for
return among all sets $S_y(i)$ sampled in the loop that satisfy $|d(S_y(i)) - \Gamma| \le \varepsilon m$, the one that $W(S_y(i), \overline{S_y(i)})$ is the maximum.

LP(1). Since the size of LP is $2^{\tilde{O}(k/\varepsilon^2)}$, in this case Algorithm 1 terminates in time $2^{\tilde{O}(k/\varepsilon^2)}$. Note that for any sample set $S_y(i)$, both $d(S_y(i))$ and $W(S_y(i), \overline{S_y(i)})$ can be computed in time $2^{\tilde{O}(k/\varepsilon^2)}$, once we know $|S_y(i) \cap P|$ for any $P \in \mathcal{P}$.

Otherwise if $|U| \gg k/\varepsilon^2$, the dependence of the running time of the algorithm on ε, k is dominated by the step where we guess the subset of $U_{S^*} = U \cap S^*$. By (4.3),

$$|U| \le \frac{m}{\Delta} = O\left(\frac{k^{1.5}}{\varepsilon^3}\right).$$

Therefore, Algorithm 1 runs in time $2^{\tilde{O}(k^{1.5}/\varepsilon^3)}$. Since it takes $\mathrm{poly}(n,k,1/\varepsilon)$ to compute the decomposition into $W^{(1)}, \dots, W^{(\sigma)}$, the total running time is $2^{\tilde{O}(k^{1.5}/\varepsilon^3)} + \mathrm{poly}(n,k,1/\varepsilon)$. This completes the proof of Theorem 2.4.

Acknowledgement.

We would like to thank an anonymous reviewer for helpful comments.

References

[1] NOGA ALON AND ASSAF NAOR: Approximating the cut-norm via Grothendieck's inequality. *SIAM J. Comput.*, 35(4):787–803, 2006. Preliminary version in STOC'04. [doi:10.1137/S0097539704441629] 248

[2] SANJEEV ARORA, BOAZ BARAK, AND DAVID STEURER: Subexponential algorithms for Unique Games and related problems. In *Proc. 51st FOCS*, pp. 563–572. IEEE Comp. Soc. Press, 2010. [doi:10.1109/FOCS.2010.59] 242

[3] BOAZ BARAK, PRASAD RAGHAVENDRA, AND DAVID STEURER: Rounding semidefinite programming hierarchies via global correlation. In *Proc. 52nd FOCS*, pp. 472–481. IEEE Comp. Soc. Press, 2011. [doi:10.1109/FOCS.2011.95] 242, 243

[4] AMIN COJA-OGHLAN, COLIN COOPER, AND ALAN M. FRIEZE: An efficient sparse regularity concept. *SIAM J. Discrete Math.*, 23(4):2000–2034, 2010. Preliminary version in SODA'09. [doi:10.1137/080730160] 243

[5] WENCESLAS FERNANDEZ DE LA VEGA, MAREK KARPINSKI, RAVI KANNAN, AND SANTOSH VEMPALA: Tensor decomposition and approximation schemes for constraint satisfaction problems. In *Proc. 37th STOC*, pp. 747–754. ACM Press, 2005. [doi:10.1145/1060590.1060701] 243

[6] ALAN M. FRIEZE AND RAVI KANNAN: The regularity lemma and approximation schemes for dense problems. In *Proc. 37th FOCS*, pp. 12–20. IEEE Comp. Soc. Press, 1996. [doi:10.1109/SFCS.1996.548459] 242

[7] ALAN M. FRIEZE AND RAVI KANNAN: Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999. [doi:10.1007/s004930050052] 242, 247

[8] MICHEL X. GOEMANS AND DAVID P. WILLIAMSON: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. Preliminary version in STOC'94. [doi:10.1145/227683.227684] 242

[9] VENKATESAN GURUSWAMI AND ALI KEMAL SINOP: Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with PSD objectives. In *Proc. 52nd FOCS*, pp. 482–491. IEEE Comp. Soc. Press, 2011. Expanded version in CoRR (2011). Merged with [11] in CoRR (2013). [doi:10.1109/FOCS.2011.36] 242, 243, 255

[10] VENKATESAN GURUSWAMI AND ALI KEMAL SINOP: Faster SDP hierarchy solvers for local rounding algorithms. In *Proc. 53rd FOCS*, pp. 197–206. IEEE Comp. Soc. Press, 2012. Expanded version in CoRR. [doi:10.1109/FOCS.2012.58] 242, 243

[11] VENKATESAN GURUSWAMI AND ALI KEMAL SINOP: Approximating non-uniform sparsest cut via generalized spectra. In *Proc. 24th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'13)*. SIAM, 2013. Expanded version, merged with [9] in CoRR (2013). [doi:10.1137/1.9781611973105.22] 242, 255

[12] ALEXANDRA KOLLA: Spectral algorithms for Unique Games. *Comput. Complexity*, 20(2):177–206, 2011. Preliminary versions in CCC'10 and ECCC (2010). [doi:10.1007/s00037-011-0011-7] 242

[13] ALEXANDRA KOLLA AND MADHUR TULSIANI: Playing Unique Games using graph spectra. manuscript, 2007. 242

[14] JEAN B. LASSERRE: Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001. [doi:10.1137/S1052623400366802] 242

[15] SHAYAN OVEIS GHARAN AND LUCA TREVISAN: A new regularity lemma and faster approximation algorithms for low threshold rank graphs. In *Proc. 16th Internat. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'13)*, pp. 303–316, 2013. [doi:10.1007/978-3-642-40328-6_22] 241

[16] PABLO A. PARRILO: *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. Ph. D. thesis, California Institute of Technology, 2000. 242

[17] ENDRE SZEMERÉDI: Regular partitions of graphs. *Colloq. Internat. CNRS: Problèmes Combinatoires et Théorie des Graphes*, 260:399–401, 1978. 242

AUTHORS

Shayan Oveis Gharan
Assistant Professor
University of Washington, Seattle, WA
shayan@cs.washington.edu
http://homes.cs.washington.edu/~shayan

Luca Trevisan
Professor
U.C. Berkeley, Berkeley, CA
luca@berkeley.edu
http://www.cs.berkeley.edu/~luca/

ABOUT THE AUTHORS

SHAYAN OVEIS GHARAN (listed alphabetically under "O") is an assistant professor in the Computer Science and Engineering department at the University of Washington. He received his Ph. D. in 2013 from the MS&E department at Stanford University under the supervision of Amin Saberi and Luca Trevisan. Before joining UW he spent a year and a half as a postdoctoral Miller Fellow at UC Berkeley where his host was Umesh Vazirani. He did his undergraduate studies at the Computer Engineering department at Sharif University. Shayan's research areas include the design and analysis of algorithms, spectral graph theory and applied probability. Shayan received an ACM doctoral dissertation award honorable mention for his Ph. D. thesis "New Rounding Techniques for the Design and Analysis of Approximation Algorithms" in 2014. He and his coauthors received the Best Paper awards at SODA 2010 and FOCS 2011 for their work on the Traveling Salesman Problem.

LUCA TREVISAN is a professor of electrical engineering and computer sciences at U.C. Berkeley and a senior scientist at the Simons Institute for the Theory of Computing. Luca received his Dottorato (Ph. D.) in 1997, from the Sapienza University of Rome, working with Pierluigi Crescenzi. After graduating, Luca was a post-doc at MIT and at DIMACS, and he was on the faculty of Columbia University, U.C. Berkeley, and Stanford, before returning to Berkeley in 2014. Luca's research is in theoretical computer science, and in the past six years he has been interested in spectral graph theory and its applications to graph algorithms. Luca lives, beyond his means, in San Francisco. When out of town, he can often be found in Rome or in Hong Kong.