Playing Random and Expanding Unique Games

Alexandra Kolla* Madhur Tulsiani†

Abstract

We analyze the behavior of the SDP by Feige and Lovász [FL92] on random instances of unique games. We show that on random d-regular graphs with permutations chosen at random, the value of the SDP is very small with probablility $1-e^{-\Omega(d)}$. Hence, the SDP provides a proof of unsatisfiability for random unique games. We also give a spectral algorithm for recovering planted solutions. Given a random instance consistent with a given solution on $1-\epsilon$ fraction of the edges, our algorithm recovers a solution with value $1-O(\epsilon)$ with high probability at least $1-e^{-\Omega(d)}$ over the inputs.

1 Introduction

A unique game is defined in terms of a constraint graph $G=(V,E)$, a set of variables $\{x_u\}_{u\in V}$, one for each vertex u and a set of permutations (constraints) $\Pi_{uv}:[k]\rightarrow[k]$, one for each edge (u,v). An assignment to the variables is said to satisfy the constraint on the edge $(u,v)\in E$ if $\pi_{uv}(x_u)=x_v$. The edges are taken to be undirected and hence $\pi_{uv}=(\pi_{vu})^{-1}$. The goal is to assign a value from the set $[k]$ to each variable x_u so as to maximize the number of satisfied constraints.

Khot [Kho02] conjectured that it is NP-hard to distinguish between the cases when almost all the constraints of a unique game are satisfiable and when very few of the constraints are satisfiable. Formally, the statement of the conjecture is the following:

Conjecture 1 *(Unique Games Conjecture) For any constants $\epsilon,\delta>0$, for any $k>k(\epsilon,\delta)$, it is NP-hard to distinguish between instances of unique games with domain size k where at least $1-\epsilon$ fraction of constraints are satisfiable and those where at most δ fraction of constraints are satisfiable.*

The Unique Games Conjecture is known to imply optimal inapproximability results for several important problems. For instance, it implies a hardness of approximation within a factor of $2-\epsilon$ for Vertex Cover [KR03] and within a factor of 0.878 for Max-Cut [KKMO04]. These results are not known to follow from any other complexity assumptions.

Several approximation algorithms using linear and semidefinite programming have been developed for approximating unique games (see [Kho02], [Tre05], [GT06], [CMM06a], [CMM06b]). These algorithms start with an instance where the value of the SDP or LP relaxation is $1-\epsilon$ and round it to a solution with value ν. Here, value of the game refers to the maximum fraction of satisfiable constraints. For $\nu>\delta$, this would give an algorithm to distinguish between the two cases. However,

*This material is based upon work supported by the National Science Foundation under grant CCF 0515231.
†This material is based upon work supported by the National Science Foundation under grant CCF 0515231.

most of these algorithms give good approximations only when ϵ is very small ($\epsilon = O(1/\log n)$ or $\epsilon = O(1/\log k)$) [1]. For constant ϵ however, only the algorithm of [CMM06a] gives interesting parameters with $\nu \approx k^{-\epsilon/(2-\epsilon)}$. We refer the reader to [CMM06a] for a comparison of parameters of various algorithms.

Is is also known (is implicit in [KV05]) that a stronger version of the unique games conjecture, where the underlying constraint graph has significant expansion, would imply hardness of the uniform version of the Sparsest Cut problem. An algorithm for solving unique games on random graphs would thus give partial evidence of a negative answer. In this paper we study the case of random and semi-random permutations, which may help in understanding how (and if) expansion can provide an algorithmic advantage.

Our results

We study the case of random unique games generated by picking a random regular graph of degree d (or a random $G_{n,p}$ graph of average degree d) and picking a random permutation for each edge. We show that with high probability over the choice of instances, the value of the SDP from [FL92] and [Kho02] is at most δ for $d = \Omega(1/\delta^4 + 1/\epsilon^4)$. Here, we think of ϵ, δ as small constants and d as a large constant.

Using techniques from the above analysis, we also study the problem of recovering planted solutions for random unique games. Specifically, we study the model where a random instance *consistent with a given solution* is chosen to start with, and an adversary then perturbs ϵ fraction of the constraints. Thus, the given instance has one planted solution with value $1 - \epsilon$. We give an algorithm which recovers w.h.p. a solution of value at least $1 - O(\epsilon)$ *even when the perturbations are adversarial*.

To obtain both the above results, we analyze the dual of the SDP. We reduce the problem of estimating the value of the SDP to estimating the eigenvalues for an associated matrix M. Since most known eigenvalue analyses are for matrices with independent entries, which does not happen to be the case with M, we adapt the analyses from [BS87] and [AKV02] to our purposes. The planted solutions are recovered by analyzing the eigenvectors of this matrix.

Remark: It is possible to prove analogous results in the $G_{n,p}$ model by using the eigenvalue analysis from [FO05]. However, in this model our current estimates only give interesting results in the range $d = \Omega(k^2)$, where $d = pn$ is the expected degree of the constraint graph and k is the size of the alphabet.

2 Preliminaries

2.1 SDPs and duality

Semidefinite programs are often used as relaxations of 0/1 quadratic programs. In obtaining the relaxations, we often replace 0/1 variables $x_1, \ldots, x_n$ by vectors $\mathbf{v_1}, \ldots, \mathbf{v_n}$. Alternatively, we may think of solving for an $n \times n$ positive semidefinite matrix Y such that $Y_{ij} = \mathbf{v_i} \cdot \mathbf{v_j}$. Then, one way

[1] It might be good to think of k as $O(\log n)$ since this is range of interest for most reductions

of writing a general SDP is

$$\begin{aligned} \text{maximize} \quad & B \bullet Y \\ \text{subject to} \quad & A_1 \bullet Y = c_1 \\ & A_2 \bullet Y = c_2 \\ & \vdots \\ & A_n \bullet Y = c_n \\ & Y \succeq 0 \end{aligned}$$

where $A_1, A_2, \ldots A_n, B$ are symmetric square matrices and $A \bullet B$ denotes the Frobenius inner product $(= \sum_{i,j} a_{ij} b_{ij})$ of the matrices.Here $Y \succeq 0$ denotes the constraint that Y is positive semidefinite. The dual of the above SDP is

$$\begin{aligned} \text{minimize} \quad & c^T x \\ \text{subject to} \quad & x_1 A_1 + x_2 A_2 \ldots x_n A_n - B \succeq 0 \end{aligned}$$

From (weak) duality, we have that $v_{\text{primal}} \leq v_{\text{dual}}$.

2.2 Spectra of graphs

In the rest of the paper we are going to investigate this SDP by looking at its dual and reducing it to estimating eigenvalues of graphs.

We remind the reader that for a graph G, the adjacency matrix $A = A_G$ is defined as :

$$A_G = \begin{cases} 1 & \text{if } (u,v) \in E \\ 0 & \text{if } (u,v) \notin E \end{cases}$$

If the graph has n vertices, A_G has n real eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \lambda_n$. The eigenvectors that correspond to these eigenvalues form an orthonormal basis of $\mathbb{R}^n$. We note that if the graph is d-regular then the largest eigenvalue is equal to d and the corresponding eigenvector is the all-one's vector.

We can use the Courant-Fisher Theorem to characterize the spectrum of A. The largest eigenvalue satisfies

$$\lambda_1 = \max_{x \in \mathbb{R}^n} \frac{x^T A x}{x^T x}$$

If we denote the first eigenvector by x_1 then

$$\lambda_2 = \max_{x \in \mathbb{R}^n, x \perp x_1} \frac{x^T A x}{x^T x}$$

Similar definitions hold for the eigenvalues λ_i, $i \geq 3$.

3 The dual of Unique-Games SDP for random graphs

We look at the SDP for Unique-Games without the triangle inequality. The SDP is

$$\begin{array}{lll} \text{maximize} & \sum_{(u,v)\in E}\sum_{i=1}^{k} \mathbf{u}_i \cdot \mathbf{v}_{\pi_{uv}(i)} & \\ \text{subject to} & \mathbf{u}_i \cdot \mathbf{u}_j = 0 & \forall u \in V, \forall i, j \\ & \sum_{i=1}^{k} \|\mathbf{u}_i\|^2 = 1 & \forall u \in V \end{array}$$

The feasible region of the dual can be expressed as $Z \succeq 0$ where Z is an $nk \times nk$ matrix. We use Z_{uv} to denote the $k \times k$ block corresponding to the vertices u and v. The blocks are given by

$$Z_{uv} = \begin{cases} 0 & \text{if } (u,v) \notin E, u \neq v \\ -\frac{1}{2}\Pi_{uv} & \text{if } (u,v) \in E \\ Z_u & \text{if } u = v \end{cases}$$

where Π_{uv} is the permutation matrix corresponding to π_{uv} and Z_u is the (symmetric) matrix of all the variables corresponding to the vertex u. The off-diagonal entries of Z_u are $(Z_u)_{ij} = (Z_u)_{ji} = \frac{1}{2}x^u_{\{i,j\}}$ - a separate variable for each pair $\{i, j\}$ and vertex u. All the diagonal entries are the same, equal to a single variable $x^{(u)}$. The objective function of the whole SDP is $\sum_{u\in V} x^u$.

We will consider dual solutions with $x^u_{\{i,j\}} = 2d/k$ for all $u \in V$ and $i, j \in [k], i \neq j$. Also, we set $x^{(1)} = x^{(2)} = \ldots = x^{(n)} = \lambda + d/2k$. Here λ is taken to be an upper bound on the second eigenvalue. Note that the first eigenvalue of M is d since M can be thought of as the adjacency matrix of a d-regular graph on nk vertices. The objective value as $nd/2k + n\lambda$. Putting in these values for the variables, we will need to show that the following equation is satisfied.

$$\lambda I + \frac{d}{2k}J - \frac{1}{2}M \succeq 0$$

where I is the $nk \times nk$ identity matrix, J is a block diagonal matrix with $k \times k$ blocks of all 1s on the diagonal and M is a block matrix with $M_{uv} = \Pi_{uv}$ if $(u,v) \in E$ and 0 otherwise.

Let z denote the all vector with all coordinates $\frac{1}{\sqrt{nk}}$. Then z is the first eigenvector of M. We prove the following in the next section

Theorem 2 *Let M be a matrix generated according to a random d-regular graph and random permutations on each edge. Then, with probability $1 - e^{-\Omega(d)}$, $\lambda_2(M) \leq Cd^{3/4}$*

Hence, we take $\lambda = Cd^{3/4}$ which is a bound on the second eigenvalue[2]. Note that z is the first eigenvector of both J and M. Since we can express any vector x and $\alpha z + \beta w$ with $w \perp z$, we have

$$\begin{aligned} x^T\left(\lambda I + \frac{d}{2k}J - \frac{1}{2}M\right)x &= (\alpha z + \beta w)^T\left(\lambda I + \frac{d}{2k}J - \frac{1}{2}M\right)(\alpha z + \beta w) \\ &= \lambda + \alpha^2\frac{d}{2} + \beta^2\frac{d}{2}w^T J w - \frac{1}{2}\left(\alpha^2 z^T M z + \beta^2 w^T M w\right) \end{aligned}$$

Since J is positive semidefinite, $z^T M z \leq d$ and $w^T M w \leq Cd^{3/4}$, we have $x^T(\lambda I + \frac{\lambda_1}{2k}J - \frac{1}{2}M) \geq 0$ for every x. This gives that the value of the SDP for random d-regular graphs is $\frac{|E|}{k} + \frac{|E|}{d^{1/4}}$ with high probability.

[2] We believe that it is possible to improve this bound to even $C\sqrt{d}$ but this is not very important for our purposes.

4 Bounding the second eigenvalue for d-regular graphs

We consider undirected random $2d$-regular graphs G_{2d} on n vertices constructed by choosing d permutations (over n elements) independently at random. For each of the chosen permutations σ and for each vertex u we add to the graph the edge $(u, \sigma(u))$. The unique game is then constructed for by then picking a random permutation π_{uv} (over k elements) for each edge $(u, v) \in E$.

The bound on the second eigenvalue is obtained in two steps. We first by first bound the expected value by examining the trace of a power of the matrix M. We then show a concentration bound using an application of Talagrand's inequality adapted from [AKV02].

4.1 Bounding the mean

In the following argument, it will be convenient to consider the normalized matrices $M^* = (2d)^{-1}M$, $A^* = (2d)^{-1}A$. For any positive integer p, we have $Trace((M^*)^p) = \frac{1}{(2d)^p} Trace(M^p)$ and same for A^*. Let $\rho_1, \rho_2, \cdots, \rho_{nk}$ the eigenvalues of M^* in order of decreasing value. Clearly, $\rho_1 = 1$. Our next goal is to upper-bound the mean value of the quantity $\rho = max\{\rho_2, |\rho_n|\}$. Let p be a large positive integer to be fixed later.

Lemma 3

$$E[\rho] \leq (E[Trace((M^*)^{2p})] - 1)^{1/2p}$$

PROOF: Because Trace$((M^*)^{2p}) = \sum_{1 \leq i \leq nk} {\rho_i}^{2p}$ and because all the eigenvalues of a symmetric matrix are real, we have :

$$\rho^{2p} \leq Trace((M^*)^{2p}) - 1$$

Taking expectations over the probability space described above,(that is, over all $2d$-regular graphs and over all permutations of k elements within each non-zero block), we have

$$E[\rho] \leq E[\rho^{2p}]^{1/(2p)} \leq (E[Trace((M^*)^{2p})] - 1)^{1/2p}$$

by Jensen's inequality. $\square$

We next relate the value of $E[Trace((M^*)^{2p})]$ to $E[Trace((A^*)^{2p})]$.

Claim 4 *Let $A = [a_{ij}]$ be the adjacency matrix of a graph G and M be a block matrix with $M_{uv} = \Pi_{uv}$ if $(u, v) \in E$ and 0 otherwise. Then $E[Trace(M^{2p})] = Trace(A^{2p})$ where p is a positive integer and the expectation on the left hand side is taken over the choice of permutations.*

PROOF: Let S be a set containing all the sequences of $2p+1$ nodes of G that begin and end at the same node. I.e $S = \{uu_1 \cdots u_{2p}u\}$. Each $s \in S$ corresponds to a walk on G of length $2p$ that begins and ends at the same node and therefore also corresponds to a sequence of blocks of the matrix M above that begins and ends at the same block.

For any matrix $Q = [q_{ij}]$ and for any positive integer n we have

$$Trace(Q^n) = \sum_{i_1, i_2, \ldots i_n} q_{i_1 i_2} q_{i_2 i_3} \cdots q_{i_n i_1}$$

Observe that when Q is the adjacency matrix of a graph, each term in the above sum is 1 if $i_1, i_2, \cdots i_n, i_1$ is a path in the graph and 0 otherwise.

Thus, for the matrices A and M we have

$$\begin{aligned} Trace(A^{2p}) &= \sum_{u_1u_2\ldots u_{2p}u_1 \in S} a_{u_1u_2}\ldots a_{u_{2p}u_1} \\ Trace(M^{2p}) &= \sum_{\substack{u_1,u_2,\ldots,u_{2p}u_1 \in S \\ i_1,i_2,\ldots,i_{2p}\in[k]}} m_{(u_1,i_1)(u_2,i_2)}\ldots m_{(u_{2p},i_{2p})(u_1,i_1)} \end{aligned}$$

where the tuple (u, i) corresponds to the index of the ith element of block u.

We can write each term $m_{(u,i)(v,j)} = a_{uv}\ldots \mathbb{I}_{\{\pi_{uv}(i)=j\}}$, where the random variable $\mathbb{I}_{\{\pi_{uv}(i)=j\}}$ is 1 when $\pi_{uv}(i) = j$ and 0 otherwise. We can now re-write the trace as

$$Trace(M^{2p}) = \sum_{u_1u_2\ldots u_{2p}u_1\in S} a_{u_1u_2}\ldots a_{u_{2p}u_1} \sum_{i_1,i_2,\ldots,i_{2p}\in[k]} \mathbb{I}_{\{\pi_{u_1u_2}(i_1)=i_2\}}\ldots \mathbb{I}_{\{\pi_{u_{2p}u_1}(i_{2p})=i_1\}}$$

and, taking expectation over all permutations

$$E[Trace(M^{2p})] = \sum_{u_1u_2\ldots u_{2p}u_1\in S} a_{u_1u_2}\ldots a_{u_{2p}u_1} \sum_{i_1,i_2,\ldots,i_{2p}\in[k]} P[\pi_{u_1u_2}(i_1) = i_2 \wedge \ldots \wedge \pi_{u_{2p}u_1}(i_{2p}) = i_1]$$

For multi-indices $U = u_1u_2\ldots u_{2p}$ and $I = i_1i_2\ldots i_{2p}$ let $E_{U,I}$ be the event $\{\pi_{u_1u_2}(i_1) = i_2 \wedge \cdots \wedge \pi_{u_{2p}u_1}(i_{2p}) = i_1\}$. For a fixed U, the events $E_{U,I}$ where I takes all possible values consist of a partition of the whole probability space. Therefore with this notation,

$$E[Trace(M^{2p})] = \sum_U a_{u_1u_2}\ldots a_{u_{2p}u_1} \sum_I P[E_{U,I}] = \sum_U a_{u_1u_2}\ldots a_{u_{2p}u_1} = Trace(A^{2p})$$

□

Hence, to bound ρ, it suffices to bound $E[Trace(A^{2p})]$. The following lemma can be found in [BS87].

Lemma 5 *Let A^* as above and $p = (2 - \epsilon')log_{d/2}n$ a positive integer. Then*

$$E[Trace((A^*)^{2p})] \le \frac{1}{n^{1-\epsilon'}} + 1 + O(\frac{(logn)^4}{n})$$

Claim 6 *Let p be as above. Then for every $\epsilon > 0$ we have the inequality :*

$$E[\rho] \le (\frac{2}{d})^{1/4}(1 + \epsilon + o(1))$$

Proof: From claim 4 we have

$$E[Trace((M^*)^{2p})] = E[Trace((A^*)^{2p})]$$

Using lemma 5 we have

$$E[Trace((M^*)^{2p})] \le (\frac{1}{n^{1-\epsilon'}} + 1 + O(\frac{(logn)^4}{n}))$$

Hence,

$$E[\rho] \le (E[Trace((M^*)^{2p})] - 1)^{1/2p} = (E[Trace((A^*)^{2p})] - 1)^{1/(2(2-\epsilon')log_{d/2}n)}$$

$$\leq (\frac{1}{n^{1-\epsilon'}})^{\frac{1}{2(2-\epsilon')log_{d/2}n}}(1+o(1)) = (\frac{2}{d})^{1/4}(1+\epsilon+o(1))$$

Which follows by the appropriate choice of ϵ'. □

From the above calculations it follows that if λ is the second largest (in absolute value) eigenvalue of M, then

$$E[\lambda] = O(d^{3/4})$$

We note that it is also possible to bound $E[\lambda]$ by $O(\sqrt{d})$ by using the (more involved) bound on $Trace((A^*)^{2p})$ from [Fri91].

4.2 Concentration of λ around the mean

We will next prove that with probability that tends to 1 as $n \to \infty$, λ deviates from its mean by at most $\sqrt{d}$. For that we will first prove concentration of λ around its median, and then use elementary probability techniques to show that the expectation and the median of λ are very close. Namely, we will prove the following theorem :

Theorem 7 *The probability that λ_2 deviates from its median by more than t is at most $4e^{-t^2/128}$. The same estimate holds for the probability that λ_{kn} deviates from its median by more than t. Therefore $Pr[|\lambda - \mu(\lambda)| \geq t] \leq 2e^{-t^2/128}$, where $\mu(\lambda)$ denotes the median of λ.*

For that reason, we will use Talagrand's inequality in a similar manner as in [AKV02].

Theorem 8 *(Talagrand's Inequality) Let $\Omega_1, \Omega_2, \cdots, \Omega_m$ be probability spaces, and let Ω denote their product space. Let $\mathcal{A}$ and $\mathcal{B}$ be two subsets of Ω and suppose that for each $B = (B_1, \cdots, B_m) \in \mathcal{B}$ there is a real vector $\alpha = (\alpha_1, \alpha_2, \cdots, \alpha_m)$ such that for every $A = (A_1, \cdots, A_m) \in \mathcal{A}$ the inequality*

$$\sum_{i:A_i \neq B_i} \alpha_i \geq t(\sum_{i=1}^{m} {\alpha_i}^2)^{1/2}$$

holds. Then

$$Pr[\mathcal{A}]Pr[\mathcal{B}] \leq e^{-t^2/4}$$

We now apply Talagrand's inequality to prove theorem 7. We will show the case for λ_2, but the same proof easily carries out for λ_{kn}. Some notation follows:
Let $\binom{m=n+1}{2}$ and consider the product space Ω of the blocks M_{ij},$1 \leq i, j, \leq n$ where each block is a $k \times k$ permutation matrix. We identify each element of Ω with the vector consisting of the corresponding m $k \times k$ blocks. Instead of i, j we will use indices u, v for the block of M corresponding to vertices u, v. Let μ denote the median of λ_2.
Let $\mathcal{A} = \{M|\lambda_2(M) \leq \mu\}$ and $\mathcal{B} = \{M|\lambda_2(M) \geq \mu + t\}$. By definition of the median, $Pr[\mathcal{A}] \geq 1/2$.

For any vector $f = (f(1), \cdots, f(nk)) \in \mathbb{R}^{nk}$ we will denote by $f_i \in \mathbb{R}^k$,$1 \leq i \leq n$ the vector that corresponds to the i-th block of k coordinates of f, i.e. $f_i = (f((i-1)k), f((i-1)k+1), \cdots, f(ik))$. Let $\|f\|$ be the euclidean norm of f.

PROOF:(Of theorem 7) Fix a vector $B \in \mathcal{B}$. Let $f^{(1)}, f^{(2)}$ denote the first and second unit eigenvector of B. We define the following cost vector $\alpha = (a_{uv})$ for B.

$$\alpha_{uu} = (\|f_u^{(1)}\| + \|f_u^{(2)}\|)$$

$$\alpha_{uv} = \sqrt{2\alpha_{uu}\alpha_{vv}}, v \neq u$$

Let $D = \{(u,v)|A_{uv} \neq B_{uv}\}$. We will show that

$$\sum_{(u,v)\in D} \alpha_{uv} \geq c \cdot t \cdot (\sum_{1\leq u\leq v\leq n} \alpha_{uv}^2)^{1/2}$$

Note that

$$\sum_{1\leq u\leq v\leq n} \alpha_{uv}^2 = (\sum \alpha_{uu})(\sum \alpha_{vv}) = (\|f^{(1)}\|^2 + \|f^{(2)}\|^2)^2 = 4$$

Let $z = c_1 f^{(1)} + c_2 f^{(2)}$ be a unit vector (i.e. ${c_1}^2 + {c_2}^2 = 1$) which is perpendicular to the first eigenvector of A. Note that such a vector can always be found, since the orthogonality of $f^{(1)}$ and $f^{(2)}$ implies that the subspace $span\{f^{(1)}, f^{(2)}\}$ is 2-dimensional. Then

$$z^T Az \leq \lambda_2(A) \leq \mu$$

and

$$z^T Az \geq \lambda_2(B) \geq \mu + t$$

which implies

$$\begin{aligned} t \leq z^T(B-A)z \leq \sum_{(u,v)\in D} {z_u}^T(B_{uv} - A_{uv})z_v &\leq \sum_{(u,v)\in D:(B_{uv}-A_{uv})_{ij}\neq 0} |z_{ui}||z_{vj}| \\ &\leq \sum_{(u,v)\in D} \sqrt{2\|z_u\|^2}\sqrt{2\|z_v\|^2} \\ &\leq \sum_{(u,v)\in D} 2\sqrt{\alpha_{uu}}\sqrt{\alpha_{vv}} = \sqrt{2} \sum_{(u,v)\in D} \alpha_{uv} \end{aligned}$$

The fourth inequality holds because each coordinate appears at most twice (each block is a permutation matrix). By combining the above, we obtain

$$\sum_{(u,v)\in D} \alpha_{uv} \geq \frac{t}{4\sqrt{2}} (\sum {\alpha_{vv}}^2)^{1/2} \quad \Rightarrow \quad Pr[B] \leq 2e^{\frac{-t^2}{128}}$$

□

We conclude by showing that the eigenvalues are also concentrated around their expectation. Namely,

Theorem 9 $Pr[|\lambda - E[\lambda]| \geq t] \leq e^{-(1-o(1))t^2/128}$

To prove this, we show that the expectation and the median of eigenvalues are very close. We show the result for λ_2 but the result holds for all eigenvalues (with different constants in the exponent).

Claim 10 $E[\lambda_2] - \mu \leq 8\sqrt{2\pi}$

PROOF:

$$E[\lambda_2] - \mu \leq E[|\lambda_2 - \mu|] = \int_0^\infty P[|\lambda_2 - \mu| > t]dt \leq \int_0^\infty 2e^{\frac{-t^2}{128}} dt = 8\sqrt{2\pi}$$

□

5 Recovering planted solutions

In this section, we consider random graphs with a planted solution satisfying most of the constraints. Consider the following model for Unique Games: we build an undirected random $d = 2s$-regular graph on n vertices by choosing independently s permutations uniformly at random among all possible permutations of n elements. For each of the chosen permutations π and for each vertex i we add to the graph the edge (i, π_i).

Let G be a random $d-$regular graph drawn according to the previous distribution and let k be the alphabet size. We construct a random satisfiable instance of Unique Games as follows:

- Pick a random d-regular graph G.
- For each node $u \in V$ pick a number n_u from 1 to k. Let $S_{uv} = \{\pi \in S_k : \pi(A(u)) = A(v)\}$ the set of constraints that are satisfied by the assignment $A(u)$ for each node u.
- For each edge $(u, v) \in E$ pick a permutation Π_{uv} uniformly from S_{uv}. We denote such a game with the triplet (G, k, M) where M is a block matrix with $M_{uv} = \Pi_{uv}$ if $(u, v) \in E$ and 0 otherwise (it is the matrix that appeared in the dual SDP). We will also use the notation M_k to emphasize the alphabet size.
- Now, let the adversary pick any $\epsilon|E|$ and perturb them by modifying the corresponding permutations as he wishes. Clearly, the previous assignment still satisfies at least $(1 - \epsilon)$ fraction of the edges. Let $(G, k, \hat{M})$ be the perturbed game.

We address the following question about the above model:

Question 11 *Given an instance of* $(G, k, \hat{M})$ *is it possible to reconstruct (w.h.p.) an assignment that satisfies* $(1 - \epsilon_1)$ *fraction of the edges for some* $\epsilon_1 = \Theta(\epsilon)$.

We exhibit an efficient algorithm that given such an instance of $(G, k, \hat{M})$ produces with high probability over the choice of the instance, an assignment that satisfies $(1-\epsilon_1)$ fraction of the edges for some $\epsilon_1 = \Theta(\epsilon)$.

5.1 The Algorithm

The algorithm can be simply described as follows :

Let M' be the perturbed matrix as above and let x be the second eigenvector of that matrix. This is a vector with n blocks of size k each, one block for each node and one entry within each block corresponding to one of the possible k assignments to that node. Denote with x_u the k-dimensional sub-vector that corresponds to node u. The algorithm is simply the following:

- Create an assignment A_1 by assigning to each node u a number $A(u)$ such that $x_{uA(u)} \geq x_{ui}$ for all $i = 1, \ldots, k$. In case of equality, break ties arbitrarily.
- Create an assignment A_2 by assigning to each node u a number $A(u)'$ such that $x_{uA(u)} \leq x_{ui}$ for all $i = 1, \ldots, k$. In case of equality, break ties arbitrarily.
- Choose the better of the two assignments A_1 and A_2.

Note that we need to create two assignments only because the negative of the eigenvector (which hopefully encodes the assignment) is also an eigenvector with the same eigenvalue. In the proof it will be sufficient to analyze only one of the above cases.

5.2 The Proof

The proof relies on two basic techniques: spectral analysis of matrix M and combination of results about matrix perturbation of matrices.

5.2.1 Spectral analysis of M

Since G is a d-regular graph and each block of M is a permutation matrix, the first eigenvector of M (with eigenvalue d) is the vector z with $z_{ui} = \frac{1}{\sqrt{nk}}$ for all (u, i). It is easy to verify that the following vector w is perpendicular to z and also has eigenvalue d.

$$w_{ui} = \begin{cases} \frac{k-1}{\sqrt{nk(k-1)}} & \text{if } i = A(u) \\ \\ \frac{-1}{\sqrt{nk(k-1)}} & \text{otherwise} \end{cases}$$

The following claim shows that w.h.p. all other eigenvalues of the matrix M are small and hence w is the only vector perpendicular to z with eigenvalue d. Hence the algorithm will at least recover the correct solution for the game (G, k, M)

Claim 12 *With high probability over the choice of M, $\lambda_i(M) \leq O(\sqrt{d})$ for all $i \geq 3$.*

Proof: Let y be a vector perpendicular to both z and w such that $\|y\| = 1$. Then, we must have that

$$\sum_u \sum_i y_{ui} = 0 \quad \text{and} \quad \sum_u \left((k-1) y_{uA(u)} - \sum_{i \neq A(u)} y_{ui} \right) = 0$$

which implies

$$\sum_u y_{uA(u)} = \sum_{i \neq A(u)} y_{ui} = 0$$

We now define y_1 as $(y_1)_{ui} = y_{ui}$ for all $i \neq A(u)$ and $(y_1)_{uA(u)} = 0$. Also, let $y_2 = y - y_1$. Then for every u, $(y_2)_{uA(u)} = y_{uA(u)}$ is the only non-zero coordinate of y_2. Also $\|y_1\|, \|y_2\| \leq 1$. We have,

$$\|My\| = \|M(y_1 + y_2)\| = \|My_1 + My_2\| \leq \|My_1\| + \|My_2\|$$

However, since all constraints are satisfied by the assignment $x_u = A(u)$, $\|My_2\| = \left\|Ay_2^G\right\|$, where y_2^G is an n-dimensional "projection" of y_2 on the graph by setting $(y_2^G)_u = (y_2)_u$, and A is the adjacency matrix of the graph. From the above equations we have that $\sum_u y_{uA(u)} = 0$, which means that y_2^G is perpendicular to the first eigenvector of A. Thus, w.h.p.

$$\|My_2\| = \left\|Ay_2^G\right\| \leq O(\sqrt{d}\left\|y_2^G\right\|) \leq O(\sqrt{d})$$

We now consider a new game with matrix M_{k-1} with alphabet size $k-1$ obtained by deleting the value $A(u)$ for each u. Note that this is a completely random unique game for alphabet size $k-1$, since we chose constraints for M randomly after fixing $\pi_{uv}(A(u)) = A(v)$. Finally, it remains to notice that $\|My_1\| = \left\|M_{k-1}y_1^{(k-1)}\right\|$, where $y_1^{(k-1)}$ is the $n(k-1)$-dimensional projection of y_1 obtained by deleting coordinates $y_{uA(u)}$ for all u. We also have

$$\sum_{u,i} (y_1^{(k-1)})_{ui} = \sum_{u,i \neq A(u)} y_{ui} = 0$$

which gives that $y_1^{(k-1)}$ is perpendicular to the first eigenvector of M_{k-1} and hence by the previous eigenvalue estimates,

$$\|My_1\| = \left\|M_{k-1}y_1^{(k-1)}\right\| \leq O(\sqrt{d})$$

□

5.2.2 Perturbation of Matrices

We now show that when we perturb the game as described, the new matrix $\hat{M}$ has second eigenvector close to w therefore we can use the same algorithm and recover the original solution. Note that M and $\hat{M}$ have the same underlying constraint graph. Let A be the adjacency matrix of this graph. With high probability, it is true that $\lambda_2(A) = O(\sqrt{d})$. We assume this condition in the rest of this section.

Let $\hat{w}$ be the second eigenvector of matrix $\hat{M}$. We assume that we normalize our eigenvectors so that they have norm 1. Also, let $\hat{\lambda}_2 = \lambda_2(\hat{M})$, $\lambda_2 = \lambda_2(M)$ and $\Delta = \min_{i \geq 3} |\lambda_i(M) - \hat{\lambda}_2|$. Since z is the first eigenvector of both M and $\hat{M}$, we can express $\hat{w}$ as $\hat{w} = \alpha w + \beta w_\perp$ where $w_\perp$ is perpendicular to both z and w.

We first relate the difference in the old and new eigenvectors to the perturbation $M - \hat{M}$. The next claim appears in [DK70] as the $\sin\theta$ theorem. We provide a proof here for self-containment.

Claim 13 $|\beta| \leq \left\|(M - \hat{M})\hat{w}\right\| / \Delta$

Proof: We have

$$(M - \hat{M})\hat{w} \;=\; \alpha\lambda_2 w + \beta M w_\perp - \hat{\lambda}_2\hat{w} \;=\; \alpha(\lambda_2 - \hat{\lambda}_2)w + \beta(Mw_\perp - \hat{\lambda}_2 w_\perp)$$

Since $Mw_\perp$ is in the eigenspace perpendicular to both z and w, we have

$$\left\|(M - \hat{M})\hat{w}\right\|^2 = \alpha^2(\lambda_2 - \hat{\lambda}_2)^2 \|w\|^2 + \beta^2 \left\|Mw_\perp - \hat{\lambda}_2 w_\perp\right\|^2 \geq \beta^2 \left\|Mw_\perp - \hat{\lambda}_2 w_\perp\right\|^2$$

However, $\left\|Mw_\perp - \hat{\lambda}_2 w_\perp\right\| \geq \Delta$ which proves the claim. □

Hence, to prove that the second eigenvector of the perturbed matrix does not change by much, we simply need to bound $\left\|(M - \hat{M})\hat{w}\right\|$. We shall also need the fact that $\hat{w}$ is somewhat "uniform" over each block. To formalize this, let $\bar{w}$ be the n-dimensional vector such that $\bar{w}_u = \|\hat{w}_u\|$ where $\hat{w}_u$ is the k-dimensional vector $(\hat{w}_{u1}, \ldots, \hat{w}_{uk})^T$. We then show that $\bar{w}$ is very close to the all-one's vector.

Claim 14 $\left\|\bar{w} - \frac{1}{\sqrt{n}}\mathbf{1}\right\| \leq 3\sqrt{\epsilon}$

Proof: It is easy to verify that for the perturbed matrix $\hat{M}$ and the vector w, which is the second eigenvector of M (as described in the previous section), we must have

$$w^T \hat{M} w \;\geq\; (1-\epsilon)d + \epsilon n d \left(\frac{-2(k-1)+(k-2)}{nk(k-1)}\right) \;=\; \left(1 - \frac{k\epsilon}{k-1}\right) d \;\geq\; (1-2\epsilon)d$$

Since $w \perp z$, we get that the second eigenvalue of $\hat{M}$ is at least $(1-2\epsilon)d$. However, this gives that $\bar{w}^T A w$ must be large, where A is the adjacency matrix of the underlying constraint graph. To see this, observe that

$$\hat{w}^T \hat{M} \hat{w} \;=\; \sum_{u,v} \hat{w}_u^T M_{uv} \hat{w}_v \;\leq\; \sum_{u,v} \|\hat{w}_u\| \|M_{uv}\hat{w}_v\| \;=\; \sum_{u,v} A_{uv} \|\hat{w}_u\| \|\hat{w}_v\| \;=\; \sum_{u,v} \bar{w}_u A_{uv} \bar{w}_v$$

which follows from the fact that M_{uv} is either a permutation matrix which does not change the norm of $\hat{w}_v$, or is the 0 matrix. Taking $\bar{w} = \gamma\left(\frac{\mathbf{1}}{\sqrt{n}}\right) + \delta \mathbf{1}_\perp$, we get that $\left\|\bar{w}^T A w\right\| \leq \gamma d + \delta \cdot C\sqrt{d}$. Combining this with the previous equations, we get that

$$\gamma d + \delta \cdot C\sqrt{d} \geq (1-2\epsilon)d$$

which gives $\gamma \geq 1 - 3\epsilon$ for $d = \Omega(1/\epsilon^2)$. This implies $\left\|\bar{w} - \frac{1}{\sqrt{n}}\mathbf{1}\right\| \leq \sqrt{6\epsilon} \leq 3\sqrt{\epsilon}$. □

Lemma 15 *If we express $\hat{w}$ as $\hat{w} = \alpha w + \beta w_\perp$ then for all perturbations of $\epsilon|E|$ edges, $|\beta| \leq 9\sqrt{\epsilon}$ and $|\alpha| \geq 1 - 9\sqrt{\epsilon}$*

Proof: Define the $n \times n$ matrix R as $R_{uv} = 1$ when the block $(M - \hat{M})_{uv}$ has any non-zero entries, and $R_{uv} = 0$ otherwise. Note that if $(M - \hat{M})_{uv}$ is non-zero, then it must be the difference of two permutation matrices. Thus, for all v $\left\|(M - \hat{M})_{uv}\hat{w}_v\right\| \leq 2R_{uv}\|\hat{w}_v\|$. We have that

$$\begin{aligned}
\left\|(M - \hat{M})\hat{w}\right\| \;&=\; \sqrt{\sum_u \left\|\sum_v (M - \hat{M})_{uv}\hat{w}_v\right\|^2} \;\leq\; \sqrt{\sum_u \left(\sum_v \left\|(M - \hat{M})_{uv}\hat{w}_v\right\|\right)^2} \\
&\leq\; \sqrt{\sum_u \left(\sum_v 2R_{uv}\|\hat{w}_v\|\right)^2} \\
&\leq\; 2\|R\bar{w}\|
\end{aligned}$$

To estimate $\|R\bar{w}\|$, we break it up as

$$\|R\bar{w}\| \leq \frac{1}{\sqrt{n}} \|R \cdot \mathbf{1}\| + \left\| R \cdot \left(\bar{w} - \frac{1}{\sqrt{n}} \mathbf{1} \right) \right\|$$

Since $\left\| \bar{w} - \frac{1}{\sqrt{n}} \mathbf{1} \right\| \leq 3\sqrt{\epsilon}$ and R has at most d 1s in any row, $\left\| R \cdot \left(\bar{w} - \frac{1}{\sqrt{n}} \mathbf{1} \right) \right\| \leq 3\sqrt{\epsilon} d$. Also, $\|R \cdot \mathbf{1}\| = \sqrt{\sum_u \left(\sum_v R_{uv} \right)^2}$. Since R has a total of $\epsilon n d$ 1s, this expression is maximized when it has d 1s in ϵn rows. This gives $\frac{1}{\sqrt{n}} \|R \cdot \mathbf{1}\| \leq \sqrt{\epsilon} d$. Combining with the above, we have that $\left\| (M - \hat{M}) \hat{w} \right\| \leq 8\sqrt{\epsilon}$.

For estimating Δ, note that $\hat{\lambda}$ is at least $(1 - 2\epsilon)d$ as witnessed by the vector w. Hence $\Delta \geq (1 - 2\epsilon)d - O(\sqrt{d})$. This gives $|\beta| \leq 9\sqrt{\epsilon}$. and $\alpha = \sqrt{1 - \beta^2} \geq 1 - 9\sqrt{\epsilon}$. □

Lemma 16 *For ϵ small enough, for at least $(1 - 199\epsilon)n$ blocks u the coordinate $w_{uA(u)}$ has the maximum value within its block.*

Proof: Within each block u, in order for coordinate $A(u)$ to be no longer the maximum one, it must happen that for some j

$$\alpha \frac{k-1}{\sqrt{nk(k-1)}} + \beta \cdot (w_\perp)_{uA(u)} \leq -\frac{\alpha}{\sqrt{nk(k-1)}} + \beta \cdot (w_\perp)_{uj}$$

This gives

$$\begin{aligned} & (w_\perp)_{uj} - (w_\perp)_{uA(u)} \geq \frac{k}{\sqrt{nk(k-1)}} \cdot \frac{\alpha}{\beta} \\ \Rightarrow \quad & [(w_\perp)_{uj}]^2 + [(w_\perp)_{uA(u)}]^2 \geq \frac{1}{2} [(w_\perp)_{uj} - (w_\perp)_{uA(u)}]^2 \geq \frac{k}{2n(k-1)} \cdot \frac{\alpha^2}{\beta^2} \\ \Rightarrow \quad & \|(w_\perp)_u\|^2 \geq \frac{k}{2n(k-1)} \cdot \frac{(1 - 9\sqrt{\epsilon})^2}{81\epsilon} \end{aligned}$$

.

We call such a block "bad". Assume that there are ηn bad blocks. Then

$$1 \geq \sum_{\text{bad } u} \|(w_\perp)_u\|^2 \geq \eta n \cdot \frac{k}{2n(k-1)} \cdot \frac{(1 - 9\sqrt{\epsilon})^2}{81\epsilon} \quad \Rightarrow \quad \eta \leq \frac{162\epsilon}{(1 - 9\sqrt{\epsilon})^2} \leq 199\epsilon$$

□

Therefore, for all but at most 199ϵ fraction of the blocks, the maximum coordinate remains at the same place. The assigment recovered by our algorithm then fails to satisfy at most $199\epsilon n d$ constraints corresponding to these blocks and $\epsilon n d$ constraints perturbed initially. Thus, the solution violates at most $200\epsilon n d = 400\epsilon|E|$ constraints and has value at least $1 - 400\epsilon$

Remark: We note that the only two properties of random instances used in the above analysis are $\lambda_1(A) \gg \lambda_2(A)$ and $\lambda_2(M) \gg \lambda_3(M)$. Thus, as long as the constraint graph has a large eigenvalue gap and the eigenvalues of M differ significantly, our algorithm succeeds in recovering planted solutions.

6 Expanding instances of Γ-max-lin

We now apply the techniques from the pevious section to solve special cases of unique games when the constraints are in the form of difference equations in a group Γ and the underlying constraint graph G has significant expansion. The special case of difference equations in a group, called Γ-max-lin was first introduced by [KKMO04] and is believed to be equally hard as the general case. We give an algorithm for solving this special case when the underlying constraint graph has good expansion.

As in the previous analysis, we assume that the graph is d-regular with second eigenvalue at most $d(1-\gamma)$. Our algorithm distinguishes instances of Γ-max-lin in which $(1-\epsilon)$ fraction of the constraints are satisfiable from those in which at most δ fraction of the cosntraints are satisfiable for $\gamma = \Omega(\epsilon^{1/3})$.

In the case of Γ-max-lin, for each edge (u,v) in the graph G, we have a constraint of the form $x_u - x_v = c_{uv}$, where x_u, x_v are variables taking values in Γ and $c_{uv} \in \Gamma$. Let k denote the size of the group Γ. As before, we consider a matrix $\hat{M}$ for the given instance, and think of it as an adversarial perturbation on ϵ-fraction of the edges of another matrix M corresponding to a fully satisfiable instance. Let A be an assignment such that the values $x_u = A(u)$ satisfy all the constraints in the instance corresponding to M.

6.1 Spectral analysis of M

For the matrix M, we define the eigenvectors $y^{(0)}, \ldots, y^{(k-1)}$ as

$$y^{(s)}_{ui} = \begin{cases} \frac{1}{\sqrt{n}} & \text{if} \quad i = A(u) + s \mod k \\ 0 & \text{otherwise} \end{cases}$$

Note that for Γ-max-lin, if $\forall u: \; x_u = A(u)$ is a satisfying assignment, then so is $\forall u: \; x_u = A(u)+s$. Hence, the vectors $y^{(0)}, \ldots, y^{(k-1)}$ correspond to satisfying assignments and are eigenvectors with eigenvalue d for the matrix M. We now show that any eigenvector which is orthogonal to all these vectors has eigenvalue at most $d(1-\gamma)$. Let x be a vector such that $x \perp y^{(s)} \; \forall s$. Then, we have

$$\forall s \in \{0, \ldots, k-1\} \quad \sum_u x_{uA(u)+s} = 0$$

We then decompose x into $x^0, \ldots, x^{(k-1)}$, where

$$x^{(s)}_{ui} = \begin{cases} x_{ui} & \text{if} \quad i = A(u) + s \mod k \\ 0 & \text{otherwise} \end{cases}$$

It is immediate from the definition that $x = \sum_s x^{(s)}$ and that $\|x\|^2 = \sum_s \left\|x^{(s)}\right\|^2$. To bound the eigenvalue corresponding to x, note that

$$x^T M x = \sum_{s,t} (x^{(s)})^T M x^{(t)}$$

Let e_i denote the ith unit vector in k-dimensions. We can then write $x^{(s)}_u$ as $x_{uA(u)+s} e_{A(u)+s}$. Using this notation, we compute the terms in the above equation as

$$(x^{(s)})^T M x^{(t)} = \sum_{(u,v)\in E} (x^{(s)}_u)^T \Pi_{uv} (x^{(t)}_v) = \sum_{(u,v)\in E} x_{uA(u)+s} x_{vA(v)+t} \cdot (e_{A(u)+s})^T \Pi_{uv} e_{A(v)+t}$$

Since the permutation maps $A(u)$ to $A(v)$ and $A(u)+s$ to $A(v)+s$ for all s, $(x^{(s)})^T M x^{(t)} = 0$ when $s \neq t$. For the rest of the terms, we have

$$(x^{(s)})^T M x^{(s)} \;=\; \sum_{(u,v)\in E} x_{uA(u)+s} x_{vA(v)+s} \;\leq\; d(1-\gamma)\left\|x^{(s)}\right\|^2 \qquad \text{(Since } \sum_u x_{uA(u)+s} = 0\text{)}$$

Hence,

$$x^T M x \;=\; \sum_s (x^{(s)})^T M x^{(s)} \;\leq\; d(1-\gamma)\sum_s \left\|x^{(s)}\right\|^2 \;=\; d(1-\gamma)\left\|x\right\|^2$$

6.2 Effect of perturbation on the eigenspace

We denote the matrix of the pertubed game, in which $(1-\epsilon)$ fraction of the constraints are satisfiable, by $\hat{M}$. Note that for all $s \in \{0,\ldots,k-1\}$, we have $(y^{(s)})^T \hat{M} y^{(s)} \geq d(1-\epsilon)$. Let $\hat{w}$ be any unit-length eigenvector of $\hat{M}$, with eigenvalue at least $(1-\epsilon)d$. We can express $\hat{w}$ as $\sum_s \alpha_s y^{(s)} + \beta y_\perp$. We will show that $|\beta| = O\left(\sqrt{\frac{2\epsilon}{\gamma^3}}\right)$, i.e. $\hat{w}$ is very close to some vector in the span of $y^{(0)},\ldots,y^{(k-1)}$. Note that this also implies that the eigenspace of vectors with eigenvalue greater than $(1-\epsilon)d$ has dimension at most k (otherwise we would find a vector othogonal to $y^{(0)},\ldots,y^{(k-1)}$ which cannot be close to their span).

By claim 13 we note that

$$|\beta| \;\leq\; \frac{\left\|(M-\hat{M})\hat{w}\right\|}{(1-\epsilon)d-(1-\gamma)d} \;=\; \frac{\left\|(M-\hat{M})\hat{w}\right\|}{(\gamma-\epsilon)d}$$

Also, let $\bar{w}$ be the n-dimensional vector such that $\bar{w}_u = \|\hat{w}_u\|$. Since, $\hat{w}$ corresponds to a large eigenvalue, we have that

$$(1-\epsilon)d \;\leq\; (\hat{w})^T \hat{M}\hat{w} \;\leq\; \sum_{u,v} \|\hat{w}_u\| A_{uv} \|\hat{w}_v\| \;=\; (\bar{w})^T A \bar{w}$$

We now show that $\bar{w}$ is very close to the all one's vector. Specifically, writing $\bar{w}$ as $\frac{a}{\sqrt{n}}\mathbf{1} + b\mathbf{1}_\perp$, we get

$$\begin{aligned} (\bar{w})^T A \bar{w} \;&\leq\; a^2 d + b^2(1-\gamma)d \\ \Rightarrow (1-\epsilon)d \;&\leq a^2 d + b^2(1-\gamma)d \;\Rightarrow\; |b| \leq \sqrt{\frac{\epsilon}{\gamma}} \end{aligned}$$

As in the previous section, we take R to be the nn adjacency matrix of the perturbed edges. Then by the argument in the proof of lemma 15, we can bound $\left\|(M-\hat{M})\hat{w}\right\|$ as

$$\left\|(M-\hat{M})\hat{w}\right\| \;\leq\; 2\left\|R\bar{w}\right\| \;\leq\; 2(a\sqrt{\epsilon}d + \sqrt{\frac{\epsilon}{\gamma}}d) \leq\; 3\sqrt{\frac{\epsilon}{\gamma}}d$$

which gives us the required bound on β, namely

$$|\beta| \;\leq\; 3\sqrt{\frac{\epsilon}{\gamma}}d \cdot \frac{1}{(\gamma-\epsilon)d} \;\leq\; 4\sqrt{\frac{\epsilon}{\gamma^3}}$$

From the above arguments we know that the eigenspace of the first k eigenvectors of $\hat{M}$ is close to the eigenspace of the first k eigenvectors of M. Also, this eigenspace contains the vectors $y^{(0)}, \ldots, y^{(k-1)}$ which encode the solutions. From the calculations in the proof of 16, we know that if v is a vector such that $v = \alpha y^{(s)} + \beta y_\perp$ for some $y^{(s)}$, then the algorithm in the previous section recovers the assignment for $(1 - \frac{2\beta^2}{\alpha^2})$ fraction of the variables.

To find a vector v close to one of the vectors $y^{(s)}$, we discretize the eigenspace of the first k eigenvectors of $\hat{M}$. Let $w^{(0)}, \ldots, w^{(k-1)}$ be the eigenvectors. We define the set S as

$$S \quad = \quad \left\{ v = \sum_{s=0}^{k-1} \alpha_s w^{(s)} \;\mid\; \alpha_s \in \frac{1}{10\sqrt{k}}\mathbb{Z},\;\; \|v\| \leq 1 \right\}$$

S contains at least one vector v such that $v = \alpha y^{(s)} + \beta y_\perp$ for some s and $\beta \leq 1/10 + 4\sqrt{\epsilon/\gamma^3} < 1/5$ for $\gamma > 20\epsilon^{1/3}$. Thus, for this vector v, the algorithm presented in the previous section recovers an assignment which agrees with $y^{(s)}$ in $(1 - \frac{1}{24})$ fraction of the block. Hence, the assignment violates at most $\frac{1}{24}nd + \epsilon nd < nd/20$ constaints. Since the total number of constraints is $nd/2$, this satisfies more than 90 percent of the constraints.

Finally, it remains to argue that the running time of the algorithm is polynomial. It can be calculated (see, for instance [FO05]) that the number of points in the set S is at most $e^{k \ln 90}$. Since $k = O(\log n)$ (this must hold for the long-code based reductions to be polynomial time), the number of points is polynomial in n. Hence, the algorithm runs in polynomial time.

Acknowledgements

We thank Luca Trevisan for suggesting the problem about the behavior of SDPs on random instances, and also for various helpful discussions. We also thank Ryan O'Donnell and James Lee for telling us about the problem of recovering planted solutions.

References

[AKV02] Noga Alon, Michael Krivelevich, and Van H. Vu. On the concentration of eigenvalues of random symmetric matrices. *Israel Journal of Mathematics*, 131:259–267, 2002.

[BS87] A. Broder and E. Shamir. On the second eigenvalue of random regular graphs. In *Proceedings of the 28th IEEE Symposium on Foundations of Computer Science*, pages 286–294, 1987.

[CMM06a] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for unique games. In *Proceedings of the 38th ACM Symposium on Theory of Computing*, 2006.

[CMM06b] Eden Chlamtac, Konstantin Makarychev, and Yury Makarychev. How to play unique games using embeddings. In *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science*, 2006.

[DK70] Chandler Davis and W.M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

[FL92] Uri Feige and László Lovász. Two-prover one round proof systems: Their power and their problems. In *Proceedings of the 24th ACM Symposium on Theory of Computing*, pages 733–741, 1992.

[FO05] Uriel Feige and Eran Ofek. Spectral techniques applied to sparse random graphs. *Random Structures and Algorithms*, 27(2):251–275, 2005.

[Fri91] Joel Friedman. On the second eigenvalue and random walks in random regular graphs. *Combinatorica*, 11(4):331–362, 1991.

[GT06] Anupam Gupta and Kunal Talwar. Approximating unique games. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*, pages 99–106, 2006.

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th ACM Symposium on Theory of Computing*, pages 767–775, 2002.

[KKMO04] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for MAX-CUT and other two-variable CSPs? In *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science*, pages 146–154, 2004.

[KR03] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. In *Proceedings of the 18th IEEE Conference on Computational Complexity*, 2003.

[KV05] Subhash Khot and Nisheeth Vishnoi. The unique games conjecture, integrality gap for cut problems and the embeddability of negative type metrics into ℓ_1. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, pages 53–63, 2005.

[Tre05] Luca Trevisan. Approximation algorithms for unique games. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.