

PCP theorem: gap amplification

CS 6810—David Steurer

Spring 2016

1 Introduction

In this note, we prove the following theorem, which is a major step toward a proof of the PCP theorem.

Recall that in the optimization problem 2-CSP_W , we are given a list of arity-2 constraints between variables over an alphabet of size W and the goal is to find an assignment to the variables that satisfies as many of the constraints as possible.

A 2-CSP_W instance φ is *d-regular* if every variable of φ appears in exactly d of the constraints of φ . It turns out that there is an efficient reduction to transform general instances of 2-CSP_W to *d-regular* ones without significantly affecting the optimal value of the instance.

Theorem (gap amplification). For all $d, t, W \in \mathbb{N}$, there exists a polynomial-time linear-blowup function f that maps every *d-regular* 2-CSP_W instance φ to a $2\text{-CSP}'_W$ instance $\varphi' = f(\varphi)$,

- **YES:** if $\text{opt}(\varphi) = 1$, then $\text{opt}(\varphi') = 1$
- **NO:** if $\text{opt}(\varphi) < 1 - \varepsilon$ for some $\varepsilon < 1/(td)$, then $\text{opt}(\varphi') < 1 - \varepsilon'$ for $\varepsilon' = t\varepsilon / \text{poly}(W, d)$
- **Efficiency:** $|W'| \leq |W|^{d^t}$ and $|\varphi'| \leq |\varphi| \cdot O(d^t)$

2 Interlude: Spectral graph expansion

A graph G is *d-regular* if every vertex of G is an endpoint of exactly d edges of G .

Let G be a *d-regular* graph G with n vertices. Its *normalized adjacency matrix* A_G is an n -by- n matrix indexed by the vertices of G such that

$$(A_G)_{u,v} = \begin{cases} 1/d & \text{if } (u,v) \in E(G), \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

We extend this definition to graphs with parallel edges and self-loops: We set $(A_G)_{u,v}$ to the fraction of u 's edges that are incident to v .

The following observation shows how A_G acts on real-valued vectors indexed by vertices of G .

Observation (averaging over neighbors). For every vector $x \in \mathbb{R}^{V(G)}$ and every vertex $u \in V(G)$,

$$(A_G x)_u = \text{average } x\text{-value of the neighbors of } u.$$

Since the matrix A_G is symmetric and real-valued, its eigenvalues $\lambda_1, \dots, \lambda_n$ are real.

Claim. All eigenvalues of A_G are contained in the interval $[-1, 1]$.

Proof. We use the following general bound on the largest eigenvalue in absolute value of a symmetric matrix,

$$\max_{i \in [n]} |\lambda_i| \leq \max_{u \in V(G)} \sum_{v \in V(G)} |(A_G)_{u,v}| = 1. \quad (2)$$

Claim. The vector $\mathbf{1} = (1, \dots, 1)^T$ is an eigenvector of A_G with value 1.

Proof. The averaging-over-neighbors observation shows that $A\mathbf{1} = \mathbf{1}$.

The previous two claims show that we can order the eigenvalues of A_G such that

$$1 = \lambda_1 \geq |\lambda_2| \geq \dots \geq |\lambda_n|. \quad (3)$$

Definition (Spectral expansion). For every regular graph, we define its *spectral expansion* $\lambda(G) = \lambda_2$ as the second largest eigenvalue in absolute value of G 's normalized adjacency matrix.

The following lemma gives a combinatorial characterization of the condition $\lambda(G) < 1$. Its proof is an exercise.

Lemma. A graph G is connected and non-bipartite if and only if $\lambda(G) < 1$.

It turns out that the spectral expansion parameter λ is a surprisingly useful way to quantify the “connectedness” and “non-bipartiteness” of a graph. Graphs with spectral expansion bounded away from 1, say $\lambda < 0.9$, behave in many ways similar to random graphs. We refer to such graphs as expander graphs.

Definition (Expander graphs). A collection $\{G_n \mid n \in \mathbb{N}\}$ of graphs is called a *family of expander graphs* if there exists a constant $\lambda < 1$ such that every graph G_n in the family satisfies $\lambda(G_n) \leq \lambda$.

For every $d \geq 3$, there exist families of d -regular expander graphs and they can be computed in polynomial time. (In fact, a random d -regular graph is an expander with high probability.)

The following property of expander graphs is useful for the analysis of our gap amplification construction. Its proof is an exercise.

Lemma (Expander mixing lemma). For every regular graph G and every set $S \subseteq V(G)$,

$$\mathbb{P}_{uv \in E(G)} \{u \in S \wedge v \in S\} \leq \mathbb{P}_{u \in V(G)} \{u \in S\} \cdot \left(\mathbb{P}_{v \in V(G)} \{v \in S\} + \lambda(G) \right). \quad (4)$$

Note that the upper bound $\mathbb{P}_{uv \in E(G)} \{u \in S \wedge v \in S\} \leq \mathbb{P}_{u \in V(G)} \{u \in S\}$ holds for every graph G and every set $S \subseteq V(G)$. The expander mixing lemma implies that this upper bound is

far from tight whenever the spectral expansion of G is bounded away from 1 and the set S excludes a constant fraction of the vertices of G .

If we choose S as a random set of size k in a regular graph G with n vertices, then

$$\mathbb{E}_{\substack{S \subseteq V(G) \\ |S|=k}} \mathbb{P}_{uv \in E(G)} \left\{ u \in S \wedge v \in S \right\} \approx \frac{k}{n} \cdot \frac{k}{n}. \quad (5)$$

The expander mixing lemma implies that every set $S \subseteq V(G)$ of size k satisfies $\mathbb{P}_{uv \in E(G)} \left\{ u \in S \wedge v \in S \right\} \approx \frac{k}{n} \cdot \frac{k}{n}$ if $\lambda(G) \ll \frac{k}{n}$.

Definition (path power). For a graph G , its t -fold path power G^t has the same set of vertices $V(G^t) = V(G)$ as G and for every t -step walk u_1, \dots, u_t in G , the graph G^t contains an edge between the walk endpoints u_1 and u_t .

If G is d -regular, then its power G^t is d^t regular. (The graph G^t may have parallel edges and self-loops.) Furthermore, the normalized adjacency matrix of G and G^t satisfy the relation,

$$A_{G^t} = (A_G)^t. \quad (6)$$

Since the eigenvalues of the power of a matrix are the powers of its eigenvalues, the spectral expansion of G^t is the t -th power of the spectral expansion of G .

Lemma (spectral expansion of path power). For every graph G and every $t \in \mathbb{N}$, the t -fold path power G^t has spectral expansion $\lambda(G^t) = \lambda(G)^t$.

This lemma shows that the path power operation improves spectral expansion. This fact underlies some of the constructions of expander graphs.

3 Gap amplification construction

Let φ be a d -regular 2-CSP $_W$ instance with variables x_1, \dots, x_n over alphabet $[W]$. Let G be the *constraint graph* of φ , i.e., the d -regular graph on $[n]$ that has an edge between vertices $u \in [n]$ and $v \in [n]$ for every constraint between variables x_u and x_v in φ . For an edge $e \in E(G)$ between u and v , let $\varphi_e: [W] \times [W] \rightarrow \{0, 1\}$ be the corresponding constraint in φ . Then, the objective value of an assignment $x \in [W]^n$ for φ , denoted $\varphi(x)$, satisfies

$$\varphi(x) = \mathbb{P}_{e=(u,v) \in E(G)} \left\{ \varphi_e(x_u, x_v) = 1 \right\}. \quad (7)$$

For every vertex $u \in V(G)$, let $N_t(u) \subseteq V(G)$ denote the t -step neighborhood of u in G , i.e., all vertices that can be reached from u by a path of length at most t .

Construction of 2-CSP $_W$ instance φ^t .

- **variables:** y_1, \dots, y_n
- **alphabet:** size $W' = W^{d^t}$; for variable y_u , we identify alphabet with an assignments to the t -step neighborhood $N_t(u)$ so that y_u takes values in $[W]^{N_t(u)}$; we refer to $(y_u)_v$ as the “opinion” of u about v .
- **constraints:** for every edge $(u, v) \in E(G^t)$, we have the following constraint $\varphi_{u,v}: W^{N_t(u)} \times W^{N_t(v)} \rightarrow \{0, 1\}$,

$$\varphi_{u,v}^t(y_u, y_v) = 1 \Leftrightarrow \left\{ \begin{array}{l} \text{the assignments } y_u \text{ and } y_v \text{ agree on } N_t(u) \cap N_t(v) \text{ and} \\ \text{satisfy all constraints of } \varphi \text{ between variables in } N_t(u) \cup N_t(v) \end{array} \right\} \quad (8)$$

This construction satisfies the desired YES property (for the gap amplification theorem): If $\text{opt}(\varphi) = 1$, then $\text{opt}(\varphi^t) = 1$. Concretely, if x_1, \dots, x_n is an assignment for φ that satisfies all constraints, then the following assignment y_1, \dots, y_n satisfies all constraints of φ^t ,

$$\forall u, v \in [n]. \quad (y_u)_v = x_v. \quad (9)$$

The construction of φ^t also satisfies the desired efficiency property: The number of constraints in φ^t is at most d^t factor larger than the number of constraints in φ . Therefore, $|\varphi^t| \leq |\varphi| \cdot O(d^t)$.

4 Analysis against consistent assignments

In order to finish the proof of the gap amplification theorem it remains to establish the NO property of the construction of φ^t .

Let φ be a d -regular instance of 2-CSP $_W$. Suppose that $\text{opt}(\varphi) < 1 - \varepsilon$ for some $\varepsilon < 1/(dt)$. We are to show that $\text{opt}(\varphi^t) < 1 - t\varepsilon/\text{poly}(d, W)$.

As a first step toward this goal, we show that all “consistent assignment” for φ^t indeed achieve objective value at most $1 - t\varepsilon/\text{poly}(d)$ as long as the constraint graph G of φ has graph expansion $\lambda(G)$ bounded away from 1. The assumption that the constraint graph has $\lambda(G)$ bounded away from 1 is without loss of generality because we can ensure this property by adding constraints to φ that are always satisfied and have an d -regular expander as constraint graph.

Definition. We say that an assignment y_1, \dots, y_n is *consistent* if there exists an assignment x_1, \dots, x_n such that $(y_u)_v = x_v$ for all $u, v \in [n]$.

Lemma. Suppose the constraint graph G of φ satisfies $\lambda(G) < 0.9$. Then, every consistent assignment y has objective value at most $1 - t\varepsilon/100d$ for φ^t .

4.1 Proof of lemma

Let y be an assignment for φ^t consistent with assignment x for φ . Suppose that x has objective value $1 - \eta$ for φ . Note that $\eta > \varepsilon$. (Recall that $\text{opt}(\varphi) < 1 - \varepsilon$.) We may also assume $\eta < 1/(dt)$. (Otherwise, we can upper bound the value of the assignment y in a simpler way.)

Let $W = e(1), \dots, e(t)$ be a random walk of length t and let φ_W^t be the corresponding constraint in φ^t . Let X_i be the indicator random variable of the event that x does not satisfy the constraint $\varphi_{e(i)}$. Let $X = X_1 + \dots + X_t$ be the number of constraints on walk W that are not satisfied by assignment x .

Note that if $X > 0$, then y does not satisfy the constraint φ_W^t . Therefore we can bound the objective value achieved by y in terms of the support of X ,

$$\begin{aligned} \mathbb{P}_{W \in E(G^t)} \left\{ y \text{ satisfies constraint } \varphi_W^t \right\} &\leq 1 - \mathbb{P}\{X > 0\} \\ &\leq 1 - (\mathbb{E} X)^2 / \mathbb{E} X^2. \end{aligned} \quad (10)$$

The second step uses Cauchy–Schwarz and the fact that $\mathbb{E} X \geq 0$ (see [homework 3](#)).

It remains to estimate the first two moments of X .

Claim 1. $\mathbb{E} X \geq t \cdot \eta$.

Proof. By linearity of expectation, $\mathbb{E} X = \sum_{i=1}^t \mathbb{E} X_i$. Therefore, it suffices to show that $\mathbb{E} X_i \geq \eta$ for all $i \in [t]$. Since the constraint graph G is regular, the distribution of $e(i)$ is uniform over the edges of G . Therefore, $\mathbb{E} X_i$ is the fraction of constraints in φ not satisfied by assignment x , which means $\mathbb{E} X_i = \eta$. \square

Claim 2. $\mathbb{E} X^2 \leq 30td\eta$.

Proof. Let S be the set of vertices adjacent to a constraint not satisfied by x . Since G is d -regular and x satisfies all but an η fraction of constraints, $\mathbb{P}_{u \in V(G)}\{u \in S\} \leq d \cdot \eta$. Let Y_i be the indicator of the event that the i -th vertex $w(i)$ of the walk W is contained in S . Since G is regular, the distribution of $w(i)$ is uniform over the vertices of G . Therefore, $\mathbb{E} Y_i \leq d \cdot \eta$. Let $Y = Y_1 + \dots + Y_t$ be the number of vertices of W contained in S . Since $X \leq Y$, we can bound the second moment of X by

$$\begin{aligned} \mathbb{E} X^2 &\leq \mathbb{E} Y^2 = \sum_i \mathbb{E} Y_i + 2 \sum_{i < j} \mathbb{E} Y_i Y_j \\ &\leq t \cdot d\eta + 2 \sum_{i < j} \mathbb{E} Y_i Y_j. \end{aligned} \quad (11)$$

The second step uses linearity of expectation and the third step uses $\mathbb{E} Y_i \leq d\eta$. Note that $Y_i Y_j$ is the indicator of the event $\{w(i) \in S, w(j) \in S\}$. Since G is regular, the distribution of $w(i), w(j)$ is uniform over edges of G^{j-i} (endpoints of walks of length $j - i$ in G). Thus, by the expander mixing lemma,

$$\begin{aligned} \mathbb{E} Y_i Y_j &\leq d\eta \cdot \left(d\eta + \lambda(G^{j-i}) \right) \\ &= d\eta \cdot \left(d\eta + \lambda(G)^{j-i} \right) \end{aligned} \quad (12)$$

The second step uses our previous lemma about the spectral expansion of power graphs.

Putting these bounds together shows

$$\mathbb{E} X^2 \leq t \cdot d\eta + 2t^2 \cdot (d\eta)^2 + 2t \cdot d\eta \sum_{k=1}^t \lambda(G)^k \leq 30td\eta \quad (13)$$

The last step uses the assumption $td\eta < 1$ and that $\sum_{k=1}^{\infty} \lambda(G)^k \leq 10$ (because $\lambda(G) < 0.9$). \square

5 Modified construction (easier analysis)

So far we know that the construction of φ^t satisfies the desired NO property if we restrict ourselves to consistent assignments. We modify the construction slightly in order to simplify the analysis against all assignments.

The modified construction of φ^t is the same as before except that we assign different weights to the new constraints $\varphi_{u,v}^t$. The modified construction has a parameter $\delta > 0$. (We can choose $\delta = 1/(4W)$.)

Weights of constraints in modified construction. The weight of the constraint $\varphi_{u(1),u(2)}^t$ in φ^t is equal to the probability of $u(1)$ and $u(2)$ being generated by the following randomized algorithm:

- choose a random path $v(1), \dots, v(t_0)$ of length δt in the constraint graph G of φ
- choose integers t_1, t_2 uniformly at random from interval $[t/2, t]$
- choose a random path P_1 from $v(1)$ of length t_1 and a random path P_2 from $v(t_0)$ of length t_2
- let $u(1)$ be the endpoint of P_1 and $u(2)$ be the endpoint of P_2




Figure 1: Generation of modified constraint distribution. A random walk of length $t_1 + \delta t + t_2$, where t_1 and t_2 are chosen uniformly between $t/2$ and t

6 Analysis against all assignments

Let y be an arbitrary assignment for φ^t (not necessarily consistent). We are to upper bound the objective value of y .

Our strategy is to construct an assignment x for φ from the assignment y for φ^t . We then use the constraints violated by x in order to bound the objective value of y for φ^t .

Plurality assignment obtained from y . Let x be the assignment for φ such that for every vertex v , we choose x_v as the most likely output of the following randomized algorithm:

- choose an integer t' uniformly at random from the interval $[t/2, t]$
- choose a random path P' from v of length t'
- let u' be the endpoint of P'
- output $(y_{u'})_v \in [W]$, that is, the *opinion of u' about v*

Note that if y is consistent with an assignment x for φ , then x is also the plurality assignment constructed above.

Let $\eta > 0$ be the fraction of constraints not satisfied by the plurality assignment x . For every $i \in [\delta t]$, let X_i be the indicator random variable such that $X_i = 1$ if and only if all of the following conditions are satisfied

- the constraint between $v(i-1)$ and $v(i)$ in φ is not satisfied by the plurality assignment x
- the opinion $(y_{u(1)})_{v(i-1)}$ of $u(1)$ about $v(i-1)$ differs from the plurality assignment $x_{v(i-1)}$
- the opinion $(y_{u(2)})_{v(i)}$ of $u(2)$ about $v(i)$ differs from the plurality assignment $x_{v(i)}$

Note that if $X_i = 1$, then the assignment y does not satisfy the constraint between $u(1)$ and $u(2)$ in φ^t . Let $X = X_1 + \dots + X_{\delta t}$. As before, we can lower bound the fraction of constraints not satisfied by y in terms of the first two moments of X ,

$$\mathbb{P}\left\{\varphi_{u(1),u(2)}^t(y_{u(1)}, y_{u(2)}) = 0\right\} \geq \mathbb{P}\{X > 0\} \geq (\mathbb{E} X)^2 / \mathbb{E} X^2. \quad (14)$$

It remains to estimate the first two moments of X . It turns out that in order to show the bound $\mathbb{E} X^2 \leq 30dt\eta$ essentially the same proof works as before. However the lower bound on $\mathbb{E} X$ requires new ideas. In particular, we will use basic properties of the statistical distance between random variables (related to the [total variation distance of probability distributions](#)).

Definition. For (finitely-supported) random variables A and B , let $\Delta(A, B)$ be their *statistical distance*,

$$\Delta(A, B) = \sum_{\omega \in \text{support}(A) \cup \text{support}(B)} |\mathbb{P}\{A = \omega\} - \mathbb{P}\{B = \omega\}| / 2 \quad (15)$$

Claim 3. For every $i \in [t]$, we have

$$\mathbb{P}\left\{\left(y_{u(1)}\right)_{v(i-1)} = x_{v(i-1)} \wedge \left(y_{u(2)}\right)_{v(i)} = x_{v(i)} \mid v(i-1), v(i)\right\} \geq \left(\frac{1}{W} - 2\delta\right)^2 \quad (16)$$

Proof. Since $u(1)$ and $u(2)$ are statistically independent conditioned on $\{v(i-1), v(i)\}$, it is enough to show that each of the events $\{(y_{u(1)})_{v(i-1)} = x_{v(i-1)}\}$ and $\{(y_{u(2)})_{v(i)} = x_{v(i)}\}$ have probability at least $1/W - 2\delta$ conditioned on $\{v(i-1), v(i)\}$. By symmetry between $u(1)$ and $u(2)$, it remains to show that

$$\mathbb{P}\left\{\left(y_{u(1)}\right)_{v(i-1)} = x_{v(i-1)} \mid v(i-1), v(i)\right\} \geq \frac{1}{W} - 2\delta. \quad (17)$$

Let (t', P', u') be the random variables as defined in the construction of the plurality assignment for x_v with $v = v(i-1)$. Since $x_{v(i-1)}$ is the most likely value for $(y_{u'})_v$ and there are only W possible outcomes, we have

$$\mathbb{P}\left\{\left(y_{u'}\right)_{v(i-1)} = x_{v(i-1)} \mid v(i-1)\right\} \geq \frac{1}{W}. \quad (18)$$

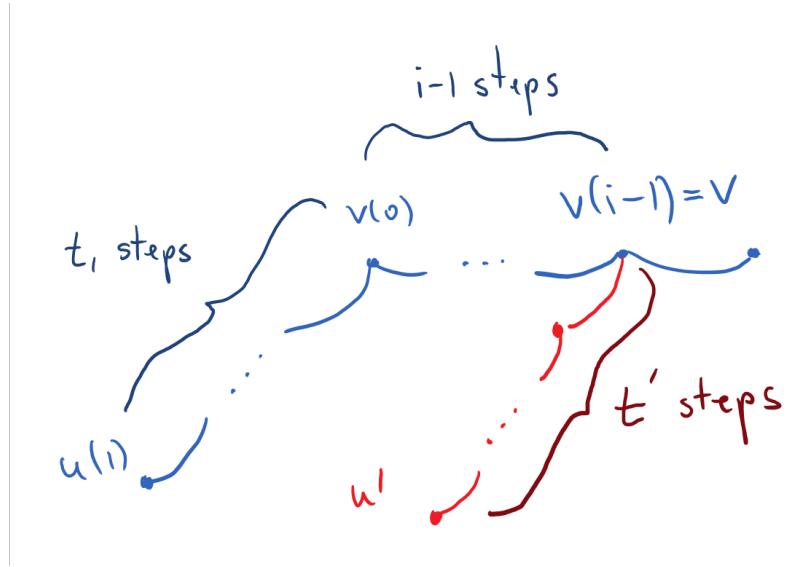


Figure 2: Conditioned on $v(i-1)$ the vertex $u(1)$ is the endpoint of a random walk of length $i-1+t_1$ from $v(i-1)$ and u' is the endpoint of random walk of length t' from $v(i-1)$. The plurality assignment for $v(i-1)$ is the most likely opinion of u' about $v(i-1)$.

We connect the bound (18) to the target (17) by estimating the statistical distance between $(y_{u'})_{v(i-1)}$ and $(y_{u(1)})_{v(i-1)}$ conditioned on $v(i-1)$,

$$\begin{aligned} \Delta\left(\left\{(y_{u(1)})_{v(i-1)} \mid v(i-1)\right\}, \left\{(y_{u'})_{v(i-1)} \mid v(i-1)\right\}\right) &\leq \Delta\left(\left\{u(1) \mid v(i-1)\right\}, \left\{u' \mid v(i-1)\right\}\right) \\ &\leq \Delta\left(t_1 + i - 1, t'\right) \\ &\leq \frac{(i-1)}{t/2}. \end{aligned} \quad (19)$$

It is instructive to verify the above computation first for the case $i = 1$. Here the statistical distance is indeed 0 because both $u(1)$ and u' are generated in the same way—by a random walk from $v(0)$ whose length is uniformly distributed in $[t/2, t]$. For larger values of i , the statistical distance arises because $u(1)$ and u' are generated by random walks with differently distributed lengths. The length of the random walk for u' is still uniformly distributed in $[t/2, t]$. However, the length of the random walk for $u(1)$ (when viewed from $v(i-1)$) is uniformly distributed between $i-1+t/2$ and $i-1+t$. The statistical distance between these two distributions for random walk lengths is $(i-1)/(t/2)$. Formally, the reasoning among uses the data processing inequality for statistical distance (see for example these [lecture notes](#) as a reference).




Figure 3: A plot of the probabilities of the two distribution of walk lengths. The area of the shaded region corresponds to the statistical distance.

We conclude the desired bound (17) by

$$\begin{aligned}
& \mathbb{P}\left\{\left(y_{u(1)}\right)_{v(i-1)} = x_{v(i-1)} \mid v(i-1), v(i)\right\} \\
& \geq \mathbb{P}\left\{\left(y_{u'}\right)_{v(i-1)} = x_{v(i-1)} \mid v(i-1)\right\} \\
& \quad - \Delta\left(\left\{\left(y_{u(1)}\right)_{v(i-1)} \mid v(i-1)\right\}, \left\{\left(y_{u'}\right)_{v(i-1)} \mid v(i-1)\right\}\right) \\
& \geq 1/W - (i-1)/(t/2) \geq 1/W - 2\delta. \quad \square
\end{aligned} \tag{20}$$

The previous claim allows us to show the desired lower bound on $\mathbb{E} X$.

Claim 4. $\mathbb{E} X \geq \delta t \eta \cdot \left(\frac{1}{W} - 2\delta\right)^2$. In particular, if we choose $\delta = 1/(4W)$, then $\mathbb{E} X \geq t\eta/(16W^2)$.

Proof. By linearity of expectation, $\mathbb{E} X = \sum_i \mathbb{E} X_i$. Thus, it suffices to show that every i satisfies $\mathbb{E} X_i \geq \eta/(100W^2)$. Indeed, by the definition of X_i ,

$$\begin{aligned}
\mathbb{E} X_i &= \mathbb{P}\left\{\varphi_{v(i-1), v(i)}(x_{v(i-1)}, x_{v(i)}) = 0\right\} \\
& \quad \cdot \mathbb{P}\left\{\left(y_{u(1)}\right)_{v(i-1)} = x_{v(i-1)} \wedge \left(y_{u(2)}\right)_{v(i)} = x_{v(i)} \mid v(i-1), v(i)\right\} \\
& \geq \mathbb{P}\left\{\varphi_{v(i-1), v(i)}(x_{v(i-1)}, x_{v(i)}) = 0\right\} \cdot \left(\frac{1}{W} - 2\delta\right)^2 \\
& = \eta \cdot \left(\frac{1}{W} - 2\delta\right)^2
\end{aligned} \tag{21}$$

The first inequality is by Claim 3. The last step uses that η is the fraction of constraints not satisfied by the plurality assignment x in φ . \square

The same proof as for Claim 2 bounds the second moment of X (again under the mild assumption that $\lambda(G) < 0.9$).

Claim 5. $\mathbb{E} X^2 \leq 30dt\eta$.

At this point, we have all components of the proof of the gap amplification theorem.

6.1 Proof of gap amplification theorem

Choose the parameter of our (modified) construction of φ^t as $\delta = 1/(4W)$. By (14), we can lower bound the fraction of constraints not satisfied by assignment y in φ^t

$$\begin{aligned} \mathbb{P}\left\{\varphi_{u(1),u(2)}^t(y_{u(1)}, y_{u(2)}) = 0\right\} &\geq \frac{(\mathbb{E} X)^2}{\mathbb{E} X^2} \\ &\geq \frac{\left(t\eta/(16W)\right)^2}{30td\eta} \\ &= t\eta \cdot \frac{1}{7680dW^2} \geq t\varepsilon \cdot \frac{1}{\text{poly}(d, W)}. \end{aligned} \tag{22}$$

The second step uses Claim 4 and Claim 5.

It follows that our construction has the desired NO property: whenever $\text{opt}(\varphi) < 1 - \varepsilon$ for some $\varepsilon > 1/(dW)$, the constructed instance φ^t satisfies $\text{opt}(\varphi^t) < 1 - \varepsilon'$ for $\varepsilon' = t\varepsilon / \text{poly}(d, W)$.

□

Footnotes