CS369E: Expanders May 31, 2005

Lecture 9: Dinur's Proof of the PCP Theorem

Lecturer: Prahladh Harsha *Scribe: Krishnaram Kenthapadi*

In this lecture, we will describe a recent and remarkable proof of the PCP theorem, due to Irit Dinur [Din]. This proof is a beautiful application of expanders for "soundness amplification" of randomized algorithms without using too many random bits.

Before describing the proof, we will first look at the PCP theorem, by relating it with the theory of NP-completeness.

9.1 Hardness of Optimization Problems

The theory of NP-completeness, as developed by Cook, Levin, and Karp, states that any language, L in NP is reducible to the Boolean satisfiability problem, 3SAT. By this, we mean that for every instance, x of the language L, we can obtain a satisfiability instance, ϕ such that $x \in L$ if and only if ϕ is satisfiable. Thus, 3SAT is at least as hard as any other problem in NP. Karp further showed that 3SAT can be reduced to other problems such as CLIQUE and 3-COLORABILITY and hence that these problems are at least as hard as any problem in NP. In other words, solving these problems *optimally* is as hard as solving any other problem in NP optimally.

However the question of the hardness of approximation was left open. For instance, can the following be true – finding a satisfying assignment for 3SAT is NP-hard, however it is easy to find an assignment that satisfies 99% of the clauses. Questions such as Other examples: can we approximate the clique size in a graph? Or, can we obtain a 3-coloring that satisfies 99% of the edge constraints? In other words, is the approximation version of some of NP-hard problems easier than the optimization versions. The PCP Theorem [FGLSS, AS, ALMSS] states that this is not the case – for several of the NP-hard problems, the approximation version is just as hard as the optimization version. The PCP theorem can be viewed as a strengthening of Karp reductions. It provides a reduction from a 3SAT instance ϕ to another 3SAT instance ψ such that if $\phi \in 3SAT$, then $\psi \in 3SAT$ and if $\phi \notin 3SAT$, then any assignment to ψ violates at least α fraction of the clauses. This provides a hardness of approximation result for MAX-3SAT (i.e., the problem of finding the assignment that satisfies the most number of clauses in a given 3CNF formula).

Theorem 9.1 *There exists a constant $0 < \alpha < 1$ such that MAX-3SAT is $(1 - \alpha)$-hard to approximate unless $P = NP$.*

Observe that the theory of NP-completeness provides $\alpha = 1/n$ instead of a constant. The PCP Theorem amplifies this sub-constant soundness $1/n$ to some constant α. Starting from the PCP Theorem, it is possible to derive hardness of approximation results for several optimization problems.

9.2 Dinur's Proof of the PCP Theorem

We define the NP-complete problem called CONSTRAINT-GRAPH (CG). An instance of CG is of the form $G = ((V, E), \Sigma, \mathcal{C})$ where (V, E) is an undirected graph, Σ a constant-sized set of colors and $\mathcal{C}$ is a set of constraint functions, one corresponding to each graph edge, i.e., $\mathcal{C} = \{c_E : \Sigma^2 \to \{0, 1\} | e \in E\}$. A coloring that assigns color c_1 to vertex v_1 and c_2 to vertex v_2 is said to satisfy the coloring constraint $c_{(v_1,v_2)}$ on edge (v_1, v_2) is $c_{(v_1,v_2)}(c_1, c_2) = 1$. An instance $((V, E), \Sigma, \mathcal{C})$ is an YES instance of CG, if there exists a coloring $\sigma : V \to \Sigma$ that satisfies all of $\mathcal{C}$. Since 3-COLORING is NP-complete, it easily follows that CG is also NP-complete.

The α-approximate version of CG is whether there exists a coloring that satisfies at least $(1 - \alpha)$-constraints. In today's lecture, we will prove the following equivalent version of the PCP Theorem. It is an easy exercise to show that these two versions are equivalent. Now the PCP theorem can be equivalently stated as follows.

Theorem 9.2 (Dinur [Din]) *There exists a constant α such that the α-approximation version of CONSTRAINT-GRAPH is NP-hard.*

We will restate this theorem in a more convenient form in terms of a reduction. For this we need some notation. For any instance $G = ((V, E), \Sigma, \mathcal{C})$ of CG, let $n = |\mathcal{C}|$ and $size(G) = |V| + |E|$. Let σ_G be the best colorings for G (i.e., it is the coloring that violates the least number of edge constraints.) If there is more than one, set σ_G to be one of them arbitrarily. Let $UNSAT(G)$, called the *unsatisfiability factor*, be the fraction of edge constraints violated by σ_G. Note that if there exists a coloring that satisfies all the constraints of G, then $UNSAT(G) = 0$. Since CG is itself a NP-complete problem, the above theorem can be restated as follows in terms of a reduction from the decision version of CG to its approximation version.

Theorem 9.3 (PCP Theorem as a reduction) *There exists a constant α and a polynomial time reduction from CG to the α-approximation version of CG that maps the instance, $G = ((V, E), \Sigma, \mathcal{C})$ to the instance, $G' = ((V', E'), \Sigma', \mathcal{C}')$ such that*

- $size(G') = poly(size(G))$
- *Completeness:* $UNSAT(G) = 0 \Rightarrow UNSAT(G') = 0$
- *Soundness:* $UNSAT(G) \geq 1/n \Rightarrow UNSAT(G') \geq \alpha$

We will prove the above theorem by applying the following *Gap Amplification* Lemma for $O(\log n)$ steps.

Lemma 9.4 (Gap Amplification Lemma) *There exists a constant $0 < \alpha < 1$, a color set Σ and a polynomial time reduction from CG to itself mapping the instance, $G = ((V, E), \Sigma, \mathcal{C})$ to the instance, $G' = ((V', E'), \Sigma', \mathcal{C}')$ such that*

- $size(G') = O(size(G))$ *and* $\Sigma' = \Sigma$
- $UNSAT(G) = 0 \Rightarrow UNSAT(G') = 0$

- $UNSAT(G') \geq 2 \cdot \min(UNSAT(G), \alpha)$

We can now prove the PCP Theorem 9.3 starting from the Gap Amplification Lemma.
Proof of Theorem 9.3: We first observe that the gap amplification increases the unsatisfiability factor of the instance G by a factor of 2 (if it is not already a constant) and in doing so it blows up the size of the instance by at most a constant factor. We can hence apply this lemma $O(\log n)$ times to improve the gap from $1/n$ to α with at most a polynomial blowup in size, thus proving the PCP Theorem 9.3. ■

Thus it suffices for us to prove the Gap Amplification Lemma 9.4 and this will be our goal for the rest of the lecture. But first we need some preliminaries regarding edge expansion of expanders.

9.3 Expanders – Edge Expansion

For a graph $G = (V, E)$, the edge expansion $\phi(G)$ is defined as follows:

$$\phi(G) = \min_{|S| \leq n/2} \frac{E(S, \bar{S})}{|S|}.$$

The following lemma provides an inverse relationship between the edge expansion, $\phi(G)$ and the spectral expansion, $\lambda(G)$.

Lemma 9.5 *For a d-regular graph G,* $\frac{\phi^2(G)}{2d} \leq 1 - \lambda(G) \leq \frac{2\phi(G)}{d}$

Thus, if we a graph with good spectral expansion, it also has good edge expansion. Since we know how to construct good spectral expanders, we can also assume the following. There exists a family of constant degree edge-expanders that can be explicitly constructed, i.e., there exists a constant ϕ_0, such that for all n, there exists a d-regular graph G_n on n vertices such that $\phi(G_n) \geq \phi_0$ and furthermore there exists an explicit construction of such graphs.

We also require the following estimate on the random-like behavior of random walk on an expander.

Lemma 9.6 *Let $G = (V, E)$ be a d-regular graph with spectral expansion λ. Let $B \subset E$ be a set of edges. The probability p that a random walk that starts at a random edge in B takes its $(i+1)^{\text{st}}$ step in B as well, is bounded above by $\frac{|B|}{|E|} + \lambda^i$.*

Note that if the edges were chosen randomly and independently (instead of choosing them along a random walk) then the above probability p is exactly $\frac{|B|}{|E|}$. The above lemma states that choosing the edges according to a random walk worsens this probability by at most λ^i.

The proof of this lemma is similar to that of the Expander Mixing Lemma. This proof is reproduced verbatim from Dinur's paper [Din].
Proof: Let K be the distribution on vertices of G induced by selecting a random edge in B and then a random vertex on which the edge is incident on. Let W be the support of the distribution K. As always, let A be the normalized adjacency matrix of G.

Let π be the vector corresponding to the distribution K. Hence, π_v is the fraction of edges incident on v that are in B, divided by 2. For any vertex v, let B_v denote the set of edges incident on v that are in B. Hence, $\pi_v = |B_v|/2|B| \leq d/2|B|$ since G is d-regular. Let y_v be the probability that a random step from v is in B, so $y_v = |B_v|/d = 2|B|\pi_v/d$. The probability p equals the probability of landing in W after i steps and then taking a step in B. Hence

$$p = \sum_{v \in W} y_v (A^i \pi)_v = \sum_{v \in V} y_v (A^i \pi)_v = \langle y, A^i x \rangle.$$

let u be all ones vector. Decomposing π along u and its orthogonal component we have $\pi = \pi^{\|} + \pi^{\perp}$. Observe that

$$\|\pi\|_2^2 \leq \left(\sum_v \pi_v \right) \cdot \left(\max_v \pi_v \right) \leq 1 \cdot \frac{d}{2|B|} = \frac{d}{2|B|}.$$

Since G has spectral expansion λ,

$$\begin{aligned} \|A^i \pi^{\perp}\|_2 &\leq \lambda^i \|\pi^{\perp}\|_2 \\ &\leq \lambda^i \|\pi\|_2 \\ &\leq \lambda^i \sqrt{\frac{d}{2|B|}} \end{aligned}$$

By Cauchy-Schwarz,

$$\langle y, A^i \pi^{\perp} \rangle > \leq \|y\|_2 \|A^i \pi^{\perp}\|_2 \leq \frac{2|B|}{d} \lambda^i \|\pi\|_2^2 \leq \lambda^i$$

Combining we have,

$$p = \langle y, A^i \pi \rangle = \langle y, A^i \pi^{\|} \rangle + \langle y, A^i \pi^{\perp} \rangle \leq \frac{2|B|}{dn} + \lambda^i = \frac{|B|}{|E|} + \lambda^i.$$

■

9.4 Proof of Gap Amplification Lemma

The reduction in the gap amplification lemma is achieved by a two-step process:

Preprocessing step This preprocessing step converts an arbitrary graph into a constant degree expander and worsens the unsatisfiability by at most a constant factor. This step blows up the size by at most a constant factor.

Powering step Assuming that the graph is a constant-degree graph, this step performs a "powering operation" that amplifies the unsatisfiability factor of the graph while blowing up the size by at most a constant factor. This step increase the size of the color-set Σ. There exist standard techniques, namely proof composition in PCP technology, that are precisely designed for reducing the size of the color-set. Proof Composition attains color-set reduction while just mildly worsening other parameters. For want of time, we will not consider this issue in lecture but for stating the required result.

9.4.1 Graph Preprocessing

The preprocessing step involves converting the graph into a constant degree expander graph. This is performed in two steps: (a) converting the graph into a constant degree graph and (b) "expanderizing" the constant degree graph.

Conversion into a constant degree graph Let G_n be a family of expander graph with degree d and edge expansion at least ϕ_0.

The given graph $G = (V, E)$ is transformed as follows: A vertex, v with degree d_v is replaced by an expander G_{d_v} on d_v vertices and the edges incident on v are now assigned to the vertices of G_{d_v}, one edge per vertex. All the vertices in the transformed graph, $G' = (V', E')$ thus have degree $d+1$, where d is the degree of any graph in the expander family. All the edges inside each expander graph have equality constraints while the external edges retain the constraint they had earlier.

$$|V'| = \sum d_v = 2|E|$$

$$|E'| = \frac{d+1}{2}|V'| = (d+1)|E|$$

Thus, the size of the new graph $G' = (V', E')$ is at most a constant factor that of G.

Clearly, if $UNSAT(G) = 0$, then so is $UNSAT(G')$.

We now need to show that if $UNSAT(G)$ is non-zero, then $UNSAT(G')$ is worsened (i.e., reduced) at most by a constant factor. The intuition is that we can try to cheat by giving different colors to the d_v vertices. However, due to the property of the expander, this will result in violating several of the equality constraints within each expander.

Let $\sigma' = \sigma'_{G'} : V' \to \Sigma$ be the best coloring for G'. From this, we can obtain a coloring $\sigma : V \to \Sigma$ for G, in which the color of a vertex v is the most popular of the colors assigned to the corresponding "cloud" of d_v vertices in G'.

Let $\mu = UNSAT(G)$. Let B be the set of edges violated by σ in G and B' be the set of edges violated by σ' in G'. Define S to be the set of vertices in G' whose color is not the popular one (in the corresponding cloud). Since every edge in B should either be in B' or contribute to S, we have $\mu|E| \leq |B| \leq |B'| + |S|$.

- *Case 1*: $|B'| \geq \mu|E|/2$

$$UNSAT(G') = \frac{|B'|}{|E'|} \geq \frac{\mu|E|}{2|E'|} = \frac{\mu}{2(d+1)} = \frac{UNSAT(G)}{2(d+1)}$$

- *Case 2*: $|S| \geq \mu|E|/2$
  Consider any vertex v in G and its corresponding cloud of vertices in G'. Let S^v be the set of vertices in the cloud which did not get the popular color. For each color a, define $S^v_a = \{u \in S^v | \sigma'(u) = a\}$. By the definition of popularity, $|S^v_a| < d_v/2$. Now, from the expansion property within each cloud, we get that $|E(S^v_a, \bar{S^v_a})| \geq \phi_0 |S^v_a|$. Note that the constraints for all the edges in $E(S^v_a, \bar{S^v_a})$ are violated. Summing over the colors and clouds,

$$|B'| \geq \frac{\sum |E(S_a^v, \bar{S_a^v})|}{2} \geq \frac{\phi_0 |S|}{2} \geq \frac{\mu\phi_0}{4}|E| \geq \frac{\mu\pi_0}{4(d+1)}|E'|$$

Thus, $UNSAT(G') \geq UNSAT(G)\frac{\mu\phi_0}{4(d+1)}$

In either case, the transformation results in at most a constant factor drop in the fraction of violated edge constraints.

The graph G is thus converted into a constant degree $(d+1)$ graph G'.

Expanderizing the graph The transformed graph G' is $(d+1)$-regular. We superimpose with a $\tilde{d}$-regular expander on $|V'|$ nodes (i.e, th new superimposed graph has the same vertex set as the original constraint edges, its edges are however the union of the two graphs – the original constraint graph and the expander). Furthermore, we add self-loops for each vertex to get G''. We then impose dummy constraints on the new edges (i.e., constraint that are always satisfied). G'' is still an expander (with constant degree, $(d+2+\tilde{d})$), but with slightly weaker spectral expansion given as follows: (this calculation uses Lemma 9.5)

$$\lambda(G'') \leq 1 - \frac{\phi^2(G'')}{2(d+\tilde{d}+2)} \leq 1 - \frac{\phi^2(G_{|V'|})}{2(d+\tilde{d}+2)} \leq 1 - \frac{\phi_0^2}{2(d+\tilde{d}+2)} = \lambda(\text{say}).$$

Observe that if G' is satisfiable, so is G''.

$$UNSAT(G') = \mu \Rightarrow UNSAT(G'') = \mu\left(\frac{d+1}{d+2+\tilde{d}}\right)$$

Thus, $G = (V, E)$ is converted into a constant-degree $\Delta = (d + \tilde{d} + 2)$ expander graph $G'' = (V'', E'')$ with spectral expansion λ. This completes the preprocessing step.

Hence, we will assume without loss of generality that the given constrain graph $G = (V, E)$ is Δ-regular and has spectral expansion λ.

9.4.2 Graph Powering

Due to the preprocessing step, we can assume without loss of generality that the given constrain graph $G = (V, E)$ is Δ-regular expander graph with self loops and has spectral expansion λ. Suppose the unsatisfiability factor of this graph is α. In other words, every coloring of the graph violates at least α-fraction of the edge-constraints. We need to double this factor to 2α without blowing up the size of the graph too much. A natural way to do this is the powering operation. i.e., we build a new *powered graph* G^t on the same set of vertices, each edge of which corresponds to a walk of length t in the original graph.

More formally, the power graph $G^t = ((V', E'), \Sigma', \mathcal{C}')$ for some parameter $t \in \mathbb{Z}^{\geq 0}$ is defined as follows.

- $V' = V$ i.e., the set of vertices is the same.

- $(u, v) \in E'$ iff there exists a walk of length t between the vertices u and v in the original graph G. I.e., there exist vertices $v_0, v_1, \ldots, v_t \in V$ such that $v_0 = u$, $v_t = v$ and $(v_{i-1}, v_i) \in E$ for $i = 1, \ldots, t$. Thus, every edge in E' is a t-walk in G. To distinguish between edges of G and G^t, we will refer to the edges of G as edges and the edges of G^t as t-walks or (just) walks.
- $\Sigma' = \Sigma^{\Delta^{\lceil t/2 \rceil}}$. More precisely, the color $\sigma' \in \Sigma'$ of any vertex v gives not only the color of the vertex v, but also v's opinion of the colors of all vertices which are reachable from v by a walk of length at most $t/2$. Note, such a coloring allows for the case that two (different) vertices u and v could have different opinions about the color of some other vertex which is within distance $t/2$ of both vertices u and v.
- The set of constraints $\mathcal{C}' = \{c_w : \Sigma' \times \Sigma' \to \{0,1\} | w \in E'\}$ is defined as follows: For any t-walk $w = (u, v) \in E'$, the constraint c_w checks that the opinion of the vertices v and u agree on their intersection and that they satisfy all the edge constraints along the walk (u, v). We call each such constraint a walk-constraint.

Below we give a very informal (and in fact wrong) argument as to why powering should amplify the unsatisfiability factor. Let us make two egregious assumptions. Suppose the coloring σ' to the vertices of $G' = (V', E')$ satisfies the property, that if two vertices have an opinion about a third vertex, then their opinions are consistent. Furthermore, let us assume that the edges along a random t-walk appear like t random edges. Then, the following calculations show that the fraction of violated constraints in the powered graph increases by a factor of t. Since by assumption one, the coloring $\sigma : V \to \Sigma'$ is consistent across the vertices, we have that σ' is actually derived from a coloring $\sigma : V \to \Sigma$ of the original constraint graph. However, we know that σ violates at least α-fraction of the edge-constraints in the original graph G. Consider any such edge. This edge occurs in exactly $t\Delta^{t-1}$ walks of the powered graph and the constraints corresponding to each of these walks is violated. Hence, the fraction of violated walk-constraints in G' is $\approx \frac{t\Delta^{t-1}\alpha|E|}{|E'|} \approx \frac{t\Delta^{t-1}\alpha|E|}{\Delta^{t-1}|E|} = t \cdot \alpha$. Thus the unsatisfiability factor increases t-fold. This argument is wrong as both the assumptions are wrong. We won't be able to completely get over the first assumption. However, we will be able to show that if a violated edge occurs in the middle of a walk (ie., between $t/2 - \sqrt{t}$ and $t/2 + \sqrt{t}$), then it is very likely that the constraint on the walk is also violated. Regarding the second assumption, we will use the fact that the underlying graph is an expander and hence the set of edges along a walk do in "some sense" look random (See Lemma 9.6). Using both these we will prove the following lemma which shows that the unsatisfiability factor is amplified $\sqrt{t}$-fold instead of t-fold as in fallacious argument above.

Lemma 9.7 *There exists $\beta > 0$ such that if $UNSAT(G) \leq 1/\sqrt{t}$, then $UNSAT(G') \geq \beta\sqrt{t}\, UNSAT(G)$.*

Proof:

Let $\sigma' : V \to \Sigma'$ be the best possible coloring satisfying the most number of walk-constraints on G'. Hence, $\alpha = UNSAT(G')$ is exactly the fraction of walk-constraints violated by σ'. From the coloring σ', we build a coloring $\sigma : V \to \Sigma$ for the original constraint graph G as follows: Define the random variable $X_{v,i}$ to be the opinion that a vertex which is i random steps away from v has about v. Let $\sigma(v) = a$ such that $\Pr[X_{v,t/2} = a]$ is maximized. In other words, this is the most popular color for v assigned by vertices which are at a distance $t/2$ far from v. By definition of popularity, we have that if $\sigma(v) = a$, then $\Pr[X_{v,t/2} = a] \geq \frac{1}{|\Sigma|}$.

Let B be the set of edges violated by σ in G. Since σ can color no better than the best coloring for G, we have $\frac{|B|}{|E|} \geq UNSAT(G) = \alpha$.

Consider any edge $e = (u, v) \in B$. Let $i \in I = [t/2 - \sqrt{t}, t/2 + \sqrt{t}]$. Consider a random t-walk $w = (v_0, v_1, \ldots, v_t)$ in G' conditioned on the fact that e is the i^{th} edge of the walk (i.e, $u = v_{i-1}$ and $v = v_i$). We will now analyze the probability that the coloring assigned by σ' to the end-vertices of the walk v_0 and v_t violates the walk-constraint c_w.

Let $\sigma(u) = a$ and $\sigma(v) = b$. We know that the coloring a and b to vertices u and v respectively violates the edge constraint c_e. Now, if the opinion of the color of $u = v_{i-1}$ held by the σ'-color of v_0 is a and that of $v = v_i$ held by the σ'-color of v_t is b, then the walk constraint c_w is violated. These events are $X_{u,i-1} = a$ and $X_{v,t-i} = b$. Hence,

$$\Pr_{w:\ (u,v)\text{ is }i\text{th edge of }w}\big[c_w \text{ is violated}\big] \quad\geq\quad \Pr[X_{u,i-1} = a] \cdot \Pr[X_{v,t-i} = b] \tag{1}$$

For simplicity, let us consider the case that $i - 1 = t/2$ and $t - i = t/2$. In other words, e is exactly the middle edge. This case actually does not arise since it assumes the walk is of length $t + 1$ as opposed to t. In this case, we have that $\Pr[X_{u,i-1} = a] = \Pr[X_{u,t/2} = a]$ which is at least $1/|\Sigma|$ since $\sigma(u) = a$ which implies a is the most popular color assigned to u by σ'. Similarly, $\Pr[X_{v,t-i} = b] \geq 1/|\Sigma|$. Thus, in this case we have that

$$\Pr_{w:\ (u,v)\text{ is }i\text{th edge of }w}\big[c_w \text{ is violated}\big] \quad\geq\quad \frac{1}{|\Sigma|^2}$$

We now have to consider the more general case when e is the i^{th} edge for some $i \in I$. Note that this i satisfies the property that $|i - t/2| \leq \sqrt{t}$, i.e., it is within one standard deviation of the mean. The general idea is that since i is at most one standard deviation of the mean, the behavior at i is similar up to certain constant factors to that at the mean. More formally, it can be shown that

$$\text{There exists } \tau > 0 \text{ such that if } |l - t/2| \leq \sqrt{t}, \text{ then } Pr[X_{u,l} = a] \geq \tau \cdot \Pr[X_{u,t/w} = a] \tag{2}$$

This is proved by considering the random walk at u using properties of the binomial distribution. The main intuition is that self loops of G make the distribution of vertices reached by a random $t/2$-step walk from u roughly the same as the distribution on vertices reached by an l-step from u, for $l \in I$. This argument can be formalized; for want of time, we do not present the proof in lecture.

Thus, it follows from (1) and (2) that for all $i \in I$,

$$\Pr_{w:\ (u,v)\text{ is }i\text{th edge of }w}\big[c_w \text{ is violated}\big] \geq \left(\frac{\tau}{|\Sigma|}\right)^2 = \mu \text{ (say)} \tag{3}$$

So far we have shown that an edge is bad, then a constant fraction of the walks in which it occurs nearly in the middle are also bad. We will now show that these bad walks do not overlap too much and hence there is not too much of over-counting. For this purpose, we define the random variable N. Let w be a random t-walk in the powered graph G' (ie., chosen by starting at a random vertex and walking t random steps),

$$N = \begin{cases} \text{Number of bad edges in } I \text{ if } w \text{ is a rejecting } t\text{-walk according to } \sigma' \\ 0 \text{ otherwise} \end{cases}$$

i.e., N is the number of bad edges encountered along the walk w around the middle of the walk if any bad edges are encountered at all and is 0 otherwise.

Clearly this definition of N satisfies,

$$UNSAT(G') \geq \Pr[N > 0].$$

So it suffices to lower bound $\Pr[N > 0]$. We do so using the following two claims.

Claim 9.8 *There exists $\mu > 0$ such that $E[N] \geq 2\mu\sqrt{t}\frac{|B|}{|E|}$.*

Claim 9.9 *There exists $C > 0$ such that $E[N^2] \leq C\sqrt{t}\frac{|B|}{|E|}$.*

We can now bound $\Pr[N > 0]$ using the second moments inequality as follows:

$$\begin{aligned} \Pr[N > 0] &\geq \frac{(E[N])^2}{E[N^2]} \\ &\geq \frac{4\mu^2}{C} \cdot \sqrt{t} \cdot \frac{|B|}{|E|} \end{aligned}$$

Choosing $\beta = \frac{4\mu^2}{C}$ completes the proof of Lemma 9.7. ■

We now need to prove Claims 9.8 and 9.9. For this purpose, we define the following two random variables. For a random walk w, let

$$\begin{aligned} Z_i &= \begin{cases} 1 & \text{if } i^{\text{th}} \text{ edge of } w \text{ is in } B \\ 0 & \text{otherwise} \end{cases} \\ Y_i &= \begin{cases} 1 & \text{if } w \text{ is a rejecting } t\text{-walk and } i^{\text{th}} \text{ edge } \in B \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Observe that $Y_i \leq Z_i$ always.

Proof of Claim 9.8: We observe that $N = \sum_{i \in I} Y_i$. Hence,

$$\begin{aligned} E[N] &= \sum_{i \in I} E[Y_i] \\ &= \sum_{i \in I} \Pr[Y_i = 1] \\ &= \sum_{i \in I} \Pr[Y_i = 1 | Z_i = 1] \cdot \Pr[Z_i = 1] \end{aligned}$$

$\Pr[Z_i = 1]$ is just the probability that the i^{th} edge of a random walk is in B and is thus $|B|/|E|$. $\Pr[Y_i = 1|Z_i = 1]$ is the probability that the random walk w violates the constraint c_w conditioned on the fact that the i^{th} edge is in B. This is precisely the probability calculated in (3). Hence,

$$\begin{aligned} E[N] &\geq \sum_{i\in I} \mu \cdot \frac{|B|}{|E|} \\ &= 2\mu\sqrt{t} \cdot \frac{|B|}{|E|} \end{aligned}$$

■

Proof of Claim 9.9: Using $Y_i \leq Z_i$, we have $N \leq \sum_{i\in I} Z_i$. Hence,

$$\begin{aligned} E[N^2] &\leq E\left[\left(\sum_{i\in I} Z_i\right)^2\right] \\ &= 2\sum_{i\in I}\sum_{j\in I, j\geq i} E\left[Z_i Z_j\right] \\ &= 2\sum_{i\in I}\Pr[Z_i = 1]\left(\sum_{j\in I, j\geq i}\Pr[Z_j = 1|Z_i = 1]\right) \\ &= \frac{2|B|}{|E|}\sum_{i\in I}\sum_{j\in I, j\geq i}\Pr[Z_j = 1|Z_i = 1] \end{aligned}$$

The probability $\Pr[Z_j = 1|Z_i = 1]$ is precisely the probability that a random walk has its $(j-i+1)$th edge in B conditioned on the fact that the first edge of the walk is in B. Here, we use the fact that G is an expander and use Lemma 9.6. We thus, have

$$\begin{aligned} E[N^2] &\leq \frac{2|B|}{|E|}\sum_{i\in I}\sum_{j\in I, j\geq i}\left(\frac{|B|}{|E|} + \lambda^{j-i-1}\right) \\ &\leq C\sqrt{t}\frac{|B|}{|E|} \qquad \text{since } \frac{|B|}{|E|} \leq \frac{1}{\sqrt{t}} \end{aligned}$$

where C is some constant. ■

This almost completes the proof of the Gap Amplification Lemma, modulo one fact – the size of the color set has expanded fro $|\Sigma|$ to $\left|\Sigma^{\lceil \Delta^{t/2} \rceil}\right|$. We now, perform what is known as Proof Composition, to reduce the size of the color set. Proof Composition is a standard procedure (introduced by Arora and Safra [AS]), to reduce the size of the alphabet while only mildly worsening other parameters. After proof composition, we have a similar statement to Lemma 9.7 only with a smaller β, but the color set being the same Σ.

Thus, by choosing an appropriate constant t, we obtain the Gap Amplification Lemma, from which the proof of the PCP theorem follows.

References

[ALMSS] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, Mario Szegedy. "Proof verification and the hardness of approximation problems". Journal of the ACM 45(3):501–555, 1998.

[AS] Sanjeev Arora, Shmuel Safra. "Probabilistic Checking of Proofs: A New Characterization of NP". Journal of ACM, 45(1):70–122, 1998.

[Din] Irit Dinur. "The PCP Theorem by Gap Amplification". ECCC Technical Report TR05-046, 2005.

[FGLSS] Uriel Feige, Shafi Goldwasser, Lázló Lovász, Shmuel Safra, Mario Szegedy. "Interactive proofs and the hardness of approximating cliques". Journal of ACM, 43(2):268–292, 1996.