CMSC 39600: PCPs, codes and inapproximability Oct 11, 16, & 18, 2007

Lecture 6,7,8: Dinur's Proof of the PCP Theorem

Lecturer: Prahladh Harsha *Scribe: Prahladh Harsha*

In the next three lectures, we will present Dinur's proof of the PCP Theorem. Let us first recall that the statement of the PCP Theorem.

Theorem 6,7,8.1 (PCP Theorem)**.** *There exists $\epsilon \in (0,1)$ and constants $Q, C \in \mathbb{Z}^{\geq 0}$ such that*

$$\textit{3-COLOR} \in PCP_{1,1-\epsilon}[C \log n, Q].$$

It will be convenient to rephrase the PCP Theorem in terms of a 2-query PCP since then we can view the PCP as a constraint graph (explained below). However, as long as the alphabet is $\{0,1\}$, any PCP for NP requires at least three (3) queries unless $NP = P$ (see Problem Set 1). For this purpose, we work over a slightly larger, but constant sized alphabet.

Theorem 6,7,8.2 (PCP Theorem (restated))**.** *There exists $\epsilon \in (0,1)$ and constant $C \in \mathbb{Z}^{\geq 0}$ and an alphabet Σ (with $|\Sigma| > 2$) such that*

$$\textit{3-COLOR} \in PCP^{\Sigma}_{1,1-\epsilon}[C \log n, 2].$$

As mentioned earlier, it will be convenient to view a 2-query PCP as a constraint graph.

Definition 6,7,8.3 (Constraint Graph ($\mathcal{CG}^{\Sigma}$))**.** *For an alphabet Σ, an instance of $\mathcal{CG}^{\Sigma}$ is of the form $G = ((V,E), \mathcal{C})$ where (V,E) is an undirected graph, 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 $\pi : V \to \Sigma$ 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) if $c_{(v_1,v_2)}(c_1,c_2) = 1$. The main problem is to find an assignment $\pi : V \to \Sigma$ that satisfies all the coloring constraints.*

We will denote the size $size(G)$ of such an instance by the number of edges $|E|$[1].

Given an instance $G = ((V,E), \Sigma, \mathcal{C})$ and a coloring $\pi : V \to \Sigma$, let

$$UNSAT_{\pi}(G) = \frac{|\{(u,v) \in E | c_{(u,v)}(\pi(u), \pi(v)) = 0\}|}{|E|}.$$

In other words, $UNSAT_{\pi}(G)$ is the fraction of edges violated by π. Let $UNSAT(G)$ denote the minimum $UNSAT_{\pi}(G)$ over all assignments $\pi : V \to \Sigma$.

The gap problem corresponding to constraint graph is as follows:

Definition 6,7,8.4. *For any $0 < \beta < \alpha < 1$, the gap problem gap-$\mathcal{CG}^{\Sigma}_{\alpha,\beta}$ has instances of the form $G = ((V,E), \mathcal{C})$ and its YES and NO instances are as defined below.*

$$\begin{aligned} YES &= \{G | UNSAT(G) \leq 1 - \alpha\} \\ NO &= \{G | UNSAT(G) \geq 1 - \beta\} \end{aligned}$$

[1] To be exact, the size of the instance description is $O(|E| \log |V|)$ where the constant hidden in the $O(\cdot)$ notation depends on the alphabet Σ, but we will ignore the logarithmic dependence.

Observe that an equivalent formulation of the PCP Theorem, in the language of constraint graphs, is the following:

$$\exists \epsilon \in (0,1) \text{ and an alphabet } \Sigma \text{ such that gap-}\mathcal{C}G^{\Sigma}_{1,1-\epsilon} \text{ is NP-hard.}$$

Furthermore, it trivially follows from the fact that 3-COLOR is NP-hard that gap-$\mathcal{C}G^{\Sigma}_{1,1-\frac{1}{n^2}}$ is NP-hard or equivalently 3-COLOR $\in \text{PCP}^{\Sigma}_{1,1-\frac{1}{n^2}}[O(\log n), 2]$ for any alphabet Σ that contains at least 3 symbols (i.e., $|\Sigma| \geq 3$).

6,7,8.1 Gap Amplification

We observed above, that to prove the PCP Theorem, we essentially need to improve the gap between the YES and NO instance from $1/n^2$ to a constant ϵ. The gap amplification lemma essentially does this in several stages.

Lemma 6,7,8.5 (Gap Amplification Lemma [Din07])**.** *There exists an alphabet* Σ *and* $\alpha \in (0,1)$ *such that*

$$PCP^{\Sigma}_{1,1-\epsilon}[r,q] \subseteq PCP^{\Sigma}_{1,1-\epsilon'}[r+O(1),2],$$

where $\epsilon' = \min\{2\epsilon, \alpha\}$*.*

OR (equivalently)

There exists a polynomial time reduction from gap-$\mathcal{C}G^{\Sigma}_{1,1-\epsilon}$ *to gap-*$\mathcal{C}G^{\Sigma}_{1,1-\epsilon'}$ *where* ϵ' *is as defined above. Furthermore, this reduction satisfies that the size of the output instance is at most linear in the size of the input instance.*

The gap amplification lemma states that the gap can be increased from ϵ to 2ϵ as long as the gap is not already a constant α at the cost of a constant increase in the randomness. We can now prove the PCP Theorem starting from the Gap Amplification Lemma.

Proof of PCP Theorem. 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 $1/n^2$ to α with at most a polynomial blowup in size, thus proving the PCP Theorem. $\square$

Thus it suffices for us to prove the Gap Amplification Lemma 6,7,8.5 and this will be our goal for the next three lectures.

First, let us recall some techniques to improve the gap that we have already seen. The first is sequential repetition (a closely related one is parallel repetition, more of which we will see later in this course). Both these types of repetition dramatically improve the gap, however at a phenomenal cost in the randomness. In fact, repeating the PCP $O(t)$ times, improves the gap from ϵ to $1-(1-\epsilon)^t$. More precisely, if $UNSAT(G) \geq \epsilon$ (i.e., for every coloring $\pi : V \to \Sigma$, we have

$$\Pr_{e=(u,v)}[c_e(\pi(u),\pi(v)) \neq 1] \geq \epsilon,$$

then repeating this t times we have

$$\Pr_{e_1=(u_1,v_1),\ldots,e_t=(u_t,v_t)}[\exists i, c_e(\pi(u),\pi(v)) \neq 1] \geq 1-(1-\epsilon)^t.$$

Thus, this improves the gap considerably however at the cost of an $O(t)$ blowup in the randomness. This increase in randomness is something we cannot afford. Is there a randomness efficient way of attaining the same improvement in gap? Dinur's key observation is that this gap improvement can be achieved if the underlying constraint graph is a constant-degree expander and the t edges are chosen by a random walk of length t along the expander. However, there is no guarantee that the underlying constraint graph to is a constant-degree expander. This necessitates an initial preprocessing phase in which we massage the constraint graph into a constant-degree graph. Furthermore, as we will see later, the PCP constructed by taking a length t walk has an exponential size in blowup. Recall, that we wanted both the input and output alphabet in the gap amplification lemma to be the same. However, proof composition is tailor-made to handle alphabet reduction. We can thus apply proof composition at the end to reduce the alphabet back to the original size. Thus, the proof of the gap amplification lemma involves the following phases.

Phase I: preprocessing In this phase, we transform the underlying constraint graph into a constant degree expander at the cost of a constant deterioration in the gap and a constant additive increase in the randomness.

Phase II: graph powering In this phase, we construct a new PCP by performing a t-length random walk on the constraint graph (which is guaranteed to be an expander by phase I). This phase dramatically increases the gap at the cost of an additive increase in randomness. However, the alphabet would have exponentially blown up from Σ to Σ^{d^t} where d is the degree of the underlying expander.

Phase III: alphabet reduction In the final phase, we reduce the alphabet back to Σ by proof composition. This stage accounts for a further constant additive increase in randomness and deteriorates the gap by a constant factor.

The deterioration in the gap in phases I and III are more than compensated by the improvement achieved in phase II. We will look at these phases in detail in the next few lectures. But, now we need some preliminaries regarding expanders.

6,7,8.2 Expanders – Preliminaries

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

$$\phi(S) = \frac{E(S,\bar{S})}{|S|}$$

$$\phi(G) = \min_{S\subseteq V,|S|\leq n/2} \phi(S)$$

We will be interested in graphs G whose edge expansion $\phi(G)$ is large (at least a constant). Such graphs are very well-connected and have no "bottlenecks". A complete graph for instance has edge expansion $\phi(G) \geq \Omega(n)$. However, we would be interested in sparse graphs, specifically d-regular graphs for some constant d.

It will be more convenient to work with the spectral notion of expansion. For any d-regular graph, let A_G denote the adjacency matrix (i.e., $(A_G)_{i,j} = 1$ if $(i,j) \in E$ and 0 otherwise). In the case multi-graphs, $(A_G)_{i,j}$ refers to the number of edges between i and j, while for weighted graphs it refers to the weight of edge (i,j). Let us look at the eigenvectors an eigenvalues of this matrix A_G. Recall that v is an eigenvector for matrix A with eigenvalue λ if $Av = \lambda v$. Since A_G is a real symmetric matrix, it has n real-valued eigenvalues, say $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. Note that the all ones vector $\mathbf{1} = (1,1,\ldots,1)$ is an eigenvector of A_G with eigenvalue d.

It is an easy exercise to check that d is in fact, the largest eigenvalue. Furthermore, the number of eigenvalues that are equal to d is precisely the number of distinct components of G. Also, the smallest eigenvalue λ_n is $-d$ iff graph is bipartite. If the graph is connected and non-bipartite, then the remaining eigenvalues are bounded in absolute value by d. This will be our alternate definition of expanders.

Definition 6,7,8.6. *A d-regular graph G on n nodes is an (n,d,λ)-expander if $\lambda = \max_i\{|\lambda_i|\} = \max\{\lambda_2, |\lambda_n|\}$ is strictly bounded above by d (i.e., $\lambda < d$). We refer to this gap $d - \lambda$ as the spectral gap.*

The following theorem shows the close relationship between the two definitions of expanders (i.e., edge expansion $\phi(G)$ and the spectral gap $d - \lambda$).

Theorem 6,7,8.7. *For a d-regular graph G,*

$$\frac{\phi^2(G)}{2} \leq d - \lambda \leq 2\phi(G).$$

Thus, if we a graph with good spectral expansion, it also has good edge expansion.

For the purpose of the proof of the gap amplification, we will assume the existence of expanders for every $n > 0$. But for this one fact, our proof of the gap amplification lemma, will be a self-contained one.

Theorem 6,7,8.8. *There exist constant d and λ with $\lambda < d$ and an explicit family of (n,d,λ')-expanders for every n with $\lambda' < \lambda$.*

Below we give some examples (without proof) of explicit expanders with a good spectral gap (and hence good edge expansion).

1. **Margulis/Gaber-Galil Expanders** These expanders have vertex set of size n^2 that we identify with the set $\mathbb{Z}_n \times \mathbb{Z}_n$ where $\mathbb{Z}_n$ is the ring of integers modulo n (to be precise, we need to use the notation $\mathbb{Z}/n\mathbb{Z}$ as $\mathbb{Z}_p$ refers to p-adic numbers, but anyway...). The vertex given by (x,y) is connected to vertices

$$(x+2y, y)(x, 2x+y)$$
$$(x+2y+1, y)(x, 2x+y+1)$$

   where all additions are done modulo n. We also add the edges corresponding to the inverse transformation. Thus, this graph is a 8-regular graph on n^2 vertices (possibly has self-loops and multiple edges). One can show that this graph satisfies $\lambda \leq 5\sqrt{2} < 8$.

2. **Lubotsky-Phillips-Sarnak (LPS) Expanders** The LPS graph has vertex set $V = \mathbb{F}_p \cup \{\infty\}$ where p is a prime. $\mathbb{F}_p$ is the finite field of the set of integers modulo p. We extend addition and multiplication in this field to the special point ∞ as follows: $0 \cdot \infty = 1, x + \infty = \infty, z \cdot \infty = \infty$ for all $x \in \mathbb{F}_p$ and $z \in \mathbb{F}_p^* = \mathbb{F}_p \setminus \{0\}$. The vertex x is then connected to the vertices x^{-1}, $x+1$ and $x-1$. Observe that this graph is 3-regular graph. In fact, this graph has a very simple structure. If one ignores the point ∞ it looks like a cycle with a set of matching edges added in. It is known that for this graph $\lambda < 3$. This graph is a considerable simplification of the original LPS graphs, which are known to have the best spectral gap (such graphs are called Ramanujan graphs).

In the next section, we prove some properties of expanders that we require in the proof of the gap amplification lemma.

6,7,8.3 Expanders – properties

The λ of a graph has the following nice characterization.

$$\lambda = \max_{x \in \mathbb{R}^n \setminus \{\mathbf{0}\}, x \cdot \mathbf{1} = 0} \frac{|\langle A_G x, x\rangle|}{\langle x, x\rangle} \tag{1}$$

This is also known as the Rayleigh quotient of the matrix A_G. This characterization has an easy proof. Since A_G is a real symmetric matrix, the eigen vectors $\mathbf{1} = v_1, v_2, \ldots, v_n$ form an orthonormal basis for the space $\mathbb{R}^n$. Any vector $x \in \mathbb{R}^n$, can thus be written as $x = \sum a_i v_i$. If $x \cdot \mathbf{1}$, we have $x = \sum_{i=2}^n a_i v_i$. Now, $A_G x$ is given by $A_G x = \sum_{i=2}^n \lambda_i a_i v_i$. Thus, $\langle A_G x, x\rangle = \sum_{i=2}^n a_i^2 \lambda_i$ which is at most $\lambda \sum_{i=2} n a_i^2 = \lambda \langle x, x\rangle$. Thus, the Rayleigh quotient is at most λ. For the other direction, observe that $\langle A_G v_2, v_2\rangle = \lambda_2 \langle v_2, v_2\rangle$ and $\langle A_G v_n, v_n\rangle = \lambda_n \langle v_n, v_n\rangle$. Hence, the Rayleigh quotient is at least λ.

We can now prove the second inequality in Theorem 6,7,8.7, which is the direction we would need in the preprocessing phase.

Lemma 6,7,8.9. *For a d-regular graph G, $d - \lambda \leq 2\phi(G)$.*

Proof. Let $S \subset V$ be any subset of vertices such that $|S| \leq n/2$. We need to show that $\phi(S) \geq (d - \lambda)/2$. We will prove this by using the Rayleigh quotient characterization of λ.

Consider the vector x defined as follows:

$$x_v = \begin{cases} -|\bar{S}| \text{ if } v \in S \\ |S| \text{ if } v \in \bar{S} \end{cases}$$

It is easy to check that $\|x\|^2 = |S||\bar{S}|^2 + |\bar{S}||S|^2 = |S||\bar{S}|n$. Let us calculate $\langle A_G x \rangle$.

$$
\begin{aligned}
\langle A_G x, x\rangle &= \sum_u x_u \sum_{(u,v)\in E} x_v \\
&= 2 \sum_{(u,v)\in E} x_u x_v \\
&= 2 \sum_{(u,v)\in E(S)} x_u x_v + 2 \sum_{(u,v)\in E(\bar{S})} x_u x_v + 2 \sum_{(u,v)\in E(S,\bar{S})} x_u x_v \\
&= (d|S| - |E(S,\bar{S})|) \cdot |\bar{S}|^2 + (d|\bar{S}| - |E(S,\bar{S})|) \cdot |S|^2 - 2|E(S,\bar{S})| \cdot |S| \cdot |\bar{S}| \\
&= d|S||\bar{S}|n - |E(S,\bar{S})|n^2
\end{aligned}
$$

Since $x \cdot \mathbf{1} = 0$, we have $\langle A_G x, x\rangle \le \lambda \|x\|^2$. Rearranging the terms, we obtain $(d-\lambda)/n \le |E(S,\bar{S}|/(|S||\bar{S}|)$. However, since $S \le n/2$, we have $|\bar{S}| \ge n/2$. Hence, $(d-\lambda)/2 \le |E(S,\bar{S}|/|S|$. Thus proved.

□

Recall that in the preprocessing phase, we need to massage the constraint graph into an expander. A simple way to convert any graph into an expander is to superimpose an expander over the given graph. The following claim that if two graphs are superimposed over one another, then the spectral gap of the final graph is at least the spectral expansion of each of the original graphs.

Lemma 6,7,8.10. *If G and H are two regular graphs on the same set of vertices V with degrees d and d' respectively,then $G' = G \cup H = (V, E(G) \cup E(H))$ is a $(d+d')$-regular graph satisfying*

$$\lambda(G \cup H) \le \lambda(G) + \lambda(H).$$

Proof. This lemma is an easy consequence of the Rayleigh quotient characterization of λ. We observe that $A_{G\cup H} = A_G + A_H$. Hence,

$$
\begin{aligned}
\lambda(G \cup H) &= \max_{x:\|x\|=1, x\cdot\mathbf{1}=0} \langle A_{G\cup H} x, x\rangle \\
&= \max_{x:\|x\|=1, x\cdot\mathbf{1}=0} (\langle A_G x, x\rangle + \langle A_H x, x\rangle) \\
&\le \max_{x:\|x\|=1, x\cdot\mathbf{1}=0} \langle A_G x, x\rangle + \max_{x:\|x\|=1, x\cdot\mathbf{1}=0} \langle A_H x, x\rangle \\
&= \lambda(G) + \lambda(H)
\end{aligned}
$$

□

We also require the following estimate on the random-like behavior of a random walk on an expander, which will be used in phase II.

Lemma 6,7,8.11. *Let $G = (V, E)$ be a (n, d, λ)-expander. Let $F \subset E$ be a set of edges. The probability p that a random walk that starts at a random edge in F takes its $(i+1)^{\text{st}}$ step in F as well, is bounded above by $\frac{|F|}{|E|} + \left(\frac{\lambda}{d}\right)^i$.*

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 $\frac{|F|}{|E|}$. The above lemma states that choosing the edges according to a random walk worsens this probability by at most $(\lambda/d)^i$.

The proof of this lemma is similar to that of the Expander Mixing Lemma. This proof is reproduced verbatim from Dinur's paper [Din07].

Proof. Let π be the distribution on vertices of G induced by selecting a random edge in F and then a random vertex on which the edge is incident on. Let W be the support of the distribution π. Hence, π_v is the fraction of edges incident on v that are in F, divided by 2. For any vertex v, let F_v denote the set of edges incident on v that are in F. Hence, $\pi_v = |F_v|/2|F| \leq d/2|F|$ since G is d-regular. Let y_v be the probability that a random step from v is in F, so $y_v = |F_v|/d = 2|F|\pi_v/d$. Or equivalently $y = (2|F|/d)\pi$.

Let A be the normalized adjacency matrix of G. The probability p equals the probability of landing in W after i steps and then taking a step in F. Hence

$$p = \sum_{v \in W} y_v (A^i \pi)_v = \sum_{v \in V} y_v (A^i \pi)_v = \langle A^i \pi, y \rangle = \frac{2|F|\langle A^i \pi, \pi\rangle}{d}.$$

let $\mathbf{1}$ be all ones vector. Decomposing π along u and its orthogonal component we have $\pi = \pi^{\parallel} + \pi^{\perp}$. Observe that

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

Since G is a (n, d, λ)-expander,

$$\begin{aligned} \|A^i \pi^{\perp}\|_2 &\leq \left(\frac{\lambda}{d}\right)^i \|\pi^{\perp}\|_2 \\ &\leq \left(\frac{\lambda}{d}\right)^i \|\pi\|_2 \end{aligned}$$

By Cauchy-Schwarz,

$$\langle A^i \pi^{\perp}, \pi \rangle \leq \|A^i \pi^{\perp}\|_2 \|\pi\|_2 \leq \left(\frac{\lambda}{d}\right)^i \|\pi\|_2^2$$

Combining we have,

$$\begin{aligned} p = \langle A^i \pi, y\rangle \quad &= \frac{2|F|\langle A^i \pi, \pi\rangle}{d} &&= \frac{2|F|}{d}\left(\langle A^i \pi^{\parallel}, \pi\rangle + \langle A^i \pi^{\perp}, \pi\rangle\right) \\ \leq \frac{2|F|}{d}\left(\frac{1}{n} + \left(\frac{\lambda}{d}\right)^i \|\pi\|_2\right) \quad &\leq \frac{2|F|}{d}\left(\frac{1}{n} + \frac{d}{2|F|}\left(\frac{\lambda}{d}\right)^i\right) &&= \frac{|F|}{|E|} + \left(\frac{\lambda}{d}\right)^i. \end{aligned}$$

$\square$

6,7,8.4 Proof of Gap Amplification Lemma

We are now ready to prove the gap amplification lemma. As indicated before, we will proceed in three phases - preprocessing, graph powering and alphabet reduction. We describe each of these phases in detail in the following sections.

6,7,8.5 Phase I: Graph Preprocessing

The preprocessing step involves converting the underlying constraint 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 $d-1$ and edge expansion at least ϕ_0.

Let $G=(V,E)$ be the underlying constraint graph The 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, where $d-1$ 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.

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

$$|E'| = \frac{d}{2}|V'| = d|E|$$

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 $\mu|E| \le |B| \le |B'| + |S|$.

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

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

- *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,

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

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

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

The constraint graph G is thus converted into a constant degree d graph G'. Or in PCP notation, $\mathrm{PCP}^{\Sigma}_{1,1-\epsilon}[r, 2] \subseteq \mathrm{PCP}^{\Sigma}_{1,1-\delta\epsilon}[r + \log_d, 2]$, where the constraint underlying the second PCP is a d-regular graph and δ is a constant (dependent on d and ϕ_0).

Expanderizing the graph The transformed graph G' is d-regular. We superimpose with a $\tilde{d}$-regular expander E on $|V'|$ nodes (i.e, the 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). We then impose dummy constraints on the new edges (i.e., constraint that are always satisfied). G'' is still an expander (with constant degree, $(d+\tilde{d})$), but with slightly weaker spectral expansion given as follows: (this calculation uses Lemma 6,7,8.10)

$$\lambda(G'') \leq \lambda(G') + \lambda(E) \leq= \lambda(\text{say}).$$

Since, $\lambda < d + \tilde{d}$, the final graph G'' is a good expander.

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

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

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

Equivalently, we have in PCP notation, we have $\mathrm{PCP}^{\Sigma}_{1,1-\epsilon}[r, 2] \subseteq \mathrm{PCP}^{\Sigma}_{1,1-\delta_I\epsilon}[r+O(1), 2]$, where the constraint underlying the second PCP is a $(n, d+\tilde{d}, \lambda)$-expander δ_I is a constant (the deterioration caused by phase I).

6,7,8.6 Phase II: Graph Powering

Due to the preprocessing step, we can assume without loss of generality that the underlying constraint graph G is d-regular expander graph and second eigenvalue expansion $\lambda < d$. In the second phase, we will perform random walks on this expander to improve the gap. More formally, we will show the following.

Lemma 6,7,8.12 (Graph Powering). *Suppose*

$$L \in PCP^{\Sigma}_{1,1-\epsilon}[r, 2]$$

and furthermore, the underlying constraint graph is a d-regular expander with second eigenvalue at most $\lambda < d$, then

$$L \in PCP^{\Sigma'}_{1,1-\epsilon'}[r + O(t \log t), 2]$$

where $\epsilon' = \Omega(t) \min\{\epsilon, \frac{1}{t}\}$ and $\Sigma' = \Sigma^{1+d+d^2+\cdots+d^t}$.

Thus, the gap is improved by a multiplicative factor of t (unless the gap is not already $1/t$) if the underlying constraint graph were a (n, d, λ)-expander. Since the preprocessing phase guarantees that this is indeed the case, combining the two phases we have the following corollary.

Corollary 6,7,8.13 (preprocessing+graph powering).

$$PCP^{\Sigma}_{1,1-\epsilon}[r, 2] \subseteq PCP^{\Sigma'}_{1,1-\epsilon'}[r + O(t \log t), 2]$$

where $\epsilon' = \Omega(t) \cdot \min\{\delta_I \cdot \epsilon, \frac{1}{t}\}$ and $\Sigma' = \Sigma^{1+d+d^2+\cdots+d^t}$.

The presentation in this lecture is different from Dinur's proof [Din07] which uses walks of fixed length. We will instead use lazy random walks along the lines of Radhakrishnan's variant of Dinur's proof [Rad06].

Given a PCP Verifier whose underlying graph is a (n, d, λ)-expander, we will construct a new PCP verifier for this language whose gap is significantly better than the original verifier. We can also construct the equivalent constraint graph of the new PCP verifier but we will not do so here. Before describing the PCP verifier, let us first describe the new alphabet Σ' and the new proof π'.

Alphabet and Proof Recall that the original alphabet Σ was a set of colors and the coloring (or proof) $\pi : V \to \Sigma$ specified the color of each vertex in the constraint graph. The new alphabet Σ' will be given by $\Sigma' = \Sigma^{1+d+d^2+\cdots+d^t}$ and the new coloring (or proof) will be given by a map $\pi' : V \to \Sigma'$. Note that $1 + d + d^2 + \cdots + d^t$ is an upper bound on the number of vertices within distance t from a given vertex since the graph is d-regular. As a result, given a color $\sigma' \in \Sigma'$ for a vertex v, we can identify for each node w in the t-neighborhood of v, a particular position in σ'. We say that the value of this position corresponds to the "opinion" about the color of w in the old constraint graph held by v. Given an assignment $\pi' : V \to \Sigma'$, we write $\pi'(v)_w$ for v's opinion about w's label. Note that if w lies in the t-neighborhood of two vertices u and v, it might not be the case that $\pi'(u)_w = \pi'(v)_w$, that is u and v might disagree on their opinion about w's label.

To describe the verifier, we need to describe the following random walk on the constraint graph.

Lazy Random Walk (LRW)

Input: $v \in V$, a vertex of the graph

1. Set $j \leftarrow 1$ and $v_1 \leftarrow v$
2. With probability $1/t$,

   Stop and output the entire sequence of vertices $(v_1 \to v_2 \to \cdots \to v_j)$ visited.
3. Else
   (a) Choose a random edge (v_j, v_{j+1}) out of vertex v_j
   (b) Set $j \leftarrow j+1$
   (c) Goto Step 2

We are now ready to describe the PCP Verifier that proves Lemma 6,7,8.12. Suppose $L \in \mathrm{PCP}^{\Sigma}_{1,1-\epsilon}[r, 2]$ and furthermore, the underlying constraint graph (V, E) is a d-regular expander with second eigenvalue at most $\lambda < d$. We now describe a new PCP verifier for L that performs a lazy random walk on the original constraint graph and checks all the edge constraints along this walk.

Graph Powering PCP Verifier (GP-Ver)

Input: $G = ((V, E), \mathcal{C})$ - constraint graph such that (V, E) is a (n, d, λ)-expander

Oracle access: Proof $\pi' : V \to \Sigma'$

1. Choose $e = (v_0, v_1) \in_R E$
2. Perform LRW(v_1) to obtain the sequence of vertices $(v_1 \to v_2 \to \cdots \to v_l)$.
3. If $l > B$, accept
4. If $l \leq B$,
   (a) Set $a \leftarrow v_0$ and $b \leftarrow v_l$
   (b) For each $i = 1, \ldots, l$

       Check if edge $e_i = (v_{i-1}, v_i)$ satisfies $c_{e_i}(\pi'(a)_{v_{i-1}}, \pi'(b)_{v_i}) = 1$
   (c) Accept if all the checks pass

We will later set $B = O(t \ln |\Sigma|)$.

GP-Ver chooses a random edge (a, v_1) of the original constraint graph (V, E). It then performs a lazy random walk to obtain the walk $v_1 \to v_2 \to \cdots \to v_l = b$. If this walk is too long (i.e., $l > B$), the verifier accepts. Else it looks at the walk $a = v_0 \to v_1 \to v_2 \to \cdots \to v_l = b$ and checks that for each edge $e_i = (v_{i-1}, v_i)$ along the walk, the opinion $\pi'(a)_{v_{i-1}}$ about the left end point v_{i-1} held by vertex a and the opinion $\pi'(b)_{v_i}$ about the right end point v_i held by vertex b satisfy the constraint c_{e_i}.

Randomness The amount of randomness required by the GP-Ver verifier is exactly the amount of randomness required to choose a random edge (v_0, v_1) in E and perform a lazy random walk of length at most B. The former is exactly r while the latter is $B(\log d + \log t) = O(t \log t)$ ($\log d$ is required to choose a random edge and $\log t$ to flip a coin that comes up head with probability $1/t$).

Completeness Suppose $x \in L$. Then there exists a coloring $\pi : V \to \Sigma$ that satisfies every constraint $c_e \in \mathcal{C}$. Clearly, setting $\pi' : V \to \Sigma'$ such that $\pi'(v)_u = \pi(u)$ for all u and v, causes GP-VER to accept with probability exactly 1.

Soundness Suppose $x \in L$. let $\pi' : V \to \Sigma'$ be the coloring that maximizes the acceptance probability of GP-VER. We need to show that $\Pr[\text{GP-VER}^{\pi'}(G) = 1] \leq 1 - \epsilon'$ where $\epsilon' = O(t)\min\{\epsilon, 1/t\}$. For this purpose, we first construct an assignment $\pi : V \to \Sigma$ of the original constraint graph as follows:

$$\pi(v) = \operatorname*{arg-max}_{\sigma \in \Sigma} \Pr\left[\pi'(w)_v = \sigma | LRW(v) = (v = v_1 \to v_2 \to \cdots \to v_l = w), l \leq t\right] \qquad (2)$$

In other words, the assignment $\pi(v)$ is defined as follows: Perform a lazy random walk starting at v to reach the final vertex w. As long as v is within a t-neighborhood of w, w has an opinion $\pi'(w)_v$ about v's label. Conditioned on the walk being less that t steps, $\pi(v)$ is the most popular opinion held by all such w.

Since $x \in L$, it must be the case that π violates at least ϵ-fraction of the constraints of $\mathcal{C}$. Let F be some ϵ fraction of edges violated by π. For technical reasons (which we encounter later), if $|F|/|E| > 1/t$, throw away edges in F such that $|F|/|E| \leq 1/t$. Thus, $|F|/|E| = \min\{\epsilon, 1/t\}$. The following claim will complete the proof of soundness for us.

Claim 6,7,8.14.

$$\Pr\left[\text{GP-VER}^{\pi'}(G) = \mathsf{rej}\right] = \Omega(t) \cdot \frac{|F|}{|E|}.$$

Consider any step $(u \to v)$ along the path $a = v_0 \to v_1 \to \cdots \to v_l = b$ visited by GP-VER verifier (i.e., $u = v_{i-1}$ and $v = v_i$ for some $1 \leq i \leq l$). We call such a step $(u \to v)$ faulty if the following three conditions are met

- $(u, v) \in F$
- $d_G(u, a) \leq t$ and $\pi'(a)_u = \pi(u)$
- $d_G(v, b) \leq t$ and $\pi'(b)_v = \pi(v)$

Observe that if any step $(v_{i-1} \to v_i)$ is faulty and the length of the walk is at most B, GP-VER rejects since $c_{(v_{i-1},v_i)}(\pi'(a)_{v_{i-1}}, \pi'(b)_{v_i}) = c_{(v_{i-1},v_i)}(\pi(v_{i-1}), \pi(v_i)) = 0$. This leads us to the following definitions: For any random walk $a = v_0 \to v_1 \to \cdots \to v_{l-1} \to v_l = b$, chosen by GP-VER , define random numbers N_F, N, N_* as follows.

$$\begin{aligned} N_F &= \#\{i | (v_{i-1}, v_i) \in F\} \\ N &= \#\{i | (v_{i-1} \to v_i) \text{ is faulty}\} \\ N_* &= \begin{cases} N & \text{if } l \leq B \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Thus, N_F is the number of steps in F, while N is the number of faulty steps and $N_* = N \cdot \mathbb{1}_{l \leq B}$ where $\mathbb{1}_{l \leq B}$ is the indicator random variable for the event "$l \leq B$". Clearly, $N_F \geq N \geq N_*$. Observe that GP-VER rejects if $N_* > 0$. The following claims bound the first and second moments of N_*.

Claim 6,7,8.15.

$$\text{Exp}\left[N_*\right] = \Omega\left(t\right) \cdot \frac{|F|}{|E|}$$

Claim 6,7,8.16.

$$\text{Exp}\left[N_*^2\right] \le \text{Exp}\left[N_F^2\right] = O\left(t\right) \cdot \frac{|F|}{|E|}$$

Before proceeding to prove these claims, we will show how these claims suffice to prove the soundness of the GP-VER verifier.

Proof of Claim 6,7,8.14.

$$\begin{aligned}
\Pr[\text{GP-VER}^{\pi'}(G) = \mathsf{rej}] \geq & \Pr[N_* > 0] \\
\geq & \frac{\text{Exp}[N_*]^2}{\text{Exp}[N_*^2]} && \text{[By Chebyshev-Cantelli's inequality]} \\
= & \left(\Omega\left(t\right) \cdot \frac{|F|}{|E|}\right)^2 \cdot \left(O\left(t\right) \cdot \frac{|F|}{|E|}\right)^{-1} \\
= & \Omega(t) \cdot \frac{|F|}{|E|}
\end{aligned}$$

□

6,7,8.6.1 $\text{Exp}[\mathbf{N}_*] = \mathbf{\Omega(t)|F|/|E|}$

In this section, we will show that the expected value of N_* is at least $\Omega(t)$ times $|F|/|E|$. The hard part will be to show that the expected number of faulty edges (i.e., $\text{Exp}[N]$) along a random walk is at least $\Omega(t)|F|/|E|$. We can then easily convert this bound to a bound on $\text{Exp}[N_*]$ since it is very unlikely that the GP-VER will take walks of length greater than B. To bound $\text{Exp}[N]$, we would need the following proposition about the random walk $a = v_0 \to v_1 \to \cdots \to v_l = b$

Proposition 6,7,8.17. *Fix any edge $(u, v) \in E$. Let walk $= (a = v_0 \to v_1 \to \cdots \to v_{l-1} \to v_l = b)$ be the random walk in the constraint graph taken by the* GP-VER *verifier. Conditioned on the fact that walk contains exactly k $(u \to v)$ steps, the distribution of the endpoints a and b of the walk satisfy the following:*

1. *The distribution of a is precisely that of an endpoint of a lazy random walk originating at u.*
2. *The distribution of b is precisely that of an endpoint of a lazy random walk originating at v.*
3. *a and b are independent.*

Proof. We will first prove (2). Suppose instead of conditioning on the fact that *walk* makes exactly k $(u \to v)$ steps, we had conditioned on *walk* making at least k $(u \to v)$ steps, then (2) is obvious is obvious. In other words, if we set $Y_{u\to v}$ to be the random variable that denotes the number of $(u \to v)$ steps, then conditioned on the event "$Y_{u\to v} \geq k$", the distribution of the right end point b of the walk is precisely that of an endpoint of a lazy random walk originating at v. Hence, for all vertices $w \in V$, the probability $\Pr[b = w|Y_{u\to v} \geq k]$ is independent of k and is precisely the probability $\Pr[b = w|LRW(v) = (v = v_1 \to v_2, \cdots \to v_l = b)]$. Let us call this probability p_w. Note that we are actually interested in the probability $\Pr[b = w|Y_{u\to v} = k]$. If we show this quantity is also p_w, we would be done. Look at the following calculation:

$$\begin{aligned} p_w &= \Pr[b = w|Y_{u\to v} \geq k] \\ &= \frac{\Pr[b = w \wedge Y_{u\to v} \geq k]}{\Pr[Y_{u\to v} \geq k]} \\ &= \frac{\Pr[b = w \wedge Y_{u\to v} = k] + \Pr[b = w \wedge Y_{u\to v} \geq k+1]}{\Pr[Y_{u\to v} = k] + \Pr[Y_{u\to v} \geq k+1]} \\ &= \frac{\Pr[b = w \wedge Y_{u\to v} = k] + p_w \cdot \Pr[Y_{u\to v} \geq k+1]}{\Pr[Y_{u\to v} = k] + \Pr[Y_{u\to v} \geq k+1]} \end{aligned}$$

Hence, it follows that

$$\Pr[b = w|Y_{u\to v} = k] = \frac{\Pr[b = w \wedge Y_{u\to v} = k]}{\Pr[Y_{u\to v} = k]} = p_w$$

This completes the proof of (2).

To prove (1), we observe that the distribution of $walk = (a = v_0 \to v_1 \to \dots tov_{l-1} \to v_l = b$ is identical to that of the reversed walk $walk^R = (b = v_l \to v_{l-1} \to \cdots \to v_1 \to v_0 = a)$. Hence, the distribution of the left endpoint of $walk$ conditioned on having taken exactly k $(u \to v)$ steps is exactly the distribution of the right endpoint of $walk^R$ conditioned on having taken exactly k $(v \to u)$ steps which from the above argument is the distribution of the endpoint of a lazy random walk originating at u.

The above argument also shows that a and b are independent since once the middle of the walk is fixed, the left and right ends are independent.

□

We are now ready to bound $\text{Exp}[N]$.

Claim 6,7,8.18. $\text{Exp}[N] \geq \frac{t}{4|\Sigma|^2} \cdot \frac{|F|}{|E|}$

Proof. Fix some edge $(u, v) \in F$. Let $N_{u\to v}$ be the number of faulty $(u \to v)$ steps along a random walk. Note that $N = 2\sum_{(u,v)\in F} N_{u\to v}$ (the 2 because $(u \to v)$ steps are different from $(v \to u)$ steps). Thus, by linearity of expectation it suffices to prove that $\text{Exp}[N_{u\to v}] \geq \frac{t}{8|\Sigma|^2} \cdot \frac{1}{|E|}$. We have

$$\text{Exp}[N_{u\to v}] = \sum_{k\geq 1} \text{Exp}[N_{u\to v} \mid \text{ exactly } k \ (u \to v) \text{ steps }] \cdot \Pr[\text{ exactly } k \ (u \to v) \text{ steps }]$$

However, if one $(u \to v)$ step along a walk is faulty so are all other $(u \to v)$ steps along the walk. Hence,

$$\text{Exp}[N_{u\to v}] = \sum_{k\geq 1} k \cdot \Pr[(u \to v) \text{ is faulty } \mid \text{ exactly } k \ (u \to v) \text{ steps }] \cdot \Pr[\text{ exactly } k \ (u \to v) \text{ steps }]$$

Now,

$$\begin{aligned} & \Pr[(u \rightarrow v) \text{ is faulty } \mid \text{ exactly } k \ (u \rightarrow v) \text{ steps }] \\ = \ & \Pr[d_G(u,a) \leq t \wedge \pi'(a)_u = \pi(u) \wedge d_G(v,b) \leq t \wedge \pi'(b)_v = \pi(v) \mid \text{ exactly } k \ (u \rightarrow v) \text{ steps }] \end{aligned}$$

Now applying the above proposition which describes the distribution of a and b conditioned on the walk taking exactly k steps we have,

$$\begin{aligned} & \Pr[(u \rightarrow v) \text{ is faulty } \mid \text{ exactly } k \ (u \rightarrow v) \text{ steps }] \\ = \ & \Pr[d_G(u,a) \leq t \wedge \pi'(a)_u = \pi(u) \wedge d_G(v,b) \leq t \wedge \pi'(b)_v = \pi(v) \mid a = \text{e-LRW}(u); b = \text{e-LRW}(v)] \\ = \ & \Pr[d_G(u,a) \leq t \wedge \pi'(a)_u = \pi(u) \mid a = \text{e-LRW}(u)] \cdot \Pr[d_G(v,b) \leq t \wedge \pi'(b)_v = \pi(v) \mid b = \text{e-LRW}(v)] \\ & \qquad \text{[By independence of } a \text{ and } b\text{]} \\ = \ & (\Pr[d_G(u,a) \leq t \wedge \pi'(a)_u = \pi(u) \mid a = \text{e-LRW}(u)])^2 \end{aligned}$$

In the above, we have used e-LRW(v) to denote the endpoint of the lazy random walk originating at v. Observe that $\Pr[\pi'(a)_u = \pi(u) \mid a = \text{e-LRW}(u); l \leq t]$ is at least $1/|\Sigma|$ since $\pi(u)$ is chosen to be the most popular vote among $\pi'(a)_u$ conditioned on $a = \text{e-LRW}(u)$ and $l \leq t$. Using this in the above, we have

$$\begin{aligned} & \Pr[d_G(u,a) \leq t \wedge \pi'(a)_u = \pi(u) \mid a = \text{e-LRW}(u)] \\ \geq \ & \Pr[d_G(u,a) \leq t \wedge \pi'(a)_u = \pi(u) \wedge l \leq t \mid a = \text{e-LRW}(u)] \\ = \ & \Pr[\pi'(a)_u = \pi(u) \wedge l \leq t \mid a = \text{e-LRW}(u)] \\ = \ & \Pr[\pi'(a)_u = \pi(u) \mid l \leq t \wedge a = \text{e-LRW}(u)] \cdot \Pr[l \leq t] \\ \geq \ & \frac{1}{|\Sigma|} \cdot \left(1 - \left(1 - \frac{1}{t}\right)^t\right) \\ \geq \ & \frac{1}{2|\Sigma|} \end{aligned}$$

Hence,

$$\Pr[(u \rightarrow v) \text{ is faulty } \mid \text{ exactly } k \ (u \rightarrow v) \text{ steps }] \geq \frac{1}{4|\Sigma|^2}$$

We can thus conclude that

$$\begin{aligned} \text{Exp}[N_{u \rightarrow v}] & \geq \sum_{k \geq 1} k \cdot \frac{1}{4|\Sigma|^2} \cdot \Pr[\text{ exactly } k \ (u \rightarrow v) \text{ steps }] \\ & = \frac{1}{4|\Sigma|^2} \sum_{k \geq 1} k \Pr[\text{ exactly } k \ (u \rightarrow v) \text{ steps }] \\ & = \frac{1}{4|\Sigma|^2} \cdot \text{Exp}[(u \rightarrow v) \text{ steps }] \\ & = \frac{1}{4|\Sigma|^2} \cdot \frac{t}{2|E|} \end{aligned}$$

where the last steps follows from linearity of expectation, the fact that the walk has t steps on expectation and a step of the walk is likely to be a $(u \rightarrow v)$ step with probability $1/2|E|$. $\square$

Having bounded $\mathrm{Exp}[N]$, we can now bound $E[N_*]$ by calculating the probability of the event "$l > B$" if $B = O(t \ln |\Sigma|)$.

Proof of Claim 6,7,8.15.

$$\begin{aligned}
\mathrm{Exp}[N_*] &= \mathrm{Exp}[N \cdot \mathbb{1}_{l \le B}] \\
&= \mathrm{Exp}[N \cdot (1 - \mathbb{1}_{l>B})] \\
&= \mathrm{Exp}[N] - \mathrm{Exp}[N \cdot \mathbb{1}_{l>B}] \\
&= \mathrm{Exp}[N] - \mathrm{Exp}[N \mid l > B] \cdot \Pr[l > B] \\
&\ge \mathrm{Exp}[N] - \mathrm{Exp}[N_F \mid l > B] \cdot \Pr[l > B] \\
&= \mathrm{Exp}[N] - \frac{|F|}{|E|} \cdot \mathrm{Exp}[l \mid l > B] \cdot \Pr[l > B] \\
&\ge \frac{t}{4|\Sigma|^2} \cdot \frac{|F|}{|E|} - \frac{|F|}{|E|} \cdot (B+t)\left(1 - \frac{1}{t}\right)^B \\
&\ge \left(\frac{t}{4|\Sigma|^2} - (B+t)e^{-B/t}\right) \cdot \frac{|F|}{|E|} \\
&\ge \frac{t}{8|\Sigma|^2} \cdot \frac{|F|}{|E|}
\end{aligned}$$

where the last inequality is obtained by setting $B = O(t \ln |\Sigma|)$. □

6,7,8.6.2 $\mathrm{Exp}[\mathbf{N_*^2}] = \mathbf{O(t)|F|/|E|}$

This is the only place in the proof where we use the fact that underlying constraint graph is a (n, d, λ)-expander. In fact, the only fact about expanders we will use is Lemma 6,7,8.11 which states that the probability that the $(i+1)$-st step of random walk whose first step is in F is also in F is bounded above by $|F|/|E| + (\lambda/d)^i$.

Proof. Since $N_* \le N_F$, it suffices to bound $\mathrm{Exp}[N_F^2]$. If we set Z_i to be the random variable that indicates if the i-th step of the random walk $a = v_0 \to v_1 \to \cdots \to v_{l-1} \to v_l = b$ is in F, we have the following expression for N_F:

$$N_F = \sum_{i=1}^{\infty} Z_i.$$

$$\begin{aligned}
\mathrm{Exp}[N_F^2] &\le \mathrm{Exp}\left[\left(\sum_{i=1\infty} Z_i\right)^2\right] = \sum_{i,j=1}^{\infty} \mathrm{Exp}[Z_i Z_j] \\
&= \sum_{i,j=1}^{\infty} \Pr[Z_i = 1] \cdot \Pr[Z_j = 1 \mid Z_i = 1] \\
&\le 2\sum_{i=1}^{\infty} \Pr[Z_i = 1]\left(\sum_{j=i}^{\infty} \Pr[Z_j = 1 \mid Z_i = 1]\right)
\end{aligned}$$

The probability $\Pr[Z_j = 1 | Z_i = 1]$ is precisely the probability that a random walk has its $(j-i)$th edge in B conditioned on the fact that the first edge of the walk is in B. This quantity is 1 if $j = i$ and bounded above by $(1-1/t)^{j-i}(|F|/|E| + (\lambda/d)^{j-i-1})$. The latter quantity is the product of the probability of the walk taking $j-i+1$ more steps and the quantity given by Lemma 6,7,8.11.

$$
\begin{aligned}
E[N_F^2] \leq & 2\sum_{i=1}^{\infty} \Pr[Z_i = 1]\left(1 + \sum_{r=1}^{\infty}\left(1-\frac{1}{t}\right)^r\left(\frac{|F|}{|E|} + \left(\frac{\lambda}{d}\right)^{r-1}\right)\right) \\
\leq & 2\sum_{i=1}^{\infty} \Pr[Z_i = 1]\left[1 + \frac{|F|}{|E|}\sum_{r=1}^{\infty}\left(1-\frac{1}{t}\right)^r + \sum_{r=1}^{\infty}\left(\frac{\lambda}{d}\right)^{r-1}\right] \\
= & 2\sum_{i=1}^{\infty} \Pr[Z_i = 1]\left(1 + t\cdot\frac{|F|}{|E|} + O(1)\right) && \text{Since } \lambda < d \\
= & O(1)\cdot\sum_{i=1}^{\infty} \Pr[Z_i = 1] && \text{Since } \frac{|F|}{|E|} \leq \frac{1}{t} \\
= & O(1)\cdot \mathrm{Exp}[N_F] \\
= & O(1)\cdot t\frac{|F|}{|E|}
\end{aligned}
$$

□

This completes the proof of the two claims, thus completing phase II, the graph powering phase.

6,7,8.7 Phase III: Alphabet Reduction

By the end of phase II, we have the following PCP transformation.

Corollary 6,7,8.13 (restated) (preprocessing+graph powering)

$$PCP^{\Sigma}_{1,1-\epsilon}[r, 2] \subseteq PCP^{\Sigma'}_{1,1-\epsilon'}[r + O(t\log t), 2]$$

where $\epsilon' = \Omega(t)\cdot\min\{\delta_I\cdot\epsilon, \frac{1}{t}\}$ *and* $\Sigma' = \Sigma^{1+d+d^2+\cdots+d^t}$.

The gap has considerably improved from ϵ to $O(t)\epsilon$ (provided the initial gap $\epsilon < 1/t$). Thus, we could choose t to be sufficiently large such that the new gap is at least twice the earlier gap and complete the proof of the gap amplification lemma 6,7,8.5. However, there is one caveat: the alphabet Σ has increased exponentially in size to $\Sigma' \approx \Sigma^{d^t}$. How, do we reduce the alphabet size? Proof Composition is tailor-made for this purpose. In fact, we showed the following consequence of proof composition in the lecture on proof composition.

Lemma 6,7,8.19.

$$PCP^{\Sigma'}_{1,1-\epsilon}[r, 2] \subseteq PCP^{\Sigma}_{1,1-\epsilon'}[r + O(\log^2|\Sigma'|) + \log_2 17, 2]$$

where $\epsilon' = \delta_{III}\epsilon$ *(for some constant* δ_{III}*, and* Σ *is any alphabet such that* $\log_2|\Sigma'| \geq 17$.

Combining this with the corollary, we have that

$$\mathrm{PCP}^{\Sigma}_{1,1-\epsilon}[r,2] \subseteq \mathrm{PCP}^{\Sigma''}_{1,1-\epsilon''}[r+f(t),2]$$

where

- $\epsilon'' = \Omega(t) \cdot \delta_{III} \min\{\delta_I \cdot \epsilon, \frac{1}{t}\}$
- $f(t) = O(t \log t) + (1 + d + d^2 + \cdots + d^t)^2 \cdot O(\log^2 |\Sigma|) + \log_2 17$.
- Σ'' is any alphabet such that $\log_2 |\Sigma'| \geq 17$.

We can choose t to be a sufficiently large constant such that $\epsilon'' = \min\{2\epsilon, \alpha\}$ (for some constant α). We can also choose the initial alphabet Σ to be Σ''. We then have that

$$\mathrm{PCP}^{\Sigma}_{1,1-\epsilon}[r,2] \subseteq \mathrm{PCP}^{\Sigma}_{1,1-\epsilon''}[r+O(1),2]$$

where ϵ'' is as described above. This completes the proof of the gap amplification lemma and thus, the PCP Theorem.

References

[Din07] Irit Dinur. *The PCP theorem by gap amplification*. J. ACM, 54(3):12, 2007. (Preliminary Version in *38th STOC*, 2006). doi:10.1145/1236457.1236459.

[Rad06] Jaikumar Radhakrishnan. *Gap amplification in PCPs using lazy random walks*. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, eds., *Proc. 33rd International Colloquium of Automata, Languages and Programming (ICALP)*, volume 4051 of *Lecture Notes in Computer Science*, pages 96–107. Springer-Verlag, Venice Italy, 10–14 July 2006. doi:10.1007/11786986_10.