6.841: Advanced Complexity Theory Fall 2012

Lecture 18 — November 13, 2012

Prof. Dana Moshkovitz *Scribe: Cheng Chen*

1 Recap

In the last lecture, we embarked on the combinatorial proof of the PCP theorem by Irit Dinur[1]. A roadmap and part of the main reduction, were given last time, and we will finish the part of the powering operation this time and finish the proof in the following lecture.

A constraint graph is a graph $G = (V, E)$ with an alphabet Σ and a collection of edge constraints $\{\phi_e\}$. Let $val(G)$ be the maximum of the fraction of constraints satisfied by some assignment. The hard direction, $NP \subseteq PCP(\log n, 1)$, of the PCP theorem is equivalent to prove that distinguishing between $val(G) = 1$ and $val(G) \leq 0.999$ is NP-hard.

Following the roadmap in [1], we start from the worst case with gap $1/|E|$, which is NP-hard from the NP-completeness of 3-coloring problem, and amplify the gap by 2 for $\log |E|$ rounds. In each round, the size of G can increase by a constant factor and the alphabet doesn't change, while the perfect completeness is maintained and the new soundness is $1 - \min\{2\delta, 0.001\}$ if the old soundness was $1 - \delta$. This is done in two steps:

1. Achieve the completeness, soundness requirements while maintaining a linear blowup in the graph size. However, the alphabet size can increase dramatically.
2. Reduce the new alphabet to the original alphabet Σ.

Last time, we talked about how to transform a constraint graph into a d-regular expander graph with self-loops at each vertex with a linear blowup in the graph size and linear decrease in gap by first reducing the degree of the constraint graph by replacing each vertex by an expander graph, then expanderize it by superimposing an expander graph on it. We will finish the first step this time by showing the powering operation to get a linear increase in gap. The second step will be left for the following lecture.

2 The Powering Operation

Definition 1. *(powering) Given a constraint graph G, the powering of G with parameter t is a constraint graph G^t such that*

1. *The vertices of G^t are V. For every length t walk (may include multiple self-loops) between u and v in G, we add an edge $\{u, v\}$ in G^t.*

   Note that we have defined an one-to-one mapping from length-t walks in G to edges in G^t. The number of edges between u and v in G^t is exactly the number of length-t walks from u to

v in G. G^t is a d^t-regular graph since there are exactly d^t walks of length-t from every vertex in G.

2. *The alphabet is Σ^{d^t}. An assignment $\sigma^t(v)$ on a vertex v in G^t corresponds to assigning a value from Σ to all the radius-t neighbors (including v itself) of v in G, i.e. vertices that can be reached by some length-t walk from v.*

   Note that since the existence of self-loops, for each radius-t neighbor u of v, there is a length-t walk from v to u. Therefore there are at most d^t neighbors of v which is a constant.

3. *The constraint $\phi_{\{u,v\}}$ on some edge $\{u, v\}$ in G^t checks if $\sigma^t(u)_w = \sigma^t(v)_w$ for every common radius-t neighbor w and $\sigma^t(u), \sigma^t(v)$ satisfy the constraints of G.*

The size of G^t is linear in the size of G since they have same vertex set and the number of edges in G^t is exactly $d^t|V|$. The perfect completeness is maintained since the assignment of G^t corresponds to the satisfying assignment of G is a satisfying assignment. So we are left with the soundness analysis.

3 Soundness Analysis

The strategy of the soundness proof is to prove its contrapositive, i.e. if we start with an assignment $\sigma^t : V \to \Sigma^{d^t}$ satisfying $1-\alpha$ fraction of constraints in G^t, we can construct an assignment $\sigma : V \to \Sigma$ satisfying $1 - \delta$ fraction of constraints in G. It suffices to show that for every constant $c > 1$ that is not too big, there exists some constant t, such that for every assignment σ^t and the constructed assignment σ, we have $\alpha \geq c\delta$.[1] The construction is defined below.

Definition 2. *(plurality assignment) Given an assignment $\sigma^t : V \to \Sigma^{d^t}$ of G^t, the plurality assignment $\sigma : V \to \Sigma$ of G is defined to be for every vertex v,*

$$\sigma(v) = \arg\max_{s\in\Sigma} |\{p \mid p \text{ is a length-}(t-1)/2 \text{ walk from } v \text{ to } u \text{ and } \sigma^t(u)_v = s\}|.$$

Note that different walks may have the same endpoint, so $\sigma(v)$ is the maximum weighted sum assignment to v of v's radius-$(t-1)/2$ neighbors, while each neighbor is weighted by the number of length-$(t-1)/2$ walks from v to u. If there are several such assignments, we can pick an arbitrary one since we only care about the following fact which simply follows the averaging lemma.

Fact 3. *Let $P_{(t-1)/2}$ be the set of all length $(t-1)/2$ walks from v, we have*

$$Pr_{p\leftarrow P_{(t-1)/2}}[\sigma^t(u)_v = \sigma(v) \text{ where } u \text{ is the other endpoint of } p \text{ other than } v] \geq \frac{1}{|\Sigma|}.$$

The main tool we use is Theorem 21.12 on pages 432-433 in the textbook [3].

Theorem 4. *(Expander walks in [3]) Let G be an (n, d, λ) graph, and let $\mathcal{B} \in [n]$ satisfying $|\mathcal{B}| \leq \beta n$ for some $\beta \in (0, 1)$. Let $X_1, \ldots, X_k$ be random variables denoting a $k-1$-step random walk in G from X_1, where X_1 is chosen uniformly in $[n]$. Then*

$$Pr[\forall_{1\leq i\leq k} X_i \in \mathcal{B}] \leq ((1-\lambda)\sqrt{\beta} + \lambda)^{k-1}$$

[1] A formal statement is $\alpha \geq \Omega(\sqrt{t}) \min\{\delta, 1/t\}$, the $\sqrt{t}$ term can be improved to t with a more careful analysis by Jaikumar Radhakrishnan and Madhu Sudan [2].

This theorem is stated for vertices, but a similar theorem holds for edges. In our settings, let $\mathcal{B}$ be the set of σ-satisfying edges in G, we have $\beta = |\mathcal{B}|/|E| \in (0,1)$ and $\delta = 1-\beta$ is very small. Recall the numerical tricks that when x is small, $\sqrt{1-x} \approx 1-x/2$ and $(1-x)^c \approx 1-cx$ for any constant c. Let P_k be the set of all length-k walks in G, we have

$$Pr_{p \leftarrow P_k}[\text{all edges in } p \text{ are } \sigma\text{-satisfying}] \leq ((1-\lambda)\sqrt{1-\delta}+\lambda)^k \approx 1-(1-\lambda)k\delta/2 = 1-O(k\delta)$$

Now we begin our proof. We need to prove that $\alpha = c\delta$ for some big enough constant $c > 1$, where α is the fraction of σ^t-unsatisfying edges in G^t. Note that in the definition of the powering operation, there is an one-to-one mapping from length-t walks in G to edges in G^t. Let P_t be set of all length-t walks in G, we have

$$\begin{aligned}\alpha &= Pr_{\{w_1,w_2\} \leftarrow G^t}[\phi^t_{\{w_1,w_2\}}(\sigma^t(w_1), \sigma^t(w_2)) = 0] \\ &= Pr_{p \leftarrow P_t}[\phi^t_{\{w_1,w_2\}}(\sigma^t(w_1), \sigma^t(w_2)) = 0 \text{ where } w_1 \text{ and } w_2 \text{ are the endpoints of } p] \\ &\geq Pr_{p \leftarrow P_t}[p \text{ passes through some } \sigma\text{-unsatisfying edge } \{v,u\} \text{ in the middle } k \text{ steps of } p \\ &\qquad\qquad \text{and } \phi_{\{v,u\}}(\sigma^t(w_1)_v, \sigma^t(w_2)_u) = 0 \text{ where } w_1 \text{ and } w_2 \text{ are the endpoints of } p]\end{aligned}$$

We say that the edge $\{v_{i-1}, v_i\}$ is in the middle k steps of a random walk $(v_0, v_1, \ldots, v_t)$ if $i \in [t/2-k/2, t/2+k/2)$. The inequality above is because the condition is more restricted. Observe that the distribution $p \leftarrow P_t$ is the same as the following distribution E^*:

1. Choose a random length-k walk q in G.

2. Extend q in both sides to a random length-t walk p while q is the middle k steps of p.

Let w_1, w_2 be the endpoints of p, where w_1 is closer to v, we have

$$\begin{aligned}\alpha &\geq Pr_{E^*}[\text{there exists some } \sigma\text{-unsatisfying edge } \{v,u\} \text{ in } q \text{ and } \phi_{\{v,u\}}(\sigma^t(w_1)_v, \sigma^t(w_2)_u) = 0] \\ &\geq Pr_{E^*}[\text{there exists some } \sigma\text{-unsatisfying edge } \{v,u\} \text{ in } q \text{ while } \sigma^t(w_1)_v = \sigma(v) \text{ and } \sigma^t(w_2)_u = \sigma(u)]\end{aligned}$$

The second inequality holds because the condition is more restricted and $\phi_{\{v,u\}}(\sigma(v), \sigma(u)) = 0$. Observe that three conditions are independent since both q and p are random walks in G.

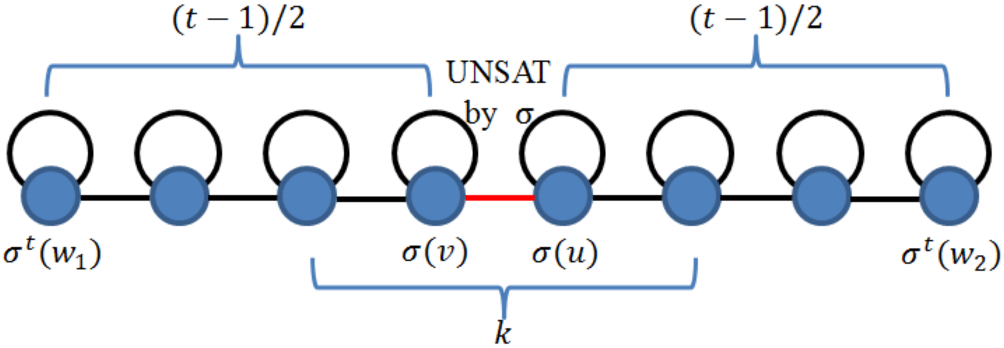

Figure 1: Illustration of distribution E^* when $\{u,v\}$ happens to be the middle of the walk

If $\{u,v\}$ happens to be the middle of the walk, then by Fact 3, we have $Pr[\sigma^t(w_1)_v = \sigma(v)] \geq 1/|\Sigma|$ and $Pr[\sigma^t(w_2)_v = \sigma(u)] \geq 1/|\Sigma|$. This situation is showed in Figure 1.

The key observation is when $k = \sqrt{t}$, $Pr[\sigma^t(w_1)_v = \sigma(v)] \geq \Omega(1/|\Sigma|)$ and $Pr[\sigma^t(w_2)_v = \sigma(u)] \geq \Omega(1/|\Sigma|)$ holds for every $\{u, v\}$ in the middle k steps. This is a property of Binomial distribution. The proof is skipped in class, but there is more details in [1] and [4].

Following this result and Theorem 4, we can lower bound α by

$$\alpha \geq \Omega\left(\frac{1}{|\Sigma|^2}\right) \cdot Pr_{E^*}[\text{there exists some } \sigma\text{-unsatisfying edge } \{v, u\} \text{ in } q] \geq \Omega\left(\frac{k\delta}{|\Sigma|^2}\right)$$

Note that $|\Sigma|$ is a constant. Choose constant t to be big enough, so $k = \sqrt{t}$ is big enough to swallow the constant hidden by Ω and $|\Sigma|$, therefore completes the proof.

References

[1] Irit Dinur, *The PCP theorem by gap amplification*, Journal of the ACM, 54(3):12, 2007.

[2] Jaikumar Radhakrishnan, Madhu Sudan, *On Dinur's proof of the PCP theorem*, Bulletin of the AMS, 44(1):19-61, 2007.

[3] Sanjeev Arora, Boaz Barak, *Computational Complexity - A Modern Approach*, Cambridge University Press 2009, isbn 978-0-521-42426-4, pp. I-XXIV, 1-579.

[4] Alan Deckelbaum, Dana Moshkovitz, *Lecture 9: The PCP Theorem Via gap Amplification*, 6.895 PCP and Hardness of Approximation, 2010.