6.895 PCP and Hardness of Approximation MIT, Fall 2010

Lecture 9: The PCP Theorem Via gap Amplification

Lecturer: *Dana Moshkovitz* Scribe: *Alan Deckelbaum and Dana Moshkovitz*

In this lecture, we discuss gap amplification, a procedure which enables us to go directly from the NP-hardness of Sat to the basic PCP theorem:

$$NP \subseteq PCP_{1,0.999}[O(\log n), 2]_{\Sigma}.$$

The proof uses the techniques we developed in the previous lecture (composition and degree reduction), but does not require the algebraic machinery we discussed before (sum check, low degree testing). The advantage of the algebraic machinery is that it in fact allows to prove a strong PCP theorem

$$NP \subseteq PCP_{1,0.001}[O(\log n), 2]_{\Sigma}.$$

while the technique we discuss in this lecture does not.

1 Graph Theoretic Formulation of PCP

We start by presenting a PCP system with two queries as a graph. The vertices correspond to locations of the proof. The edges correspond to tests made by the verifier on the two endpoints. Each vertex is assigned a symbol from an alphabet Σ by the prover. Each edge is associated with a constraint on pairs in $\Sigma \times \Sigma$. Note that the definition of "degree" we had before coincides with the definition of the maximal degree in the graph.

We consider the following gap problem: given a graph as above, distinguish between the following two cases:

- There exists an assignment to the vertices which satisfies all edges.
- For all assignments to the vertices, at most s fraction of the edges are satisfied.

Note that showing that this problem is NP-hard is equivalent to proving the PCP theorem:

$$NP \subseteq PCP_{1,s}[O(\log n), 2]_{\Sigma}.$$

We define $gap(G)$ to be $1 - s$.

2 Iterative Construction

We observe that, since the graph 3-coloring problem is NP-hard, we already have a PCP with a very small gap $1/poly(n)$ (corresponding to the NP-hardness of deciding whether a graph can be 3-colored so that there is no edge that is monochromatic, or every 3-coloring leaves at least one edge monochromatic). The overall goal of gap amplification is to enlarge the gap to 0.001 without significantly harming the other parameters.

The gap amplification operation ("powering") doubles the gap (as long as the gap is not too large, say at most 0.001) with a tolerable damage to the other parameters. After doubling the gap $\Theta(\log n)$ times, the gap will be at least 0.001.

3 Properties of Powering

The powering operation with parameter t maps G to a new graph G'. As long as G meets certain restrictions specified in the next section, we are guaranbteed that:[1]

$$gap(G') \geq \Omega(\sqrt{t}) \min\{gap(G), \frac{1}{t}\}, \tag{1}$$

where the $\Omega(\cdot)$ hides a constant that is smaller than 1.

Hence, by taking t to be large enough with respect to the constant in the $\Omega(\cdot)$, we can use the powering operation to double the gap. This is as long as the gap is not too large already (i.e., as large as $\frac{1}{t}$).

The powering operation has the following effects on the other parameters:

- The number of queries remains $q = 2$, and the completeness remains $c = 1$.
- $size(G') = C \cdot size(G)$, where C is a constant depending on t. (We call the "size" of G' the number of its vertices plus the number of its edges.) Therefore, after $\Theta(\log n)$ iterations of the powering operation, the resulting graph will have polynomial size.
- The alphabet of G' is $\Sigma^{d^{t/2}}$, where $d = deg(G)$.

4 Requirements of Powering

The powering operation works as specified above only assuming certain restrictions on G:

- **Constant Degree:** The graph is regular and its degree $d = deg(G)$ is a constant D.
- **Constant Alphabet:** The alphabet Σ is constant.
- **Expansion:** The graph has a self-loop on every vertex, and has second eigenvalue $\lambda < D$.

We remark that the constants in inequality (1) and in the bound on $size(G')$ depend on the degree of G and on the alphabet size.

While relying on the requirements above, the powering operation does not preserve them: it increases the degree and the alphabet significantly (it in fact preserve the expansion requirement). Hence, in the next sections we explain how to transform any PCP to a PCP that satisfies the requirements, using degree reduction, composition and a new trick. The transformation will have to take place before *every* application of powering. It, too, comes with a price: it slightly increases the size of the graph, and slightly decreases the gap. However, this modest cost is swallowed in the parameters of powering as described in Section 3.

4.1 Reducing Degree

The degree may be non-constant to begin with, and powering may increase it further. However, in Lecture 8 we saw how to set the degree to any desired D at the price of increasing the soundness error by $\approx 1/D$ and multiplying the size by $\Theta(D)$. Perfect completeness is preserved. Unfortunately, the transformation of Lecture 8 also multiplies the number of queries by a factor of D. Nonetheless, the number of queries can be brought back to 2 at the expense of decreasing

[1] A more careful analysis of the powering procedure, as in [RS07], can improve the $\sqrt{t}$ term to t.

the gap to $\approx poly(\frac{1}{D})$ (see first homework). By choosing as D an appropriate constant, the decrease in the gap after degree-reduction is swallowed in inequality 1, and the increase in the size is swallowed in the size bound following it.

4.2 Reducing Alphabet

The alphabet is initially constant, but the powering operation makes its size increase dramatically, as it raises it to the power of $D^{t/2}$. Applying the powering operation i times will result in an alphabet of size $|\Sigma|^{(D^{t/2})^i}$. For non-constant i (e.g., $i = \Theta(\log n)$), the alphabet size becomes non-constant (and even exponential).

We solve this problem by using the composition operations from Lecture 8 together with a Hadamard-based construction (see Lecture 8). Composition multiplies the size by a term that depends only on $|\Sigma|^{D^{t/2}}$ (a constant), and preserves perfect completeness. The gap multiplies by a constant. While the number of queries is not preserved, as before, it can be brought back to 2 at an acceptable cost.

4.3 Forming Expansion

We can guarantee that there are self loops on all the vertices, if we simply add them to the graph. We can similarly guarantee that the graph's second eigenvalue is small by adding an expander H (over the same vertex set) on top of our graph G. We make the constraints on self loops and on the edges from H be trivial (satisfied by all assignments). This procedure decreases the gap by a multiplicative factor of $\frac{D}{D+1+deg(H)}$, where D is the original degree of our graph. This is because we have the same number of unsatisfied edges, but the degree of each vertex increases from D to $D + 1 + deg(H)$, since every vertex has an additional self-loop and edges from H.

We claim that by adding self-loops and H to our graph in the second bullet above, the resulting graph has small second-eigenvalue. We prove this by using the Rayleigh quotient.

Claim 4.1 (Rayleigh Quotient). *Let A be the adjacency matrix of a regular graph. Then the second eigenvalue satisfies*

$$\lambda(A) = \max_{x \neq \vec{0},\ x \perp \vec{1}} \frac{|\langle x, Ax \rangle|}{\langle x, x \rangle} = \max_{||x||=1,\ x \perp \vec{1}} |\langle x, Ax \rangle|.$$

The Rayleigh quotient formula follows from the fact that the second eigenvalue is the largest eigenvalue when we restrict A to act on the space perpendicular to the first eigenvector. (Since A is a real symmetric matrix, it has an orthonormal eigenbasis.)

Corollary 4.2. *Let A and B be adjacency matrices of regular graphs on the same vertex set. Then*

$$\lambda(A + B) \leq \lambda(A) + \lambda(B).$$

The above corollary implies that $\lambda(A + I + H) \leq deg(A) + 1 + \lambda(H)$. Since $\lambda(H) < deg(H)$, we conclude that $\lambda(A + I + H) < deg(A + I + H)$, where $A + I + H$ is the graph we get by adding self-loops and H to our original graph.

5 The Powering Operation

We now describe the graph powering operation. We start with a graph $G = (V, E)$, an alphabet Σ and constraints associated with the edges of E. We pick some parameter t, and construct a new graph G' as follows:

- The vertices of G' are V.
- The alphabet is $\Sigma^{d^{t/2}}$, where d is the degree of G. An assignment to a vertex in G' corresponds to assigning a value from Σ to all of its radius-$t/2$ neighbors. (Note- We needed to insist that the degree of G was small so that no vertex has more than $d^{t/2}$ radius-$t/2$ neighbors.)
- We have an edge (u, v) in G' for every length-t walk between u to v in G.
- The constraint on an edge (u, v) in G' checks that all the assignments on the intersection of the radius-$t/2$ neighborhood about u and the radius-$t/2$ neighborhood about v are consistent. (That is, we check that they obey all appropriate constraints from G, and that any vertex in the overlapping neighborhood is assigned the same symbol from u and from v.)

We notice that the above construction preserves the perfect completeness property and preserves 2 queries. Moreover, the size is multiplied by $O(d^{t-1})$.

Therefore, the main analysis which remains is to study the soundness of G'. We notice that any assignment σ' to G' naturally induces an assignment σ to G by taking the "plurality vote": each vertex picks an assignment that repeats the maximal amount of times in other vertices whose labels contain assignments for it (we break ties arbitrarily):

$$\sigma(u) = \arg\max_{s \in \Sigma} |\{\frac{t}{2}\text{-step walk from } u \text{ ending at } v \text{ satisfies } \sigma'(v)_u = s\}|.$$

Notice that the following can be deduced immediately.

$$P\left(\frac{t}{2}\text{-step random walk from } u \text{ ending at } w \text{ satisfies } \sigma'(w)_u = \sigma(u)\right) \geq \frac{1}{|\Sigma|}. \tag{2}$$

Recall that while $1/|\Sigma|$ may be small, in our setup it is a constant.

Suppose that G is not satisfiable and has gap α. Further assume that $\alpha \leq \frac{1}{t}$ (otherwise, we are done). Then, by soundness of G, we know that σ has at least an α fraction of rejecting edges. The intuition is that, since G is an expander, at least $\approx 1 - (1 - \alpha)^t \approx \alpha t$ fraction of all t-step walk in G pass through at least one rejecting edge of σ. Moreover, at least $\alpha\sqrt{t}$ fraction of the t-step walks pass through a rejecting edge of σ in their "middle", i.e., in step i for $\frac{t}{2} - \sqrt{\frac{t}{2}} \leq i \leq \frac{t}{2} + \sqrt{\frac{t}{2}}$. We will show that a constant fraction of those walks correspond to rejecting edges for σ'. This will prove inequality (1).

The intuition for the above claim is that, if we let (u, v) be a rejecting edge of G in the middle of the walk, and w_1, w_2 be the endpoints of the walk, then the events $\sigma'(w_1)_u = \sigma(u)$ and $\sigma'(w_2)_v = \sigma(v)$ are: (i) independent events; and (ii) (less obviously) each happens with probability $\Omega(1/|\Sigma|)$ (a constant). The point (ii) follows from inequality 2 and the realization that due to self-loops, walks of length $\frac{t}{2} \pm \sqrt{\frac{t}{2}}$ are essentially equivalent to walks of length exactly $\frac{t}{2}$.

6 Soundness Analysis

In this section we formally prove the soundness of the graph powering operation.

Fix an assignment σ' for G', and let the corresponding plurality vote assignment σ for G be as in the previous section. By soundness of G, we know that an α fraction of the edges of G reject the assignment σ. Denote the set of rejecting edges F. We may assume that $\alpha \leq \frac{1}{t}$. We show that $\Omega(\sqrt{t})\alpha$ fraction of the edges in G' are rejecting.

6.1 Walks in G

We say that the "middle" of a t-step walk is

$$I = \{t/2 - \sqrt{t/2} \leq i \leq t/2 + \sqrt{t/2}\}.$$

We say that a walk $(v_0, v_1, \ldots, v_t)$ is *hit* by the i^{th} edge if

1. $(v_{i-1}, v_i) \in F$
2. $\sigma'(v_0)_{v_{i-1}} = \sigma(v_{i-1})$
3. $\sigma'(v_t)_{v_i} = \sigma(v_i)$.

For a t-step random walk, we let N_i be the indicator that the walk is "hit" by the i^{th} edge. We set $N = \sum_{i \in I} N_i$. We show that

$$Pr\left[N > 0\right] \geq \Omega(\alpha\sqrt{t}).$$

In order to prove the above result, we use the "second moment method". The second moment of N is $\mathbf{E}\left[(N - \mathbf{E}\left[N\right])^2\right] = \mathbf{E}\left[N^2\right] - \mathbf{E}\left[N\right]^2$. We lower bound $\mathbf{E}\left[N\right]^2$ and upper bound $\mathbf{E}\left[N^2\right]$:

Claim 6.1. $\mathbf{E}\left[N\right] \geq \Omega(\alpha\sqrt{t})$.

Claim 6.2. $\mathbf{E}\left[N^2\right] \leq O(\alpha\sqrt{t})$.

Since we have the inequality

$$P[N > 0] \geq \frac{E[N]^2}{E[N^2]},$$

the two above claims will imply that $P[N > 0] \geq \Omega(\alpha\sqrt{t})$, as desired.

6.2 Proof of Claim 6.1

We obtain a lower bound on $\mathbf{E}\left[N_i\right]$ for all $i \in I$. The result of the claim then follows by linearity. Since G is undirected and regular, we can view a choice of a random walk as follows:

- Pick (v_{i-1}, v_i) uniformly from the set of all edges in G.
- Do a random walk of length $i-1$ starting from v_{i-1}. This will give the vertices $(v_{i-2}, \ldots, v_0)$ of the walk.
- Do a random walk of length $t - i$ starting from v_i. This gives the vertices $(v_{i+1}, \ldots, v_t)$ of the walk.

Therefore, we have

$$P[N_i > 0] = \alpha \cdot P[\sigma'(v_0)_{v_{i-1}} = \sigma(v_{i-1})] \cdot P[\sigma'(v_t)_{v_i} = \sigma(v_i)].$$

We prove that for every vertex v, for any length $l \in I$, it holds that

$$P[l\text{-step random walk from } v \text{ to } w \text{ satisfies } \sigma'(w)_v = \sigma(v)] \geq \frac{\Omega(1)}{|\Sigma|}.$$

The lower bound on $P[N_i > 0]$ follows.

Notice that for $l = t/2$, this result is clear, since σ is the plurality vote over all $t/2$-step walks. Because the graph has self-loops, we see that the above probability is equal to

$$\sum_k P(\text{stay in place } (l-k) \text{ times}) \cdot P(\text{length-}k \text{ r.w. without self-loops has } \sigma^t(w)_v = \sigma(v))$$

where k in the above expression represents the number of times that the random walk takes an edge other than a self-loop, and where w and v are the start and end vertices of the length-k random walk without self-loops, respectively.

Since each vertex in G has degree d and exactly one edge at each vertex is a self-loop, we know that the number of times that an l-step random walk takes a self-loop is distributed as a $Binomial(l, 1/d)$ distribution.

We now make a technical claim about the properties of the binomial distribution, without proof:

Claim: For any c, l_1, l_2 such that

- $l_1 - \sqrt{l_1} \leq l_2 \leq l_1 + \sqrt{l_1}$
- l_1 is sufficiently large
- $|k - pl_1| \leq c\sqrt{l_1}$

there exists a constant τ (depending on c) such that

$$\tau \leq \frac{P(B(l_1, p) = k)}{P(B(l_2, p) = k)} \leq \frac{1}{\tau}.$$

The above claim follows from the properties of the binomial distribution. Intuitively, we think of the first condition as saying that $l_1 \approx l_2$, and the third condition as saying that k is approximately equal to the expectation of $B(l_1, p)$.

We now pick c such that $P(|k - \frac{pt}{2}| > c\sqrt{\frac{t}{2}}) < \frac{1}{2|\Sigma|}$ where $p = 1 - \frac{1}{d}$ (the probability of a step of the random walk not being a self-loop), k is a $Binomial(t/2, p)$ random variable, $l_1 = t/2$, and $l_2 = l$. We denote the set of all values of k such that $|k - \frac{pt}{2}| \leq c\sqrt{\frac{t}{2}}$ by K. Using the above claim, we know that

$$\forall k \in K : P(B(l, p) = k) \geq \tau \cdot P(B(t/2, p) = k).$$

We can now lower-bound the probability that a length-l random walk satisfies $\sigma^t(w)_v = \sigma(v)$ (where v is the start vertex of the walk and w is the end vertex) by

$$\sum_{k\in K} P(\text{stay put } l-k \text{ times})\cdot P(\text{a rand. walk of length k with no self-loops has } \sigma^t(w)_v=\sigma(v))$$

$$\geq \tau \sum_{k\in K} P(B(t/2,p)=k)\cdot P(\text{a rand. walk of length k with no self-loops has } \sigma^t(w)_v=\sigma(v))$$

$$\geq \tau\cdot\left(P\left(\frac{t}{2}\text{-step walk satisfies } \sigma^t(w)_v=\sigma(v)\right)-\frac{1}{2|\Sigma|}\right)\geq \tau\left(\frac{1}{|\Sigma|}-\frac{1}{2|\Sigma|}\right)=\frac{\tau}{2|\Sigma|}.$$

Thus, since we've shown that

$$P[l\text{-step random walk from } v \text{ to } w \text{ satisfies } \sigma^t(w)_v=\sigma(v)]\geq \frac{\tau}{2|\Sigma|},$$

the proof of the claim is complete.

6.3 Proof of Claim 6.2

We will define the random variable Z to be the number of steps in I that the random walk reaches an edge in F. We clearly have $N \leq Z$. Therefore, to upper bound $E[N^2]$, it suffices to upper bound $E[Z^2]$. Let Z_i be the indicator for the i^{th} step of the walk hitting F. We now compute

$$E[Z^2]=\sum_{i,j\in I}E[Z_iZ_j]=\sum_{i\in I}E[Z_i]+2\sum_{i<j;\ i,j\in I}E[Z_iZ_j].$$

Since $|I|=O(\sqrt{t})$, it follows that $\sum_{i\in I}E[Z_i]=O(\alpha\sqrt{t})$.

In the third problem set, we show that the expander property of G implies that

$$P[Z_j=1|Z_i=1]\leq \alpha+\left(\frac{\lambda}{d}\right)^{j-i+1}.$$

Therefore, we compute

$$E[Z_iZ_j]=P[Z_i=1]\cdot P[Z_j=1|Z_i=1]=\alpha\cdot P[Z_j=1|Z_i=1]\leq \alpha\cdot\left(\alpha+\left(\frac{\lambda}{d}\right)^{j-i-1}\right)$$

and hence

$$\sum_{i<j;\ i,j\in I}E[Z_iZ_j]\leq \alpha\sum_{i<j;\ i,j\in I}\left(\alpha+\left(\frac{\lambda}{d}\right)^{j-i-1}\right)<|I|^2\alpha^2+|I|\alpha\sum_{i=1}^{\sqrt{t}}\left(\frac{\lambda}{d}\right)^i=O(\alpha^2t)+O(\alpha\sqrt{t}).$$

Since $\alpha\leq\frac{1}{t}$, we conclude that $\sum_{i<j;\ i,j\in I}E[Z_iZ_j]=O(\alpha\sqrt{t})$ (where the constant depends on $\frac{\lambda}{d}$.) This concludes the proof of the claim.

References

[RS07] J. Radhakrishnan and M. Sudan. On Dinur's proof of the PCP theorem. *Bulletin of the AMS*, 44(1):19–61, 2007.