CS 880: Advanced Complexity Theory 3/27/06

Lecture 21: Satisfiability Gap Amplification

Instructor: Dieter van Melkebeek Scribe: Siddharth Barman and Anand Sinha

In the last two lectures we saw that the PCP theorem follows the existence of a polynomial time transformation which amplifies the unsatisfiability gap of constraint satisfaction problems. We modeled constraint satisfaction problems as graphs wherein vertices were to be assigned with labels and each edge had a constraint dictating permissible pair of values at the either end of the edge. Remember that the procedure worked in stages, the first one of which was to convert a given constraint graph to another such that the resulting graph was a constant degree expander. The transformation had the property that it preserved the unsatisfiability gap up to a constant factor. The next stage is the main step wherein we power the graph by taking random walks of a chosen length. This was done with an aim of increasing the unsatisfiability gap. The edges in the new graph are paths in the original graph and the constraints are the conjunction of the constraints of the edges on the path. To ensure that the constraints are well defined, we need that the value taken by a particular vertex has to contain an opinion of the values of all the vertices in a particular neighborhood of the original graph. The neighborhood was defined as a ball of radius equal to the length of the random walks under consideration. The details have already been discussed in the previous lectures. We will recapitulate the main ideas behind graph powering in this lecture and give a proof of the gap amplification lemma.

1 A brief Recapitulation of Graph Powering

In this section we formally define the powering operation.

Definition 1. *(**GAP**: Gap Amplifying Powering) : Let $G(V, E, \Sigma, C)$ be a constraint graph. We define s power of this graph as $G'(V', E', \Sigma', C')$, where $V' = V$, $E' = \{< e_1, e_2, \ldots, e_s > \mid < e_1, e_2, \ldots, e_s >$ is a valid path in G $\}$, $\Sigma' = \Sigma^{1+d(G)+d(G)^2+\ldots d(G)^s}$, $C' = \{\wedge_{i=1}^{s} c_{e_i} \mid < e_1, e_2, \ldots e_s > \in E'\}$.*

Since G is a constant degree graph (Remember that the first step of the transformation maps constraint graphs to constraint graphs which are constant degree expanders), the alphabet size is still a constant. The degree blows up to the s^{th} power. In order to apply the powering recursively, we first need to reduce the alphabet size. This can be done by the idea of proof composition. To reduce the degree we go back to the first step. Each step might lead to an increase in the size of the graph, but we will see that this blow up is only up to some fixed constant. As far as gap is concerned, the only step where we increase the gap is powering. As it will turn out later, this increase depends on s, the power to which the graph is raised. The other two steps will possibly lead to a decrease in the gap. We will choose s in such a fashion that the cumulative effect of all three steps is to increase the gap by some fixed constant factor. Logarithmic number of iterations will ensure that the gap increases up to a fixed constant. The details can be seen in the previous lecture.

2 Gap Amplification Lemma

2.1 The Statement

In this section we state the gap amplification lemma.

Lemma 1. *There is some alphabet Σ with $|\Sigma| \geq 3$ and some constant $\gamma > 0$ such that for every constraint graph G and G', the s^{th} power of G, $UNSAT(G') \geq \gamma\sqrt{s}(1-\lambda_2(G))MIN\{UNSAT(G), \frac{1}{\sqrt{s}}\}$.*

Note that we insist that the alphabet size is greater than 3. This is because we know that when this is the case, the constraint graph problem is complete for NP. Hence, we can use this problem itself to prove the PCP theorem by gap amplification. Another issue is that the constant γ might depend on $|\Sigma|$. However, it is good enough for our purposes as we will take alphabet size to be a fixed constant. Next, we give an intuition of why such a theorem should hold followed by a rigorous proof of the same.

2.2 Some Intuition

In this section, we will prove that the graph powering increases the unsatisfiability gap of constraints graphs. We will first give an intuition as to why that should be the case. Then, we will prove that our intuition is indeed correct. However, the proof will not follow intuition. It remains an open question if one can get an intuitive proof the gap amplification lemma. For the sake of simplicity, let us consider a simpler scenario. Here the $s-$powering operation just transforms the edge set into a $s-$fold Cartesian product of the original edge set. The new alphabet will be the $n-$fold Cartesian product of the original alphabet, where $n = |V|$. It is easy to see that in this case, if an assignment of labels to vertices in G fails to satisfy an α fraction of the constraint, the same assignment will fail to satisfy at least $1-(1-\alpha)^s$ fraction of constraints in G'. How about the inverse problem? I.e, if every assignment fails to satisfy at least a β fraction of constraints in the powered graph (Note that we are talking about the Cartesian product graph right now), what is maximum fraction of constraints that can be satisfied in the original graph. It can be shown without much effort that this fraction is just $1 - \beta^{1/s}$. Thus the result is tight in both the directions.

We know that and have already seen a number of times that many statistical experiments on expander graphs have similar success rates as for a pure random tests. We hope that we are lucky one more time! First note that the reason why we would like to take an expander graph is that it does not have an ugly blow up in the alphabet size. The new alphabet size is still a constant. Apart from that the the size of resulting constraint set also increases mildly (It rose to s^{th} for Cartesian product). Now let us try to see why expander graph should increase gap. Suppose, there is an assignment which fails on a subset of constraints (or equivalently edges). Let that set be B. When we take random walk of length expander, we aim to hit this set and hit it quite often. If we hit it once, the natural extension to the assignment for the original graph will fail for the constraint corresponding to this path (or edge in the powered graph corresponding to this path). We know that the probability that a random walk of length s hits a given set increases exponentially with s. However, this intuition is not enough. We just looked at the reverse problem, i.e if an assignment is bad for original graph, it is even worse for the powered graph. We need to prove that if there is a good assignment for the powered graph then there is even better one for the original graph. One way of proving this would be to show that given an assignment for powered graph, one can get an assignment for the original graph with lower unsatisfiability gap than what was assumed of it. This is exactly what we will prove in the next section.

2.3 Proof of Gap Amplification Lemma

As we already pointed out, we would like to show that if some assignment is good for the powered graph then it is good for the original one as well. Let $\sigma' : V' \rightarrow \Sigma'$ be an assignment of vertices in G'. We will come up with an assignment, $\sigma : V \rightarrow \Sigma$, for the graph G such that $UNSAT_\sigma(G) < \frac{1}{\gamma\sqrt{s}(1-\lambda_2(G))} UNSAT_{\sigma'}(G')$. There is a small catch here. It follows from the statement of the lemma that we can prove that such a gap preserving reduction exists only if $UNSAT(G) \leq \frac{1}{\sqrt{s}}$. Therefore, we assume that such in the case. Consider the assignment σ'. For every vertex in G', it dictates some value to all the vertices that are at a distance at most s from it. We would like to use this information to come up with the assignment σ. Consider an edge e in the graph G. Every vertex that is at a distance at most s from it dictates a value to the end points of this edges in the assignment σ'. We would like to have some kind of voting scheme to get the value assigned to the vertices under σ. Let $B(u, r)$ be the set of vertices that are at distance exactly r from it. We choose $\sigma(u)$ to be the plurality of the opinion of the vertices in $B(u, s/2)$ under the assignment σ'.

Now let us convert our intuition into a rigorous proof.

Proof. Let $\sigma'(x)|_u$ be the opinion of the assignment to x about u. Note that this is defined only if u is at a distance at most s from x. Let $W_{u,r}$ denote the distribution of end points of length r random walks starting from the vertex u. Let B denote the set of edges of G that are unsatisfied under the assignment σ. Since, $UNSAT(G) \leq \frac{1}{\sqrt{s}}$, we can assume that $\mu(B) \geq \frac{1}{\sqrt{s}}$. If that is not the case, one can show that $UNSAT(G')_\sigma \geq \frac{1}{\sqrt{s}}$ and we already have the promised unsatisfiability gap for the powered graph. From our definition, it follows that

$$Pr_{x \leftarrow W_{u,s/2}}[\sigma(u) = \sigma'(x)|_u] \geq \frac{1}{|\Sigma|}.$$

Also note that for any edge (u, v) in G,

$$Pr_{x \leftarrow W_{u,s/2}, y \leftarrow W_{v,s/2}}[\sigma(u) = \sigma'(x)|_u \wedge \sigma(v) = \sigma'(y)|_v] \geq \frac{1}{|\Sigma^2|}.$$

To see this note that the event that u gets the same values as the opinion of x about it is independent of the event that v gets the same value as the opinion of y about it. This is because x and y are sampled independently from their corresponding distributions. Now, suppose that some edge $(u, v) \in B$. Then any path, say (x, y), for which this is the $s/2^{th}$ edge will be violated under σ' if $\sigma(u)$ and $\sigma(v)$ are same as $\sigma(x)|_u$ and $\sigma(y)|_v$ correspondingly. We already know that the fraction of these paths is exactly $\frac{1}{|\Sigma|^2}$. Hence the fraction of paths that are violated is at least $\frac{1}{\Sigma^2}\mu(B)$. However this is not enough. Hence we need to consider more edges than just the middle edge. Remember that the every vertex in G has a self loop. This ensures that the distribution of vertices that are at distance $l \in [s/2 - \sqrt{s}, s/2 + \sqrt{s}]$ from u is too far away from the distribution $W_{u,s/2}$. In particular, it can be shown that for every $l \in [s/2 - \sqrt{s}, s/2 + \sqrt{s}]$

$$Pr_{x \leftarrow W_{u,l}}[\sigma(u) = \sigma'(x)|_u] \geq \frac{\epsilon}{\Sigma},$$

where ϵ is some fixed constant. The reason for this is that adding self loops decreases the conductance of the graph and hence it can not deviate much from the distribution, $W_{u,s/2}$, in $\sqrt{s}$ steps on either side. Therefore, we have that,

$$Pr_{x \leftarrow W_{u,l}, y \leftarrow W_{v,s-l}}[\sigma(u) = \sigma'(x)|_u \wedge \sigma(v) = \sigma'(y)|_v] \geq \frac{\epsilon^2}{\Sigma^2}.$$

Hence, we now have that a path can be violated possibly in more than one position. Let ν be the random variable denoting the number of edges in it that are in the bad set and the values at the end points under σ agrees with the opinion of the end points of the paths. To be more precise, let $\chi_i, \forall i : s/2 - \sqrt{s} \leq i \leq s/2 + \sqrt{s}$ be indicator random variables which are 1 if the i^{th} edge of a random walk is in the set B and that the end points of the i^{th} edge have the property that they get the same values as the opinion of the vertices on the either end of the path. Then ν is sum $\sum_i X_i$. We already know that $Pr[\chi_i = 1] \geq \frac{\mu(B)\epsilon^2}{\Sigma^2}$. The factor $\mu(B)$ comes because of the fact that we only consider the edges in the set B. Therefore,

$$E[\nu] \geq \frac{2\sqrt{s}\epsilon^2}{|\Sigma|^2}. \tag{1}$$

We want the probability that at least one constraint remains unsatisfied. That is equal to the probability of the event that $\nu > 0$. Let χ_ν be an indicator random variable which is one if ν is greater than zero and zero otherwise. Then we have that,

$$E[\nu] = E[\nu\chi_\nu].$$

Using Cauchy-Schwartz inequality we have that $E[AB]^2 \leq E[A^2]E[B^2]$. Therefore,

$$\begin{aligned} E[\nu]^2 &= E[\nu\chi_\nu]^2 \\ &\leq E[\nu^2]E[\chi_\nu^2] \\ &= E[\nu^2]E[\chi_\nu]. \end{aligned}$$

The last step uses the fact that for every indicator random variable χ, $E[\chi^2] = E[\chi]^2$. Since $E[\chi_\nu] = Pr[\nu > 0]$, we have that,

$$Pr[\nu > 0] \geq \frac{E[\nu]^2}{E[\nu^2]}.$$

We already have a lower bound on $E[\nu]$. We need to upper bound its second moment. For this note that,

$$\begin{aligned} \nu &= \sum_i \chi_i \\ \Rightarrow E[\nu^2] &= \sum_i \sum_{j \geq i} E[\chi_i\chi_j]. \end{aligned}$$

Let us analyze the random variable $\chi_i\chi_j$. Note that χ_i is one if the i^{th} in a length s random walk is in the set B and its end points get the same values as the opinions of the start and the end vertices of the random walk. Suppose we drop the second condition and define new random variables ${\chi'}_i$s, which are one if the i^{th} edge is in the set B. Clearly, $\chi_i = 1$ implies that ${\chi'}_i = 1$. Hence, $E[{\chi'}_i{\chi'}_j] \geq E[\chi_i\chi_j]$. Once again we have that

$$\begin{aligned} E[{\chi'}_i{\chi'}_j] &= Pr[{\chi'}_i{\chi'}_j = 1] \\ &= Pr[{\chi'}_j = 1 | {\chi'}_i = 1] Pr[{\chi'}_i = 1]. \end{aligned}$$

The first term of the above expression is just the conditional probability that the $(j-i)^{th}$ edge of a random walk is in the set B, given that the first edge was in B. We already bounded this quantity from above in the previous lecture. Using lemma3 from lecture 19, we get that,

$$Pr[{\chi'}_j = 1 | {\chi'}_i = 1] \leq \mu(B) + \lambda_2(G)^{j-i}.$$

We know that the distribution of i^{th} edges of a random walk is uniform for all values of i. Hence $Pr[\chi'_i = 1] = \mu(B)$. Thus we get that

$$
\begin{aligned}
E[\chi'_i \chi'_j] &= Pr[\chi'_j = 1 | \chi'_i = 1] Pr[\chi'_i = 1] \\
&\leq \mu(B)\left(\mu(B) + \lambda_2(G)^{j-i}\right) \\
\Rightarrow E[\nu^2] &= \sum_i \sum_{j \geq i} E[\chi_i \chi_j] \\
&\leq \sum_i \sum_{j \geq i} E[\chi'_i \chi'_j] \\
&\leq \sum_{i=s/2-\sqrt{s}}^{s/2+\sqrt{s}} \sum_{j \geq i}^{s/2+\sqrt{s}} \mu(B)\left(\mu(B) + \lambda_2(G)^{j-i}\right) \\
&\leq \sum_{i=s/2-\sqrt{s}}^{s/2+\sqrt{s}} \mu(B) \sum_{k=1}^{2\sqrt{s}} \left(\mu(B) + \lambda_2(G)^k\right).
\end{aligned}
$$

Note that we assumed that $\mu(B) \leq \sqrt{s}$. Therefore,

$$\sum_{k=1}^{2\sqrt{s}} \mu(B) = O(1).$$

At the same time, we also have that,

$$\sum_{k=1}^{2\sqrt{s}} \lambda_2(G)^k = \frac{O(1)}{1 - \lambda_2(G)}.$$

Thus we get that

$$E[\nu^2] = 2\sqrt{s}\mu(B)\left(O(1) + \frac{O(1)}{1 - \lambda_2(G)}\right) \leq 2\sqrt{s}\mu(B)\frac{\beta}{1 - \lambda_2(G)}, \tag{2}$$

where β is some fixed constant greater than zero. Now we already have that $Pr[\nu > 0] \geq \frac{E[\nu]^2}{E[\nu^2]}$. Using (1) and (2), we get that,

$$Pr[\nu > 0] \leq \gamma\sqrt{s}(1 - \lambda_2(G))\mu(B),$$

for some constant $\gamma > 0$. Note that γ can depend upon $|\Sigma|$. □

This completes the proof of the gap amplification lemma.