Approximation Algorithms and Hardness of Approximation	May 7, 2013
Lecture 19	
Lecturer: Ola Svensson	*Scribes: Bogdan Alexandru Stoica*

We continue from where we left in the previous lecture.

3 Gap amplification of nice $2CSP_w$ instances

We have seen in the last lecture that a $2CSP$ instance φ is said to be nice if and only if:

- $G(\varphi) = (V, E)$ is a d-regular graph, where $G(\varphi)$ denotes the constraint graph corresponding to φ with variables as vertices and constraints as edges (see *Definition 20, Lecture 18*)
- For every vertex $v \in V$, at least half of the incident edges are self-loops
- $G(\varphi)$ is an expander

Since one can reduce arbitrary CSP instances to nice $2CSP$ instances (*Section 2.6, Lecture 18*), it is sufficient to prove Dinur's gap amplification lemma for nice $2CSP$ instances.

Lemma 27 (Powering lemma) *Given a nice $2CSP_w$ instance ψ and a positive integer t we can produce a $2CSP_{w'}$ instance of ψ^t in time $poly(n, w^{d^{5t}})$ – where d is the degree of ψ's constraint graph – such that*

- *the alphabet size of w' is $\leq |w|^{d^{5t}}$ and the number of constraints in ψ^t is $\leq d^{2t+2} \cdot n$*
- *ψ is satisfiable $\Longrightarrow \psi^t$ is satisfiable*
- *for every $\epsilon < \frac{1}{d\sqrt{t}}$, $val(\psi) \leq 1 - \epsilon \Longrightarrow val(\psi^t) \leq 1 - \frac{\sqrt{t}}{10^5 d w^5}\epsilon$ (gap amplification)*

Proof

Let us first argue the way ψ^t is constructed:

Variables: Let ψ^t be a nice $2CSP$ instance such that it has the same number of variables as ψ. In other words, ψ^t has a "new" variable y_i for each "old" variable u_i in ψ. We shall think of the domain of the new variables to be a d^{5t} tuple over $\{0, 1, ..., w-1\}$, i.e., of size $[w]^{d^{5t}}$.

The intuition behind the new variables is that y_i claims a value of each old variable u_j if j lies in the ball centered in i and with radius $t + \sqrt{t}$ in G. The number of such vertices is $d^{t+\sqrt{t}+1} \leq d^{5t}$ since G is a d-regular graph and therefore the alphabet size w' is sufficiently large for encoding this information.

Constraints: For each $(2t+1)$-step path $p_i = \langle i_1, ..., i_{2t+2} \rangle$ there is a constraint C_p in ψ_t that depends only on y_{i_1} and y_{2t+2} (so it is a $2CSP_{w'}$) and outputs false iff $\exists j \in \{1, ..., 2t+1\}$ such that:

1. y_{i_1} claims value ω of i_j
2. $y_{i_{2t+2}}$ claims a value ω' of i_{j+1}
3. (ω, ω') does not satisfy the original constraint in ψ

Notice that constructing ψ^t takes $poly(n, w^{d^{5t}})$, $w' \leq |w|^{d^{5t}}$ and the number of paths/constraints are at most $d^{2t+2}n$ (since the constraint graph of ψ is d-regular). Also, an assignment to the variables $u_1, ..., u_n$ in ψ is naturally "expanded" to $y_1, y_2, ..., y_n$ by assigning to each d^{5t}-tuple y_i all values of u_j that lie in the ball centered in i and with radius $t + \sqrt{t}$ in $G(\varphi)$. In this way, if u_j is a satisfying assignment for ψ, then so is y_i for ψ^t since it satisfies all constraints encountered on paths of length

$\leq 2t+1$. Thus, the only property of the Powering lemma that we need to prove is the soundness property (i.e., the gap amplification).

In what follows we will argue how to transform an assignment $\mathbf{y}$ of ψ^t into an assignment $\mathbf{u}$ of ψ such that if $\mathbf{u}$ violates a few constraints, then $\mathbf{y}$ violates a fraction $\Omega(\frac{\sqrt{t}}{10^5 dw^5}\epsilon)$ of ψ^t's constraints.

The analysis has the following outline:

1. Consider an arbitrary assignment $\mathbf{y}$ and decode it to get assignment $\mathbf{u}$ to ψ
2. Since $\mathbf{u}$ leaves $m\epsilon$ constraints unsatisfied this implies that $\mathbf{y}$ leaves a ϵ' fraction of constraints unsatisfied where $\epsilon' = \frac{\sqrt{t}}{10^5 dw^5}\epsilon$

Plurality assignment. For every y_i consider a random variable Z_i which represents the outcome of the following process:

1. Start from vertex i and take a t-step random walk to reach a vertex k
2. Output the value that y_k's claims for u_i (this is intuitive since we want to determine the value of i)

Let z_i be the most likely outcome of Z_i, thus the plurality assignment of the variables $y_1, \ldots, y_n$ is $z_1, ..., z_n$. Since $val(\psi) \leq 1 - \epsilon$ there exists a subset F of constraints in ψ such that $|F| = m\epsilon = \frac{nd\epsilon}{2}$ and the plurality assignment violates these constraints.

Consider a $(2t+1)$-step path $\langle i_1, ..., i_{2t+1}\rangle \in G(\varphi)$ (or constraint $\in \psi^t$) picked uniformly at random. An edge $e = (i_j, i_{j+1}), j \in \{1, 2, ..., 2t+1\}$ is truthful if i_1 claims the plurality value of i_j and i_{2t+2} claims the plurality value of i_{j+1}. Thus, if $e \in F$ is truthful, then the constraint corresponding to this path is not satisfied. Through the next sequence of claims we show that at least a $\frac{\sqrt{t}}{10^5 dw^5}\epsilon$ fraction of the paths have such truthful edges.

Claim 28 (Random edge) *Let $G = (E, V)$ be an arbitrary graph. Then it follows that for each edge $e \in E$ and each $j \in \{1, ..., 2t+1\}$*

$$Pr[e \textit{ is the } j\textit{'th edge on path } p] = \frac{1}{|E|}$$

where the probability is taken over a random walk.

Proof We prove this using an example.

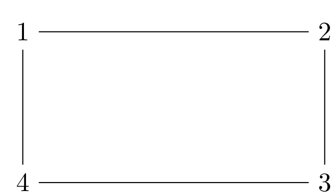

	1	**2**	**3**	**4**
1	0	$\frac{1}{2}$	0	$\frac{1}{2}$
2	$\frac{1}{2}$	0	$\frac{1}{2}$	0
3	0	$\frac{1}{2}$	0	$\frac{1}{2}$
4	$\frac{1}{2}$	0	$\frac{1}{2}$	0

Start from 1. With probability 1/2 we end up on 2 or on 4. Then $G \cdot (1\ 0\ 0\ 0)^T$ is the 1-step probability distribution.

For the general case, $G \cdot (1/n\ ...\ 1/n)^T = (1/n\ ...\ 1/n)^T$ where $n = |V|$. ■

The probability that the j-th edge in the path belongs to F is $\frac{|F|}{|E|}$. Thus, a claim is truthful if it agrees with the plurality assignment. If an edge lies roughly in the middle of a path (i.e. within $d \in \{t - \delta\sqrt{t}, ..., t + \delta\sqrt{t}\}$ steps), then it is very likely to be truthful.

Claim 29 (Middle-truthful) *Let $\delta < \frac{1}{100w}$. For each $e \in E$, each $j \in \{t, t+1, ..., t+\delta\sqrt{t}\}$. Then*

$$\gamma = Pr[\text{ } j\text{'th edge is truthful} \mid e \text{ is } j\text{'th edge }] \geq \frac{1}{2w^2}$$

Proof Let $e = (i_j, i_{j+1})$. The key idea is to launch two independent random walks – one j-step walk from i_j and one $(2t - j)$-step walk from i_{j+1}.

If $j = t$ then start a t-step random walk from i_j and a t-step random walk from i_{j+1}; this is exactly the process for defining plurality. Finally, since the two walks are independent $\gamma \geq \frac{1}{w^2}$.

However, if j varies in $\{t, t+1, ..., t+\delta\sqrt{t}\}$ we produce a ℓ-step random walk where $\ell \in \{t - \delta\sqrt{t}, t - \delta\sqrt{t} + 1, ..., t + \delta\sqrt{t}\}$ as follows

- Flip ℓ coins and let S_ℓ be the number of times it turned out heads
- Walk S_ℓ "real" edges (no self-loops)

S_ℓ follows a binomial distribution and thus similar S_t since ℓ differs from t by an additive factor $\delta\sqrt{t}$ which is small ($\delta \leq \frac{1}{100w}$). It can be shown that

$$\frac{1}{2}\sum_{\alpha} |Pr[S_\ell = \alpha] - Pr[S_t = \alpha]| \leq 10\delta$$

In other words, the distribution of a t-step walk is statistically close to both the distribution of a $(t-\delta\sqrt{t})$-step random walk and the distribution of a $(t+\delta\sqrt{t})$-step random walk. Therefore, considering the definition of a truthful edge,

$$\gamma = Pr[y_{i_1} \text{ claims the plurality value of } i_j \wedge \text{ claims the plurality value of } i_{j+1}] \geq$$

$$\geq \left(\frac{1}{w} - 10\delta\right)\left(\frac{1}{w} - 10\delta\right) \geq \frac{1}{2w^2}$$

■

Let V be the set of truthful edges in the middle of the path that belong to F. The probability that $j \in \{t, t+1, ..., t+\delta\sqrt{t}\}$ edge is in F and truthful is $\frac{|F|}{|E|} \cdot \frac{1}{2w^2}$. Using the linearity of expectation

$$E[V] = \frac{|F|}{|E|} \cdot \frac{1}{2w^2} \cdot \delta\sqrt{t} = \frac{\epsilon\delta\sqrt{t}}{2w^2}$$

It is expected to have many truthful edges, however some paths may contain only violated constraints. Note that $Pr[V > 0] \geq \frac{(E[V])^2}{E[V^2]}$. Our final goal is to prove that this probability is at least $\frac{\sqrt{t}}{10^5 dw^5}\epsilon$ (which implies that ψ^t satisfies a fraction of at most $1 - \frac{\sqrt{t}}{10^5 dw^5}\epsilon$ constraints). Claim 29 shows that $E[V]$ is high, however $Pr[V > 0]$ could potentially still be 0 with a large probability. To bound $Pr[V > 0]$ we upper bound the variance of V (actually $E[V^2]$). This is the only step where we use that $G(\varphi)$ is an expander.

Claim 30 $E[V^2] \le 25\epsilon\delta d\sqrt{t}$

Proof Let V' is the number of middle edges (i.e., edges within $t \pm \delta\sqrt{t}$ steps) that are in F. Clearly $V' \ge V$ since V counts the number of edges in F that are truthful and so $E[V'^2] \ge E[V^2]$.

Let $\mathcal{I}_j$ be the indicator function for F. Then

$$V' = \sum_{j\in\{t,\ldots,t+\delta\sqrt{t}} \mathcal{I}_j \Longrightarrow E[V'] = \sum_{i,j} E[\mathcal{I}_i \cdot \mathcal{I}_j] = \underbrace{\sum_i E[\mathcal{I}_i^2]}_{\le \epsilon\delta\sqrt{t} \text{ since } E[\mathcal{I}_i]=\epsilon} + 2\sum_{i<j} E[\mathcal{I}_i \mathcal{I}_j]$$

Also,

$$E[\mathcal{I}_i \cdot \mathcal{I}_j] = Pr[i^{th} \text{ edge is in } F \text{ and } j^{th} \text{ edge is in } F] \le$$

$$\le Pr[\text{the end point of the } i^{th} \text{ edge is in } S \text{ and the end point of the } i^{th} \text{ edge is in } S] \le$$

$$\le \frac{|S|}{|G|}\left(\frac{|S|}{|G|} + \lambda(G)^{j-1}\right),$$

where S is the set of endpoints of an edge in F. Clearly we have $|F| = \frac{\epsilon n d}{w}$ and $|S| \le \epsilon n d$. Thus,

$$E[V^2] \le \sum_{i\le j} E[\mathcal{I}_i \cdot \mathcal{I}_j] \le \sum_{i\le j}\left(\frac{|S|}{|V|}\right)^2 + \frac{|S|}{|V|}\lambda(G)^{j-1} \le \sum_{i\le j}(\epsilon d)^2 + \sum_{i\le j}\epsilon d\lambda(G)^{j-1} \le$$

$$\le \epsilon^2 d^2 \delta^2 t + \sigma\delta d\sqrt{t}\sum \lambda(G)^k \overset{*}{\le} \epsilon^2 d^2\delta^2 t + 10\delta\sqrt{t}\epsilon d \le 25\epsilon\delta d\sqrt{t} \text{ (by assumption } \epsilon < \frac{1}{d\sqrt{t}})$$

where $*$ holds since we use a good enough expander such that $\lambda(G) \le \frac{1}{10}$. ■

Finally using the fact that $Pr[V > 0] \ge \frac{(E[V])^2}{E[V^2]} \ge \frac{(E[V])^2}{E[V'^2]} = \frac{\epsilon^2\delta^2 t}{4w^4} \cdot \frac{1}{25\epsilon\delta d\sqrt{t}} = \frac{\delta\sqrt{t}}{100dw^4}\epsilon$. Hence, ψ^t satisfies at most $1 - \frac{\delta\sqrt{t}}{100dw^4} \overset{\delta=\frac{1}{100w}}{=} 1 - \frac{\sqrt{t}}{10^5 dw^5}$ and so Lemma 27 follows.

■