Lecture 6: Hardness via "dictator vs. quasirandom" tests

Feb. 1, 2007

Lecturer: Ryan O'Donnell *Scribe: Amitabh Basu*

In this lecture we show that constraint satisfaction is a hard problem. In particular, even if all the constraints are linear, it is NP-hard to distinguish between the case that there exists an assignment such that almost all of the constraints are satisfied and the case that for all assignments only about half the constraints are satisfiable.

1 Hardness of Constraint Satisfaction

Theorem 1.1 $\forall \eta > 0$ *: Given a CSP* $\mathcal{C}$ *where all constraints are of the form* $v_{i_1}v_{i_2}v_{i_3} = 1$ *or* $v_{i_1}v_{i_2}v_{i_3} = -1$*, it is NP-hard to distinguish* $val(\mathcal{C}) \geq 1 - \eta$ *and* $val(\mathcal{C}) \leq \frac{1}{2} + \eta$

A CSP of the above type is called *3-Lin* (denoting 3 Linear). The above theorem is actually optimal in the following sense. It is easy to distinguish $val(\mathcal{C}) = 1$ v/s $val(\mathcal{C}) < 1$ - since the constraints are linear equations, a solution can be found by Gaussian elimination. Moreover, we can always achieve $val(\mathcal{C}) = \frac{1}{2}$, by simply assigning -1 to all variables if the majority of the constraints have -1 in the right hand side, otherwise assigning 1 to all variables.

The proof of the theorem proceeds as follows : *3-SAT* reduces to *Gap 3-SAT* using the PCP theorem. *Gap 3-SAT* reduces to *Label cover* using a parallel repetition theorem and finally *Label cover* reduces to *3-Lin*.

In this lecture, however, we use the Unique Games Conjecture to prove hardness of *3-Lin*. Although this doesnt give a rigorous proof of hardness (since the UGC is a conjecture), the proof is easier to appreciate.

2 Unique Games Conjecture

Definition 2.1 *A two variable constraint over alphabet* $[k]$ *:* $\phi : [k] \times [k] \to \{T, F\}$ *is called unique if* $\exists$ *a permutation* σ *on* $[k]$ *such that* $\forall i \in [k], \phi(i, j) = T \Leftrightarrow j = \sigma(i)$.

In other words, $\forall i$ $\exists$ unique j so that $\phi(i, j) = T$.

Conjecture 2.2 *"Unique Games Conjecture" :* $\forall \lambda > 0, \exists k$ *s.t. given a CSP* $\mathcal{G}$ *with unique 2-variable constraints over* $[k]$*, its NP-hard to distinguish* $val(\mathcal{G}) \geq 1 - \lambda$ *and* $val(\mathcal{G}) < \lambda$

Given a CSP $\mathcal{G}$ with 2-variable unique constraints, we can associate a corresponding labeled graph with it. We have nodes for each variable and we put edges for each 2-variable constraint. The edges are labeled according to the possible pairs that satisfy that constraint, as well as the weight of that constraint. See Figure 1.

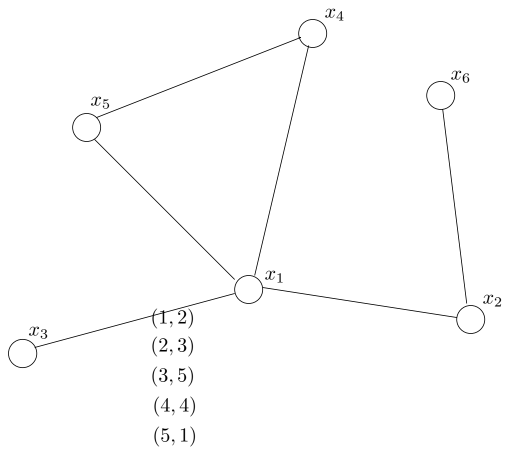

Figure 1: Constraint Graph for UGC, k = 5

Notation 2.3 *We will use $\sigma_{v\to w}$ to denote the permutation which satisfies edge (v, w).*

Remark 2.4 *We assume the graph for the UGC is regular and the weights of the constraints are the same.*

Theorem 2.5 *The conjecture is true if the uniqueness condition is dropped.*

Theorem 2.6 *It is easy to distinguish $val(\mathcal{G}) = 1$ v/s $val(\mathcal{G}) < 1$.*

Proof: Assume a label for x_1. Then deduce the labels of the remaining vertices in breadth-first order (because of the uniqueness condition). If there is a conflict at some vertex, then choose another label for x_1. Iterate through all labels for x_1, until all the edges can be satisfied. $\square$

3 Hardness of CSP via Unique Games Conjecture

Theorem 3.1 *Suppose $\forall n$, $\exists$ a function tester T, making $O(1)$ queries for $f : \{-1, 1\}^n \to \{-1, 1\}$ such that*

1. *all n dictators pass with prob $\geq C$*
2. *if $h : \{-1, 1\}^n \to [-1, 1]$ is (ϵ, δ)-quasirandom, it passes T with prob $< S$*

Then U.G.C $\Rightarrow \forall \eta > 0$, it is NP-hard given a CSP $\mathcal{C}$ of the same "type" as the test, to distinguish $val(\mathcal{C}) \geq C - \eta$ v/s $val(\mathcal{C}) < S + \eta$

Remark 3.2 *Such a tester always exists, e.g. the* Hast-Odd test.

Proof: The idea of the proof is to take an instance of the UGC, and from its constraint graph build a tester which works over $2^k n$ variables - for each node $v \in V$ in the graph, we have a function $f_v : \{-1,1\}^k \to \{-1,1\}$, where k is the size of the alphabet in UGC. So the tester will work on the strings which are truth tables of these n functions. This tester is equivalent to a CSP as shown in the previous lecture. We will show the following two properties of our reduction:

1. If $\exists$ labeling $L : V \to [k]$ that satisfies $\geq 1-\lambda$ fraction of the edges, then $\exists f_v$ such that tester accepts with prob $\geq C - O(\lambda)$.

2. If $\forall$ labelings $L : V \to [k]$, at most $\eta\delta^2\epsilon^3/64$ fraction of the edges are satisfied, then $\forall f_v$, the tester accepts with prob $< S + \eta$.

Given these two properties, we just need to set the parameters right to complete the proof. Given any $\eta > 0$, pick $\lambda < \eta\delta^2\epsilon^3/64$ then UGC shows that $\exists$ large enough $k = k(\lambda)$ such that its NP-hard to distinguish $val(\mathcal{G}) \geq 1-\lambda$ and $val(\mathcal{G}) < \lambda$ for $\mathcal{G}$ with alphabet k and hence a CSP of the same "type" as the tester has the property stated.
Proof of 1 : Given some $L : V \to [k]$ satisfying $\geq 1-\lambda$ fraction of the edges, let $f_v : \{-1,1\}^k \to \{-1,1\}$ be the $L(v)^{th}$ dictator function.

Definition 3.3 *Given $v \in V$, for each neighbour $w \sim v$, define $g_v^w : \{-1,1\}^k \to \{-1,1\}$ by $g_v^w = f_w \circ \sigma'_{v\to w}$, where $\sigma'_{v\to w}(x) = y$ implies $y_i = x_{\sigma_{v\to w}(i)}$.*

In other words, g_v^w is w's opinion on what dictator v should have.

Given a labeling L satisfying at least $1-\lambda$ fraction of the constraints, the tester T's actions are the following :

1. Pick $\boldsymbol{v} \in V$ uniformly at random.

2. Pick q random neighbours $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_q$ of $\boldsymbol{v}$ and apply T to the collection $\{g_v^{w_1}, \ldots, g_v^{w_q}\}$.

Note that since the graph is regular, this is equivalent to choosing q uniformly random edges from the graph. Therefore, by the union bound, L satisfies all $(\boldsymbol{v}, \boldsymbol{w}_i)$ with prob $\geq 1 - q\lambda = 1 - O(\lambda)$. So by definition, all the $g_v^{w_i}$'s are the same dictator function $\chi_L(v)$. So T gets applied consistently to one dictator function and hence passes with probability $\geq C$. Hence, it passes overall with prob $\geq (1 - O(\lambda)C = C - O(\lambda)$.
Proof of 2 : We prove the contrapositive of property 2. We show that given f_v's such that the tester passes with prob $\geq S+\eta$, $\exists$ labeling for the original graph $L : V \to [k]$ which satisfies $\geq \eta\delta^2\epsilon^3/64$ fraction of the constraints.

For each vertex $v \in V$, define a set of candidate labels

$$\mathcal{L}(v) = \{i : Inf_i^{(1-\delta)}(f_v) \geq \frac{\epsilon}{2} \quad OR \quad Inf_i^{(1-\delta)}(h_v) \geq \frac{\epsilon}{2}\}$$

where $h_v = \mathrm{avg}\{g_v^w\}, w \sim v$. Using the proposition from the previous lecture, $|\mathcal{L}(v)| \leq \frac{4}{\epsilon\delta}$.

By an averaging argument, at least $\frac{\eta}{2}$ fraction of v's are such that $\mathbf{Pr}[\text{tester accepts}|v] \geq S + \frac{\eta}{2}$. Call such v's "good".

Since the tester passes (given v) with prob $\geq S + \frac{\eta}{2} > S$, therefore h_v cannot be (ϵ, δ)-quasirandom

$$\Rightarrow \exists i \quad \text{such that} \quad Inf_i^{(1-\delta)}(h_v) \geq \epsilon$$
$$\Rightarrow i \in \mathcal{L}(v)$$

Moreover,

$$\begin{aligned} \epsilon &\leq Inf_i^{(1-\delta)}(h_v) \\ &= \sum_{S \ni i} (1-\delta)^{|S|-1} \widehat{h_v}(S)^2 \\ &= \sum_{S \ni i} (1-\delta)^{|S|-1} \widehat{\mathbf{E}_{w \sim v}[g_v^w]}(S)^2 \\ &= \sum_{S \ni i} (1-\delta)^{|S|-1} (\mathbf{E}_{w \sim v}[\widehat{g}_v^w(S)])^2 \end{aligned}$$

using the Cauchy Shwartz inequality,

$$\leq \sum_{S \ni i} (1-\delta)^{|S|-1} \mathbf{E}_{w \sim v}[\widehat{g_v^w}(S)^2]$$

$$\because g_v^w = f_w \circ \sigma_{v \to w}, \quad \therefore \widehat{g_v^w}(S) = \widehat{f_w}(\sigma_{v \to w}^{-1}(S))$$

$$\begin{aligned} \therefore \sum_{S \ni i} (1-\delta)^{|S|-1} \mathbf{E}_{w \sim v}[\widehat{g_v^w}(S)^2] &= \mathbf{E}_{w \sim v}[\sum_{S \ni i} (1-\delta)^{|S|-1} \cdot \widehat{f_w}(\sigma_{v \to w}^{-1}(S))^2] \\ &= \mathbf{E}_{w \sim v}[\sum_{T \ni \sigma_{v \to w}(i)} (1-\delta)^{|T|-1} \widehat{f_w}(T)^2] \qquad (T = \sigma_{v \to w}^{-1}(S)) \\ \therefore \epsilon &\leq \mathbf{E}_{w \sim v}[Inf_{\sigma_{v \to w}(i)}^{(1-\delta)}(f_w)] \end{aligned}$$

By another averaging argument, at least $\frac{\epsilon}{2}$ fraction of $w \sim v$ have $Inf_{\sigma_{v \to w}(i)}^{(1-\delta)}(f_w) \geq \frac{\epsilon}{2}$. Therefore, $\sigma_{v \to w}(i) \in \mathcal{L}(w)$. Call such neighbours "good".

We have shown that at least $\frac{\eta}{2} \cdot \frac{\epsilon}{2}$ fraction of the edges are "good-good". For any such "good-good" edge,

$$\exists i \in \mathcal{L}(v) \text{ s.t. } \sigma_{v \to w}(i) \in \mathcal{L}(w)$$

Also recall that $|\mathcal{L}(v)|, |\mathcal{L}(w)| \leq \frac{4}{\epsilon\delta}$. Construct the labeling $L : V \to [k]$ by choosing $L(v)$ randomly from $\mathcal{L}(v)$. For each "good-good" edge (v, w), with prob $\frac{1}{|\mathcal{L}(v)|} \cdot \frac{1}{|\mathcal{L}(w)|} \geq (\frac{\epsilon\delta}{4})^2$, we will choose the correct labels.

Then, $\mathbf{E}[\text{fraction of edges satisfying } \mathcal{G}] \geq \frac{\eta\epsilon}{4}(\frac{\epsilon\delta}{4})^2 = \eta\delta^2\epsilon^3/64$.
$\square$