CS 153 Current Topics in Theoretical Computer Science **Spring 2010**

Solution Set 2

Posted: June 3

If you have not yet turned in this problem set, you should not consult these solutions.

1. (a) The expected payoff is

$$\begin{aligned} &= \mathbf{E}_{u\in U}\mathbf{E}_{z_1,\ldots,z_k}\mathbf{E}_{v_1,\ldots,v_k\in N(u)}[P(f_v(\pi_{u,v_1}(z_1)),\ldots,f_v(\pi_{u,v_k}(z_k)))] \\ &= \mathbf{E}_{u\in U}\mathbf{E}_{z_1,\ldots,z_k}[P(\mathbf{E}_{v_1\in N(u)}[f_v(\pi_{u,v_1}(z_1))],\ldots,\mathbf{E}_{v_k\in N(u)}[f_v(\pi_{u,v}(z_k))])] \\ &= \mathbf{E}_{u\in U}\mathbf{E}_{z_1,\ldots,z_k}[P(f_u(z_1),\ldots,f_u(z_k))] \\ &= \mathbf{E}_{u\in U}[V(f_u)] \end{aligned}$$

where the second equality used the fact that P is multilinear and that the v_i are independently chosen neighbors, to move the expectation inside.

(b) In the proof associated with such an assignment, we have, for "good" $u \in U$ (i.e., one for which all incident edges are satisfied by A):

$$f_u(z) = \mathbf{E}_{v\in N(u)}[f_v(\pi_{u,v}(z)] = \mathbf{E}_{v\in N(u)}[(\pi_{u,v}(z))_{A(v)}] = \mathbf{E}_{v\in N(u)}[z_{\pi^{-1}_{u,v}(A(v))}] = z_{A(u)}$$

where the last equality used the fact that all edges incident to u are satisfied by assignment A. By the previous part, the expected payoff of the verifier is $\mathbf{E}_{u\in U}[V(f_u)]$. Since at least $(1-\delta)$ fraction of $u \in U$ are "good" (resulting in payoff at least c by the completeness of the dictatorship test), and $V(f_u)$ is never smaller than -1 (by the boundedness of the payoff function P), we get that $\mathbf{E}_{u\in U}[V(f_u)] \geq (1-\delta)c - \delta$, as required.

(c) By the first part, the expected payoff returned by the verifier is $\mathbf{E}_{u\in U}[V(f_u)]$. If this is larger than $s+\eta$, then by an averaging argument, it must be that for at least an η fraction of $u \in U$ (which we will call "good"), $V(f_u) > s$. Then by the soundness of the dictatorship test, there exists an i for which $I_i(f_u) > \tau$, and thus $L(u)$ is non-empty for "good" $u \in U$.

Consider some u and i for which $I_i(f_u) \geq \tau$. If we define $g_v(z) = f_v(\pi_{u,v}(z))$, we get (by replacing f_u with its definition)

$$I_i(\mathbf{E}_{v\in N(u)}[g_v]) \geq \tau.$$

By convexity of I_i, we find that

$$\mathbf{E}_{v\in N(u)}[I_i(g_v)] \geq \tau.$$

Since $\sum_i I_i(g_v) \leq t$ we have in particular that for every i, $I_i(g_v) \leq t$. Thus, by an averaging argument, at least a $\tau/(t+1)$ fraction of $v \in N(u)$ satisfy $I_i(g_v) \geq \tau/(t+1)$.

Call the associated edges (u, v) "good". Since $I_i(g_v) = I_{\pi_{u,v}(i)}(f_v)$, we see that for good edges (u, v), the label $\pi_{u,v}(i)$ is a member of $L(v)$.

Thus when choosing a random labeling from the candidate label sets (choose an arbitrary label if a set is empty), we see that for "good" edges (u, v), the probability they are satisfied is $|L(u)|^{-1}|L(v)|^{-1}$ (because as we've just argued, $i \in L(u)$ and $\pi_{u,v}(i) \in L(v)$). How large can the candidate label sets be? By the bounded sum property of I_i, we have that $|L(u)| \le t/\tau$ and $|L(v)| \le t/(\tau/(t+1)) = t(t+1)/\tau$. We conclude that the fraction of edges satisfied, in expectation, is at least

$$\eta \cdot \frac{\tau}{t+1} \cdot \frac{\tau}{t} \cdot \frac{\tau}{t(t+1)},$$

and thus some assignment achieves at least this fraction of edges satisfied.

(d) Our instance of GCSP has one variable for each coordinate in the intended proof described above (so there are $2^R \cdot |V|$ variables). Each possible k-tuple S queried by the verifier corresponds to a k-tuple of variables. We set the associated weight w_S to be the weight of S under distribution μ. Then is is easy to see that every assignment to the variables of this GCSP corresponds to a possible proof supplied to the verifier. Part (b) shows that in the YES case, there is an assignment with payoff at least $(1 - \delta)c - \delta$; part (c) shows that in the NO case, there is no assignment with payoff more than $s + \eta$, provided

$$\delta < \eta \cdot \frac{\tau}{t+1} \cdot \frac{\tau}{t} \cdot \frac{\tau}{t(t+1)}.$$

Thus by taking $\eta > 0$ and $\delta > 0$ sufficiently small, we can make the ratio of soundness to completeness approach s/c. Then any approximation algorithm that does better than this ratio can be used in the standard way to distinguish between the YES and NO cases of the original Label Cover instance.

2. (a) A linear function f satisfies $f(x)f(y) = f(xy)$ for all x, y, and so in this case the acceptance probability is 1. For the soundness direction, as suggested, we begin by observing that $\Pr[\text{accept}] = \mathbf{E}_{x,y}[1/2 + (1/2)f(x)f(y)f(xy)]$. We have $\mathbf{E}_{x,y}[f(x)f(y)f(xy)]$

$$\begin{aligned} &= \mathbf{E}_{x,y}\left[\left(\sum_S \hat{f}(S)\chi_S(x)\right) \cdot \left(\sum_T \hat{f}(T)\chi_S(y)\right) \cdot \left(\sum_U \hat{f}(U)\chi_U(x)\chi_U(y)\right)\right] \\ &= \sum_{S,T,U} \hat{f}(S)\hat{f}(T)\hat{f}(U)\left[\mathbf{E}_x[\chi_S(x)\chi_U(x)] \cdot \mathbf{E}_y[\chi_T(y)\chi_U(y)]\right] \end{aligned}$$

The first expectation above is 1 if $S = U$ and 0 otherwise; the second expectation above is 1 if $T = U$ and 0 otherwise. So the above expression simplifies to

$$\sum_S \left[\hat{f}(S)\right]^3 \le \epsilon \cdot \sum_S \left[\hat{f}(S)\right]^2,$$

where we have used the assumption that $\max_S \hat{f}(S) \le \epsilon$. Finally, by Parseval, the above expression is at most ϵ. Plugging back into the original expectation, we conclude that the acceptance probability is at most $1/2 + \epsilon/2$.

(b) Let us proceed as in the previous part. We begin by observing that Pr[accept] = $\mathbf{E}_{x,y}[1/2 + (1/2)g(x)h(y)f(xy)]$. We have $\mathbf{E}_{x,y}[g(x)h(y)f(xy)]$

$$\begin{aligned} &= \mathbf{E}_{x,y}\left[\left(\sum_S \hat{g}(S)\chi_S(x)\right)\cdot\left(\sum_T \hat{h}(T)\chi_S(y)\right)\cdot\left(\sum_U \hat{f}(U)\chi_U(x)\chi_U(y)\right)\right] \\ &= \sum_{S,T,U} \hat{g}(S)\hat{h}(T)\hat{f}(U)\left[\mathbf{E}_x[\chi_S(x)\chi_U(x)]\cdot\mathbf{E}[\chi_T(y)\chi_U(y)]\right] \end{aligned}$$

The first expectation above is 1 if $S = U$ and 0 otherwise; the second expectation above is 1 if $T = U$ and 0 otherwise. So the above expression simplifies to

$$\sum_S \hat{g}(S)\hat{h}(S)\hat{f}(S) \le \epsilon \cdot \sum_S \hat{g}(S)\hat{h}(S),$$

where we have used the assumption that $\max_S \hat{f}(S) \le \epsilon$. By Cauchy-Schwarz,

$$\sum_S \hat{g}(S)\hat{h}(S) \le \left[\left(\sum_S \hat{g}(S)^2\right)\cdot\left(\sum_S \hat{h}(S)^2\right)\right]^{1/2}$$

Finally, by Parseval, the last expression above equals 1 (using the fact that g and h are Boolean). Plugging back into the original expectation, we conclude that the acceptance probability is at most $1/2 + \epsilon/2$.

(c) As suggested, we analyze the terms

$$\mathbf{E}_{x_1,x_2,\ldots,x_k}\left[\prod_{(i,j)\in P} f(x_i)f(x_j)f(x_ix_j)\right]$$

one by one, for non-empty sets P. Choose some i^*, j^* such that $(i^*, j^*) \in P$. Fix each x_ℓ for $\ell \notin \{i^*, j^*\}$ to the value $c_\ell \in \{-1,+1\}^n$ that *maximizes* the expectation (one can imagine fixing the variables in this fashion one by one). After this fixing, the product in the expectation,

$$\prod_{(i,j)\in P} f(x_i)f(x_j)f(x_ix_j),$$

simplifies considerably. For example, for pairs $\{i,j\} \in P$ that don't intersect $\{i^*, j^*\}$, $f(x_i)f(x_j)f(x_ix_j) = f(c_i)f(c_j)f(c_ic_j)$, which is a constant in $\{-1,+1\}$. Let C be the product (over pairs $\{i,j\} \in P$ that don't intersect $\{i^*, j^*\}$) of these constants. The remaining pairs $\{i,j\}$ have intersection with $\{i^*, j^*\}$ of size 1 or 2. Consider the pairs with intersection size 1. The contribution of $f(x_i)f(x_j)f(x_ix_j)$ from such pairs depends only on x_{i^*}, or only on x_{j^*}, so we can capture it in the definition of g and h as follows (this is not our final definition):

$$\begin{aligned} g'(x) &= \prod_{(i,j)\in P: i=i^*, j\neq j^*} f(x)f(c_j)f(xc_j) \\ h'(x) &= \prod_{(i,j)\in P: i\neq i^*, j=j^*} f(c_i)f(x)f(c_ix). \end{aligned}$$

For the single pair $(i^*, j^*) \in P$, we need g to capture the $f(x_{i^*})$ term and h to capture the $f(x_{j^*})$ term, and we need to capture the constant C. So our final definition of g, h is

$$\begin{aligned} g(x) &= f(x) \cdot g'(x) \\ h(x) &= C \cdot f(x) \cdot h'(x). \end{aligned}$$

We now have that, after the fixing, the product in the expectation equals $g(x_{i^*})h(x_{j^*})f(x_{i^*}x_{j^*})$. Thus, by our choice of the c_ℓ, we have

$$\mathbf{E}_{x_1,x_2,\ldots,x_k}\left[\prod_{(i,j)\in P} f(x_i)f(x_j)f(x_ix_j)\right] \leq \mathbf{E}_{x_{i^*},x_{j^*}}[g(x_{i^*})h(x_{j^*})f(x_{i^*}x_{j^*})].$$

Applying the analysis in part (b), we know that each such term is bounded above by ϵ. When $P = \emptyset$, the term equals 1 (this is not an arbitrary convention but rather the correct answer when one considers how the original expectation was expanded). Altogether we get:

$$\sum_{P\subseteq\binom{[k]}{2}} \mathbf{E}_{x_1,x_2,\ldots,x_k}\left[\prod_{(i,j)\in P} f(x_i)f(x_j)f(x_ix_j)\right] \leq 1 + (2^{\binom{k}{2}} - 1)\cdot\epsilon < 1 + 2^{\binom{k}{2}}\cdot\epsilon$$

from which we get the desired bound by plugging back into the original expression.

(d) The linearity test described in the previous part makes $q = k + \binom{k}{2}$ queries, and it achieves soundness error $2^{-\binom{k}{2}} + \epsilon = s + \epsilon$. Thus the amortized query complexity $q/\log_2(1/s)$ approaches 1 as k approaches infinity.