Theory Lunch, October 19, 2005 Scribe: Sam Riesenfeld

Approximating Unique Games

Kunal Talwar

The subject of this talk is work done by the speaker at Microsoft Research in Redmond in collaboration with Anupam Gupta.

Feige and Lovasz introduced unique games in 1992. In the Unique Games problem, we are given a graph $G = (V, E)$, a set of labels $[k]$, and a set of constraints represented by a permutation $\pi_{uv} : [k] \to [k]$ for each edge $(u, v) \in E$. Given a labeling $f : V \to [k]$, we call edge $e = (u, v) \in E$ *satisfied* if $\pi_{uv}(f(u)) = f(v)$. The goal is to define a labeling $f : V \to [k]$ for G that satisfies as many edge constraints as possible.

It is straightforward to check if a game is satisfiable by using propagation. The more interesting question is whether the following conjecture is true. Khot initiated much of the interest in this conjecture by showing that many hardness results stem from it. It basically states that it is NP-Hard to distinguish when many or only few edges are satisfiable. Let $m = |E|$.

Proposition 1 (Unique Games Conjecture) *For every $\epsilon > 0$, there is a k such that given a Unique Game on the label set $[k]$, it is NP-Hard to distinguish between the case when at least $(1 - \epsilon)m$ edges are satisfiable and the case when at most ϵm edges are satisfiable.*

If the optimal labeling violates at most ϵm edge constraints, Khot shows how to find a labeling that violates at most $\epsilon k^2 m$ constraints. Trevisan improves this result to $(\epsilon \log m)^{\frac{1}{2}} m$.

On the other side, Feige and Reichman show $\Omega(2^{\log^{1-\delta} n})$-hardness for the MAX-SAT version of the Unique Games problem if OPT violates at least $(1 - \epsilon)m$,

Also Charikar and the Makarychev twins recently showed that if OPT violates ϵm of the constraints, then there is an algorithm that satisfies $\frac{1}{k^{\epsilon}}$ constraints.

The speaker and Gupta give an $O(\log n)$-approximation (specifically, a $(\epsilon \log n)m$-approximation) for MIN-UNSAT, which is better for small ϵ than the result of Trevisan. In contrast to previous results, their solution is LP-based rather than SDP-based.

For each node $u \in V$ and label $l \in [k]$, let $x(u, l)$ be a variable such that $x(u, l) = 1$ if $f(u) = 1$ and 0 otherwise. Here is an Integer Program for the Unique Games problem:

$$\min \sum_{(u,v)} w_{uv} \sum_{l \in [k]} d(u, v, l),$$

where w_{uv} is the weight of edge $e = (u, v)$, such that

$$\sum_{l \in [k]} x(u, l) = 1$$

$$d(u, v, l) \geq |x(u, l) - x(v, \pi_{uv}(l))| \qquad \forall\, (u, v) \in E,\ l \in [k].$$

The first constraint guarantees that every node gets a label. The second constraint ensures that $d(u, v, l) = 1$ exactly when the edge (u, v) constraint is violated and is 0 otherwise.

The problem with the LP relaxation of this IP is that at cost 0, a labeling assignment can get away without giving any opinion at all. For example, if the label set is $\{0,1\}$ and G is the triangle on three points, with the inequality constraint on each edge, then every integral labeling violates at least one constraint. However, the fractional labeling that assigns $x(u,l) = \frac{1}{2}$ and $x(u,0) = \frac{1}{2}$ to each node $u \in V$ violates no constraints.

Notice that if G is a tree, then the Unique Game on it is always satisfiable. So consider a cycle c, with nodes $u = v_0, v_1, \ldots, v_t = u$, and an assignment to a node in the cycle, say u. Let l_0^S be the label assigned to u. Let $l_1^S, \ldots, l_t^S$ be the labels assigned to the other nodes in the cycle by propagating. If it is the case that l_t^S may or may not be l_0^S, then we say that l_0^S is *Bad* for (u,c) in this case.

In order to eliminate these kinds of Bad assignments, we add the following constraints to the LP for each cycle c:

$$\sum_{i=1}^{t} d(v_{i-1}, v_i, l_{i-1}^S) \geq x(u, l_0^S) \qquad \forall u, l_0 \in \text{Bad}(u,c)$$

There are exponentially many such constraints, of course, but elementary graph theory can be used to get a separation oracle. (No details on this here.)

Let us define a distance d by

$$d(u,v) = \sum_l d(u,v,l).$$

Distance d gives a sense of the extent to which the LP violates edge (u,v). Let $\hat{d}$ be the shortest path metric of G with distances $d(u,v)$ on edges.

The results of Fakcharoenphol, Rao, and Talwar show that we can approximate $\hat{d}$ by a tree metric T such that $T(u,v) \geq \hat{d}(u,v)$ for every edge $(u,v) \in E$ and

$$\sum_{(u,v)\in E} w_{uv} T(u,v) \leq (\log n) \sum_{(u,v)\in E} w_{uv} \hat{d}(u,v).$$

We will show that $\mathbf{P}[(u,v) \text{ violated}] \leq T(u,v) + d(u,v)$ (times a constant), and then the expected cost analysis gives the $\log n$ approximation.

Let us consider just one edge (u,v), with label set $\{0,1\}$ so we can assume that π_{uv} is the identity. Suppose $l(u,0) = \frac{1}{4}$ and $l(u,1) = \frac{3}{4}$, and $l(v,0) = \frac{1}{3}$ and $l(v,1) = \frac{2}{3}$. The violation probability is

$$\frac{1}{4} \cdot \frac{2}{3} + \frac{3}{4} \cdot \frac{1}{3} = \frac{5}{12},$$

but the LP only pays $\frac{2}{12}$. We would like to force the LP to make just one choice. We could accomplish this somewhat by correlating the distributions so that they have the same marginals but the "transportation distance" between them is minimized. For example, in the previous example, the transportation distance is the minimum cost of transforming the distribution $(\frac{1}{4}, \frac{3}{4})$ at u into $(\frac{1}{3}, \frac{2}{3})$ at v.

We define the variable $c_{l,l'}(u,v)$, which is 0 if $\pi_{uv}(l) = l'$ and 1 otherwise. The LP for computing the transportation distance between $x(u,\cdot)$ and $x(v,\cdot)$ is:

$$\min \sum_{l,l' \in [k]} c_{l,l'}(u,v) y(u,v,l,l')$$

such that

$$\sum_{l' \in [k]} y(u,v,l,l') = x(u,l) \qquad \forall\, l \in [k]$$
$$\sum_{l \in [k]} y(u,v,l,l') = x(v,l') \qquad \forall\, l' \in [k]$$
$$y(u,v,l,l') \geq 0.$$

Observe that

$$\sum_{l,l'} c_{l,l'}(u,v) y(u,v,l,l') = \frac{1}{2} \sum_{l} |x(u,l) - x(v,\pi_{uv}(l))| \,.$$

(This may be off by a factor of 2.) To find a labeling, we first choose a label for u by rounding according to $x(u,l)$. If u gets label l_u, then we label v according to the distribution $y(u,v,l,l')$, i.e.

$$\mathbf{P}[\, v \text{ gets labeled } l' \,] = y(u,v,l,l').$$

Because of the constraints in the transportation distance LP, we get the correct marginals for $x(u,l)$ and $x(v,l')$. Now let us compute

$$\begin{aligned}
\mathbf{P}[(u,v) \text{ is violated}] &= \sum_{l \in [k]} \mathbf{P}[\, u \text{ gets } l \,] \mathbf{P}[\, \text{violation} \mid u \text{ gets } l \,] \\
&= \sum_{l} x(u,l) \sum_{l'} c_{l,l'}(u,v) \frac{y(u,v,l,l')}{x(u,l)} \\
&= \frac{1}{2} \sum_{l} |x(u,l) - x(v,\pi_{uv}(l))| \\
&\leq d(u,v).
\end{aligned}$$

So we can start from any node and propagate the assignments in the tree T. This analysis shows that for every $e \in T$, the LP pays exactly for the violation of that edge. Now what about for edges *not* in T? Consider the cycle c induced by adding (u,v) to T. There are two cases. In the first case, u gets a "good label" in which case if (u,v) gets violated that implies that some other edge in the cycle is violated, which happens with probability at most $T(u,v)$.

In the second case, u gets a Bad label for c. In that case

$$\begin{aligned}
\mathbf{P}[\,(u,v) \text{ is violated}\,] &\le \sum_{l \in \mathrm{Bad}(u,c)} x(u,l) \\
&\le \sum_{l^S \in \mathrm{Bad}(u,c)} \sum_{i=1}^{t} d(v_{i-1}, v_i, l^S_{i-1}) \\
&= \sum_{i=1}^{t} \sum_{l^S \in [k]} d(v_{i-1}, v_i, l^S_{i-1}) \\
&= \sum_{i=1}^{t} \sum_{l^S_{i-1} \in [k]} d(v_{i-1}, v_i, l^S_{i-1}) \\
&= \sum_{i=1}^{t} d(v_{i-1}, v_i) \\
&= T(u,v) + d(u,v)
\end{aligned}$$

Note that we have left out some things: there may be edges in T that are not in the original graph G. It turns out that it works out in the end.

Open Questions: This LP is not comparable with the SDP of Khot and Trevisan because it contains cycle constraints. What happens if you add cycle constraints to the SDP?