CS 6840: Algorithmic Game Theory **Spring 2017**

Lecture 27: March 31

Lecturer: Éva Tardos *Scribe: Sujit Rao*

We decided last time that we would go back and look at the solutions of the last two problems on PS2. The solutions on CMS are alright, but they're not particularly good at advising you on how you could have thought of this.

27.1 Problem 5

Problem 5 involved a Hotelling game. We had a graph $G = (V, E)$ and a total of N people sitting at different vertices. Players locate themselves at vertices, possibly with overlaps. Each person chooses a player who is closest to themselves. If one vertex has the same distance to multiple players, then the people at that vertex divide themselves evenly among the closest players. The utility of player i is $u_i(s) = \#$ people selecting i.

The standard version of this game takes place on the line segment $[0, 1]$. There is a continuum of people, one person at each point $x \in [0, 1]$. Each votes in a two-person election to the closest candidate, and the two candidates want to locate themselves along the line so as to get the maximum number of votes. The Nash equilibrium of this version with two players is for both candidates to place themselves at $\frac{1}{2}$, since a candidate can always get more votes by moving closer to the other. With 3 candidates, there is no pure Nash equilibrium. (This is not necessarily obvious.)

In the graph version, we think of the players as coffee shops trying to locate themselves in a city.

The homework problem would have been maybe better if it had three parts instead of the two it did have:

(a) In a pure strategy Nash equilibrium with k players, we have $u_i \geq \frac{N}{2k}$ for all i.

(a′) In a mixed Nash equilibrium, each player's expected utility is at least $\frac{N}{2k}$. (Part (a) for a mixed Nash equilibrium.)

(b) Part (a) for a coarse correlated equilibrium.

27.1.1 Part (a)

For part (a), we have $\sum_i u_i(s) = N$ for all s. Thus on average each player has utility $\frac{N}{k}$, so there must be some player i such that $u_i(s) \geq \frac{N}{k}$. Any player could just copy player i and get utility at least $\frac{N}{2k}$, so to have a Nash equilibrium each player must already be getting utility at least $\frac{N}{2k}$. (Starbucks often does this – for example, in Collegetown they set up a shop right across the street from Stella's.) This argument can be adapted for (a′), but doesn't quite get you to (b).

27.1.2 Part (a′)

Let σ be the mixed Nash distribution. We have $\sum_i E_{s\sim\sigma}[u_i(s)] = N$, so there is some i such that $E_{s\sim\sigma}[u_i(s)] \geq \frac{N}{k}$. But we were made to pick a pure strategy, not a mixed strategy, so we can't just directly copy player i. So let $p_{is} = P(i \text{ chooses } s)$, which gives

$$E[u_i(s)] = \sum_{s_i} p_{is_i} E_{s_{-i}\sim\sigma}[u_i(s_i, s_{-i})].$$

This is a convex combination of the $E_{s_{-i}\sim\sigma}[u_i(s_i, s_{-i})]$, and it is a NAsh equilibrium, so i is only playing her/her best options, so it must be the case that for s_i played with positive probability $E_{s_{-i}\sim\sigma}[u_i(s_i, s_{-i})] = E_{s\sim\sigma}[u_i(s)] \geq \frac{N}{k}$. What happens if we select of one these s_i? If the other player also chooses s_i, then the two of them together get $E_{s_{-i}\sim\sigma}[u_i(s_i, s_{-i})] \geq N/k$, and so we get at least half of this. If the person i play somewhere else, the two of them together get at least N/k, as playing two of these strategies each guaranteeing at least N/k can only help the total. However, not clear which gets more.

The idea to solve this last problem of what the expectation is if a player j switches to a strategy s_i or player i, while i is player a different one of their strategies is to select the strategy s_i using the same distribution as i uses. With this new strategy, j and i play the same startegy distribution, so should get the same reward. By the argument above the two of their reward is at least N/k, so hence this deviation must get reward at least $N/2k$ for j in expectation. Selecting the best performing of these strategies then gives a deterministic strategy s_i that also has expected utility at least $N/2k$.

27.1.2.1 Correlated distributions

Let's see why this doesn't work for a correlated distribution. If σ is correlated, it's still true that $\sum_i E_{s_i\sim\sigma}[u_i(s)] = N$, so there is still some i such that $E_{s\sim\sigma}[u_i(s)] \geq \frac{N}{k}$. But we don't know that copying i will be useful (although we also don't have a proof that it's definitely not useful).

To see why the previous argument breaks down, consider a graph with c connected components and people evenly distributed across all of the vertices. In any reasonable solution, there will be a coffee stand in each component. Consider the following correlated distribution: player 1 chooses a uniformly random connected component, and the other players (at least $c-1$) fear player 1 so much that they evenly distribute themselves among the other $c-1$ connected components. This isn't an equilibrium. There's only one player who's doing very well, and that's player 1 – they have expected utility $\frac{N}{c}$, where c is the number of components.

We claim that blindly copying player 1 isn't good, or at least, it is not resulting in utility 1/2 of what player 1 used to get. With probability $\frac{1}{c}$, you go to the component player 1 is in, and with probability $\frac{c-1}{c}$ you land in another connected component and get some small fraction of what remains. The total expected utility is

$$\approx \frac{1}{c}\frac{N}{2c} + \left(1 - \frac{1}{c}\right)\frac{N(c-1)}{ck}.$$

The issue is that player 1 benefits from some correlation, and you're trying to copy that benefit. But selecting a strategy unilaterally destroys the correlation, and thus the advantage player 1 has, and thus the advantage you're trying to get.

27.1.3 Part (b)

Let's try to choose a random i instead of a single i.

Claim 27.1 *Let σ be a coarse correlated equilibrium. Let $s' \sim \sigma$ and i be a uniformly random player, giving a random strategy s'_i. Suppose that j uses strategy s'_i. Then we claim that $E_{s',s\sim\sigma,i}[u(s'_i, s_{-j})] \geq \frac{N}{2k}$. Note that s' and s are independent.*

Proof: We have

$$E_{s',s\sim\sigma}[u(s_i', s_{-j})] \geq E_{s',s\sim\sigma}[u_i(s', s)] = \frac{N}{2k}.$$

where there are $2k$ players in the second expectation. ■

In the original question, we have several iterations $s^1, \ldots, s^T$. Then there is a correlated distribution $\tilde{\sigma}$ where we first select a uniformly random $1 \leq t \leq T$ and then output s^t. The strategy is to first select t, then select a random i, and then do s_i^t. So we're interested in

$$\sum_{\ell} u(s_i^t, s_{-j}^{\ell}).$$

27.2 Problem 6

Due to limited time, we'll not talk about the parts which are pure probability. One issue with multiplicative weights is that it randomizes at each time step, which is not what people do in real life. Follow the Perturbed Leader is a more realistic algorithm.

Regular FTL is a greedy algorithm. The cost of a at time t is $c^t(a)$. At each time t, we select

$$\operatorname*{argmin}_a \sum_{l=1}^{t-1} c^l(a).$$

Note that this is deterministic. As we saw, purely deterministic algorithms can't work.

27.2.1 Part (d)

We flip some coins independently, and have the random variables $X_a = (\#$ number of flips till H with $P(H) = \epsilon)$. We should at least know the expected value of X_a, which is

$$E[X_a] = \sum_{k=1}^{\infty} P(X_a \geq k) = \sum_{k=1}^{\infty} (1-\epsilon)^{k-1} = \frac{1}{1-(1-\epsilon)} = \frac{1}{\epsilon}.$$

Part (d) asks to show that $E[\max_a X_a] = O(\frac{1}{\epsilon}\log(n))$. It uses a similar trick, but is harder to use.

27.2.2 Part (b)

We want to look at the cost

$$\sum_{l=1}^{t-1} c^l(a) - X_a.$$

Consider the imaginary algorithm which chooses

$$\operatorname*{argmin}_a \sum_{l=1}^{t} c^l(a) - X_a.$$

The main idea is that if we could do this, the algorithm would work (with or without the X_a). The next part asks us to show that with probability at least $1-\epsilon$, including or not including $c^t(a)$ in the sum makes no difference.

If we could select $\min_a \sum_{l=1}^t c^l(a)$ at time t, giving the action a^t, then we have

$$\sum_{l=1}^{t} c^l(a^t) \leq \sum_{l=1}^{t} c^l(a^{t+1}).$$

Summing this inequality from $t=1$ to T gives the desired inequality, since the sum telescopes.

Doing this with the random term X_a makes these values telescope also, so we get by the choice of a^t

$$\sum_{l=1}^{t} c^l(a^t) - X_{a^t} \leq \sum_{l=1}^{t} c^l(a^{t+1}) - X_{a^{t+1}}.$$

Summing over all $t=1,\ldots,T$, and using an arbitrary action a as a^{T+1} (just to unify notation), we get that all terms $c^l(a^t)$ for $l<t\leq T$ occur on both sides, and hence cancel, and all terms X_{a^t} also occur on both sides and hence cancel, except X_{a^1} and $X_{a^{T+1}}$. So we get

$$\sum_{t=1}^{t} c^t(a^t) - X_{a^1} \leq \sum_{l=1}^{T} c^l(a^{T+1}) - X_{a^{T+1}}.$$

Rearranging terms, and recalling that a^{T+1} was an arbitrary action we got

$$\sum_{t=1}^{t} c^t(a^t) \leq \min_a \sum_{l=1}^{T} c^l(a) - X_a + X_{a^1}.$$