March 14 - Price of Anarchy for GSP

Instructor:Eva Tardos Jonathan DiLorenzo (jd753)

Administrative details:

The fourth problem set will be out after break.
That will be the last one before the final.

Generalized Second Price (GSP)

Def:
n is the number of slots.
$\forall i$. we have some α_i corresponding to the click-through rate at slot i.
m is the number of ads (or advertisers).
v_j is the value per click for ad j.
γ_j is the quality factor for ad j.
The probability of someone clicking ad j in slot i is $\alpha_i \times \gamma_j$.

For today, we will assume that $\forall j.\ \gamma_j = 1$. This is a common assumption, which mostly serves to simplify notation.

In GSP, we ask advertisers for some bid b_j and sort by $b_j \times \gamma_j$ (i.e. sort b_j given our assumption). Note that bids are given at a per-click rate, not a total.

We can safely assume $b_1 \geq ... \geq b_m$ (on account of the sorting noted above), which means that $\forall i.\ p_i = b_{i+1}$ because it's a second price auction (well, excluding that last i, where we say $p_i = 0$).

The slots have a total ordering based on α, so also assume WLOG that $\alpha_1 \geq ... \geq \alpha_n$.

We can set $m = n$ by adding phantom slots (if $n < m$, where they have an $\alpha = 0$), or by adding phantom bidders (if $m < n$, where they have a $b = 0$), so for simplicity, we'll consider this situation.

On a side note, apparently Google invented the γ and Yahoo did not initially use it. The γ helps the search company get more money.

Price of Anarchy

Theorem 1. *Price of Stability of GSP is* 1 *(i.e. in a full info game there exists a Nash Eq. that is optimal).*

We will come back to this theorem a bit after break as it needs a topic that hasn't yet been covered.

So now to the part we're actually dealing with:

Theorem 2. *Price of Anarchy: all Bayesian Nash of GSP have* $SW(NE) \geq \frac{1}{4} SW(Opt)$ *assuming* $\forall i.\ b_i \leq v_i$ *where* i *are advertisers.*

In fact, it turns out that $SW(NE) \geq \frac{1}{2(1-\frac{1}{e})} SW(Opt)$, but we won't prove this today. Also, note that the second condition, $\forall i.\ b_i \leq v_i$ tends to be accurate since you don't want to bid more than your value as bidding above your value is dominated by bidding the value itself (see further explanation at the end).

It's best to think of the value/click as not being random. You can sort of figure this out if you're an advertiser. In actuality, the real randomness comes from γ which turns out to be super random, but in our case we're assuming it to be 1.

We prove Theorem 2:

Proof. Recall that $u_i = (v_i - p_i) * \alpha_{k_i}$ where k_i is the slot that i gets with bid b_i.

Firstly, we choose some $b_i^* = \frac{v_i}{2}$ because this happens to be convenient for our proof.

If b is the Bayesian Nash vector and b^* is the bid vector from above, then:

$E_{v_{-i}}(u_i(b_i^*, b_{-i})|v_i) \leq E_{v_{-i}}(u_i(b_i))|v_i)$

by the definition of a Nash. We take the expectation over v_i and sum over i:

$\sum_i E_v(u_i(b_i^*, b_{-i})) \leq \sum_i E_v(u_i(b_i))$

And so we get our standard Bayesian Nash.

Now, suppose in Opt, ad i goes to slot j_i. In that case, i contributes $v_i \times \alpha_{j_i}$ to $SW(OPT)$. Note that this is the value times the number of clicks, since $\gamma = 1$.

Let β_j be the bid that actually wins slot j in GSP. Note that this is a random variable. Also, recall that $b_i^* = \frac{v_i}{2}$. Then: $u_i(b_i^*, b_{-i}) \geq \frac{1}{2} v_i \alpha_{j_i} - \beta_{j_i} \alpha_{j_i}$.

Here is the intuition for why this is true: We want to claim that i 'wins' if he gets slot j_i (the slot he gets in optimum) or better. He loses if he gets a slot lower than j_i. If he wins, then the price $p_i \le b_i^* = \frac{v_i}{2}$. Then, $v_i - p_i \ge \frac{v_i}{2}$ and the number of clicks is greater than or equal to α_{j_i} (since the slots are ordered by α and he did at least as well as slot j_i). Thus, the above inequality holds (since u_i must be greater than the first term on the right side of the inequality).

If he loses, then it's still true because $\frac{v_i}{2} \le \beta_{j_i}$ so it just says that $u_i \ge 0$ (or some negative number).

Now, we sum over all players (explanations of some steps below the equations):

$$\begin{aligned}
\sum_i u_i(b_i^*, b_{-i}) &\ge \frac{1}{2}\sum_i v_i\alpha_{j_i} - \sum_i \alpha_{j_i}\beta_{j_i} && (1)\\
&= \frac{1}{2}OPT(v) - \sum_i \alpha_{j_i}\beta_{j_i} && (2)\\
&= \frac{1}{2}OPT(v) - \sum_i \alpha_{k_i}b_i && (3)\\
&\ge \frac{1}{2}OPT(v) - \sum_i \alpha_{k_i}v_i && (4)\\
&= \frac{1}{2}OPT(v) - SW(b(v)) && (5)
\end{aligned}$$

Of note:
k_i in steps (3) and (4) are meant to denote the slot that player i gets with bid b_i.
The equation in (3) is true because if we sum over i, we cover all the values whether we use the j_i notation or not.
The inequality in (4) is true because $\forall i. v_i \ge b_i$.

Thus, we now know both of these things:
$\sum_i E_v(u_i(b_i^*, b_{-i})) \ge \frac{1}{2}E_v(OPT(v)) - E_v(SW(b(v)))$
$\sum_i E_v(u_i(b_i^*, b_{-i})) \le \sum_i E_v(u_i(b_i)) \le E_v(SW(b(v)))$

and so we get:

$2E_v(SW(b(v))) \ge \frac{1}{2}E_v(OPT(v))$

And so we've proven what we want.

Final claim: Bidding above your value is a dominated strategy. bid $b_i > v_i$ is dominated by $b_i = v_i$. If you're bidding above your value either you pay more and you're hosed or you pay less than your value and then you may as well have bid the same as your value.

Thus, the assumption made that $b_i \le v_i$ is a decent assumption. Of course, in the real world we may want to drive neighbors out of business or make sure that they don't get business at least, in which case bidding above our value is perhaps worth it. Though arguably you could include that in your value.

$\square$