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

Lecture 22: March 20

Lecturer: Éva Tardos *Scribe: Michael Luo*

Today we discuss basic generalized second-price (GSP) auction mechanisms and price of anarchy bounds. These are the mechanisms actually used in the real-world ad auctions, like the ones Google and Yahoo host.

22.1 Generalized second-price auctions

Intuitively, we can think of GSP auctions in terms of a search engine's ad auctions. For a given search query, there are many slots for ads on the results page. These slots have some intrinsic click probability, and the advertisers bid to be placed in the better slots. More formally:

Definition 22.1 (Generalized second-price auction) *There are ad slots $1, 2, ...$ sorted in decreasing order by α_i, the probability that a click will occur at slot i. Each advertiser j bids b_j, and assume that the advertisers are sorted in decreasing order by their bids. Each advertiser is assigned the ad slot associated with her rank-order. Each advertiser only pays when she receives a click, and she only pays b_{i+1}.*

By "sorted" we mean that the ad slots are enumerated such that $\alpha_1 \geq \alpha_2 \geq ...$ and the advertisers are enumerated such that $b_1 \geq b_2 \geq ....$ Notice that this implies advertiser 3 is assigned to ad slot 3 – and that she only has to pay b_4 if her ad is clicked.

We assume each player j has a value v_j for receiving a click. We then write her utility as the difference between her value and her cost for each click, multiplied by the probability of receiving a click.

$$u_j = (v_j - b_{j+1})\alpha_j$$

22.2 (Extended) generalized second-price auctions

In the above formulation for GSP auctions, we assumed that the probability of receiving a click is solely determined by where an ad is placed. More realistically, we can also consider the inherent clickability of a given ad – the likeliness that the ad attracts a click.

In this model, we introduce a relevance parameter γ_j for each advertiser j. The probability of advertiser j receiving a click in position i is $\alpha_i\gamma_j$. The advertisers are now sorted by $b_j\gamma_j$ instead of just b_j (because the auctioneer would like to fill its best ad slots with highly-expensive-and-clickable ads, not simply highly-expensive ads).

In order to calculate how much each advertiser should pay for a click, we consider the "critical value" – the minimum bid it takes for an advertiser to maintain her position. Recall that advertisers are ordered such that $b_1\gamma_1 \geq b_2\gamma_2 \geq ....$ Thus for each advertiser j, this critical value is the minimum b_j such that $b_j\gamma_j \geq b_{j+1}\gamma_{j+1}$. In other words, for each click, advertiser j must pay

$$p_j = b_{j+1}\frac{\gamma_{j+1}}{\gamma_j}$$

Using this cost variable, we can now easily write the utility as the difference between an advertiser's value for a click and her cost, multiplied by the probability of receiving a click.

$$u_j = (v_j - p_j)\alpha_j\gamma_j$$

22.3 Other possible extensions

- Implementing a reserve price to participate in the auction.
- Considering that there is value in just being seen, and charging for that.
- Accounting for externalities, where being surrounded by more attractive ads can hurt your chances of being clicked.
- Modeling value as a function in the number of clicks, acknowledging that advertisers may have diminishing returns or necessary critical masses.

Unfortunately, our analysis doesn't easily carry over to these extensions.

22.4 Relating to the real world

- Auctioneers don't simply know γ. Google spends time, energy, and (therefore) money to learn these relevance parameters, which is a whole other problem.
- Why does Google charge per click, rather than per view or per conversion? This is for historical reasons, and because charging per click is seen as a middle ground. The click is observable to the bidder (as the viewer landed on their page) and indicates some real interest; at the same time, clicks from the Google search are also observable by Google.
- For a while Google used the extended GSP above and Yahoo used regular GSP above, and you could see clear differences between the ad results. For example, eBay was super well-known (thus not very clickable for things not sold on eBay), so they bid crazy high amounts on Yahoo to land top spots on Yahoo essentially for free. Google could learn this low clickability and thus lower eBay's assigned position accordingly.

22.5 Price of anarchy bound

We show a derivation for a bound on the price of anarchy for the regular GSP mechanism described above. If you are willing to carry around a bunch of γ's, you could show that the same proof strategy applies for the extended GSP mechanism. Additionally, in our derivation the bid function depends only on one individual's value (rather than everyone else's), so the same proof also applies to bound the price of anarchy in Bayesian equilibria.

Before we begin, we first make three important notes:

1. If there is only one ad slot, GSP simplifies to single-item second-price, where we know the POA is 1.
2. In this proof we assume $b_i \leq v_i$ for all i, which isn't true for some pathological Nash equilibria. We think this is a reasonable assumption because bidding $b_i > v_i$ is dominated by bidding $b_i = v_i$.

3. Truthful bidding is clearly not always a dominant strategy in GSP. Consider three bidders with values 100, 99, and 2 and two ad slots with click probabilities 1.00 and 0.99. The top bidder doesn't actually want to get the top slot, because the second slot is practically just as good (for far less cost).

We also make two less important notes:

1. GSP auctions are second-price-esque because players don't actually pay their own bid; they pay a lesser amount. Specifically, each player pays the minimum – the "critical value" – it takes to stay in her current assigned position.
2. GSP auctions are first-price-esque because players are incentivized to shade their true values, rather than to reveal their true values.

The rest of the lecture is devoted exclusively to proving this bound:

Claim 22.2 *GSP has a POA bound of 4.*

Proof:

Consider a strategy vector b from a Nash equilibrium in a GSP auction. We would like to bound the price of anarchy by comparing this equilibrium to OPT, but we don't actually know OPT. Inspired by our success in previous lectures, we (arbitrarily) choose the alternative bidding strategy $b_i^* = v_i/2$ where each advertiser considers bidding half of her true value.

As a notational convenience, we say that bidder i in b would have gotten slot j_i in OPT.

Much like we have accomplished other proofs of price of anarchy bounds, we start by considering a single player's utility. Because there are multiple items to win, we cannot follow the case-by-case if-she-wins-and-if-she-doesn't analysis that worked in single-item second-price auctions. However, we can still prove that

$$u_i(b_i^*, b_{-i}) \geq \left(\frac{v_i}{2} - b_{j_i}(b)\right)\alpha_{j_i}$$

is true, where $b_j(b)$ is denoting the bid that won the slot j with bid vector b, so we compare $v_i/2$ to the bid that won i's rightful slot in OPT. The inequality is true because if i gets slot j_i or better with bid b_i^* then $u_i(b_i^*, b_{-i}) \geq \frac{v_i}{2}\alpha_{j_i}$, and otherwise there are j_i or more higher bidders.

Summing this inequality over all players, we get that

$$\sum_i u_i(b_i^*, b_{-i}) \geq \frac{1}{2}\sum_i v_i\alpha_{j_i} - \sum_i b_{j_i}(b)\alpha_{j_i}$$

1. The first summation on the RHS is clearly half of the optimal social welfare, as each value is multiplied by it's "rightful slot" in OPT.
2. The last summation can be re-indexed to be $\sum_i b_i(b)\alpha_i$ which (using our assumption that $b_i \leq v_i$) is less than or equal to $\sum_i v_i(b)\alpha_i = SW(b)$.

Adding that players do not regret not bidding b_i^*, and possibly taking expectations over bids (and maybe also values) we get

$$\mathbb{E}\Big[SW(b,v)\Big] = \mathbb{E}\Big[\sum_i u_i(b,v)\Big] \geq \mathbb{E}\Big[\sum_i u_i(b_i^*(v_i), b_{-i})\Big] \geq \frac{1}{2}\mathbb{E}\Big[OPT(v)\Big] - \mathbb{E}\Big[SW(b,v)\Big]$$

Altogether, we can write this as

$$\mathbb{E}\Big[SW(b,v)\Big] \geq \frac{1}{4}\mathbb{E}\Big[OPT(v)\Big]$$

proving the claim. ■