CS 6840: Algorithmic Game Theory | Spring 2017

Lecture 21: March 17

Lecturer: Éva Tardos | *Scribe: Angela Zhou*

21.1

Last time (Monday) we pointed out that the second price auction, even in a single-item setting, has an issue because it has many bad equilibria. People with low value bid very high; other people bid very low, and this equilibrium is very inefficient.

We like to assume that bidding above your value, $b_i > v_i$, is unrealistic. In a single item auction, this would be a bad idea because it's always better to just bid $b_i = v_i$. More formally, $b_i > v_i$ is dominated by bidding $b_i = v_i$, that is $u_i(b_i, b_{-i}) \geq u_i(v_i, b_{-i})$ for any $b_i > v_i$ and any b_{-i}.

Consider the setup from last time where we had multiple items, simultaneously bidding on all items with unit demand bidders. $v_i(S) = \max_{j \in S} v_{ij}$. To start with, we'll assume that for whatever reason, it's physically impossible to bid on multiple things.

What's the right way to bid? It remains valid that bidding above your value is a strictly dominated strategy. If you choose to bid on item j (which is a strategic choice - everyone needs to choose which item to bid on), what is the optimal amount to bid? Recall that we are running second-price on each item: give item j to the max bidder on j on the second price.

Is it still optimal to bid your valuation? Yes; by the time you make the decision you're participating in a normal second price auction. So the strategic part is deciding which item to bid on.

We'll analyze the Price of Anarchy for this game of second price simultaneous auction; full information (or Bayes-Nash). The way the proof goes is a smoothness-style argument.

We'll start with the full-information case. Everyone has their v_{ij} valuation and σ is the Nash equilibrium of the full-information game. Consider the optimum allocation vector: $x_{ij} = 1$ if i gets j in OPT and 0 otherwise. In OPT, social welfare is $\sum_{i,j} x_{ij} v_{ij}$. Given the definition, this picks out the assignments made in OPT and gives the total welfare.

We want to write down the usual inequality. We want $\mathbb{E}_{b \sim \sigma}[u_i(b)] \geq \mathbb{E}_{b \sim \sigma}[u_i(b_i^*, b_{-i})]$. Note that $b_i^* =$ bid v_{ij} on item j assigned in OPT. This is the usual inequality corresponding to Nash equilibrium conditions. For all valuations v and all possible bid vectors b, we need a smoothness-style inequality, which we'll show next:

Lemma 21.1 *Smooth-ness type inequality:*

$$\sum_i u_i(b_i^*(v), b_{-i}) \geq OPT - \sum_j \max_k b_{kj}$$

Proof: Consider a player i winning item j in OPT, so $x_{ij} = 1$. Now condition on the event of a win or loss when bidding $(b_i^*(v), b_{-i})$. Suppose he wins item j, then his utility is $v_{ij} - p_j(b) \geq 0$, where $p_j(b)$ is a price that is independent of b_i^*. More specifically, $p_j(b) = \max_{k \neq i} b_{kj}$, in the second-price setting, he pays the highest of the other bids on this item.

What if he loses? This inequality is actually true if he wins or loses! Since he then has 0 utility, someone must have outbid him, so something nonnegative is greater than something negative.

So

$$u_i(b_i^*(v), b_{-i}) \geq v_{ij} - p_j(b)$$

A similar version is also true for all other items, and even the sum of items

$$u_i(b_i^*(v), b_{-i}) \geq \sum_j (v_{ij} - p_j(b)))x_{ij}$$

as this simply adds a number of 0 terms to the right hand side.

Note also we can eliminate the maximum over $k \neq i$ since we are only making the term smaller by maximizing over more items.

Then we do our favorite thing, sum over all users:

$$\sum_i u_i(b_i^*(v), b_{-i}) \geq \sum_{ij} [v_{ij} - \max_k b_{kj}]x_{ij}$$

$$\sum_i u_i(b_i^*(v), b_{-i}) \geq OPT - [\sum_j \max_k b_{kj}] \sum_i x_{ij}$$

$$\geq OPT - [\sum_j \max_k b_{kj}]$$

since $\sum_i x_{ij} \leq 1$. ■

Theorem 21.2 *Price of Anarchy of a full-information mixed Nash:*

$$PoA \leq 2$$

e.g. you are guaranteed at least half the social welfare at a Nash equilibrium.

Proof:

$$\mathbb{E}_{b \sim \sigma}[\sum_i u_i(b)] \geq \mathbb{E}_{b \sim \sigma}[\sum_i u_i(b_i^*(v), b_{-i})]$$

$$\geq OPT(v) - \mathbb{E}_{b \sim \sigma}[\sum_j \max_k b_{kj}]$$

This is the standard approach we have been taking but this time, we have this maximum bid term $\sum_j \max_k b_{kj}$ which is not in fact auction revenue! Note that we can't relate this to revenue since this is the maximum bid; but because we assume no one is bidding above their value, we can relate this to welfare.

Fact:

$$\sum_j \max_k b_{kj} \leq SW(v)$$

where $SW(v) =$ sum of values of bidders for what they get.

This follows from the fact that $b_{ij} \leq v_{ij}$ and the fact that people bid on a single item only. Now rearranging the inequality, we get

$$2\mathbb{E}[SW(v,b)] \geq \mathbb{E}[\sum_i u_i(b) + \mathbb{E}[\sum_j \max_k b_{kj}]] \geq OPT$$

$\blacksquare$

Remark: This factor is coming from something different from last time; last time the factor of 2 came from bid shading. Here we cannot make direct claims on revenue of the auction.

21.1.1 General smoothness variant for second-price

Assume special bids b, we have an inequality for (b_i^*, b_{-i}):

$$\sum_i u_i(b_i^*(v), b_{-i}) \geq OPT(v) - \sum_i b(S_i(b))$$

where $S_i(b)$ is the subset of items upon which they happen to bid and win with bidding strategy b. We look at the bids for what they win, and sum them up.

21.1.2 Generalization of bidding strategy

How general is this proof? Let's think about removing this rule where you bid on a single item. Keep the definition of b_i^* as is; e.g. bid only on optimal assignment. You can think about it for the first-price and second-price setting; recall that for the first-price setting, bid $\frac{v_{ij}}{2}$ for OPT item and for second price bid v_{ij} on item assigned by OPT.

The smoothness inequality from first price isn't affected at all: recall we had

$$\sum u_i(b_i^*(v), b_{-i}) \geq \sum_j (\frac{v_{ij}}{2} - \max_k b_{kj})x_{ij})$$

and for second price:

$$\sum u_i(b_i^*(v), b_{-i}) \geq \sum_j (v_{ij} - \max_k b_{kj})x_{ij}$$

In the first-price setting, $\max_k b_{kj}$ is revenue for item j of the auction, and this remains true no matter how many items each bidder is bidding on.

So we still get $SW = \sum_i u_i + \text{Rev}$, which then is at least $\geq \frac{OPT}{2}$ in expectation, if the bids are at Nash equilibrium.

However, in the case of the second price auction we have $SW + \sum_j \max_k b_{kj} \geq OPT$.

Do we still have a nice argument for bounding $\max_k b_{kj} x_{ij}$ in the case of multiple items? Note that in the case of unit demand, you don't want to bid your value on everything. It may be a good idea to bid on multiple items; just don't bid so much (e.g. shade value appropriately), as winning too many things is dangerous, as winners have to pay.

Much of the literature is making the following useful (but somewhat unjustified) assumption:

Assume that for all players i and all S subsets of items $\sum_i b_{ij} \leq v_i(S) = \max_{j \in S} v_{ij}$, this is, bidders are conservative, as they do not risk paying more than their values ever. Note that this strategy is not dominating: this is not guaranteed to be a good idea depending on your probability of winning. This assumption implicitly implies that everyone is afraid of negative utility.

Under this assumption of conservative bidding, the price of anarchy of symmultanous second price auction is at most 2, as the assumption implied that $\sum_i \max_k b_{kj} \leq SW(b)$ even when bidders can bid on multiple items.