CS 6840 Algorithmic Game Theory 09/27/2024

Lecture 14: Auctions with multiple items

Instructor: Eva Tardos *Scribe: Mona Mohammadi*

In this lecture, we explore the price of anarchy in multi-item auctions, a natural extension of single-item auctions like first-price or all-pay auctions. The basic structure of such a problem is as follows: we have n items, each sold separately via a first-price auction. Each player i has a value v_{ij} for item j, and we assume each player desires at least one of the items and would like at most one item.

For a set of items A, the player's valuation is as follows, $v_i(A) = \max_{i \in A} v_{ij}$. There is free disposal, meaning that if a player receives an item they did not want, they can discard it without any cost. Although more sophisticated multi-item auction settings exist, we will focus on this basic valuation framework. One even simpler assumption that many of multi-item auctions have is that the valuation function of a set of items is additive.

Imagine there are n houses and m buyers. The socially optimal solution would be to match each person to at most one house and each house to one person. Formally, we are interested in finding a matching that maximizes the sum of values across all assignments:

$$\max_{M \in \text{Matching}} \sum_{(i,j) \in M} v_{ij}$$

In Homework 1, we studied how to compute such an optimal matching. However, we are not focused on algorithmic computation today, but rather on how to measure efficiency loss, specifically through the lens of the **Price of Anarchy** (PoA).

Given an arbitrary strategy profile S, we offer each player i a strategy S_i^* such that they do not regret deviating to it. If player i is matched to item j in the optimal matching, we define their strategy as:

$$S_i^* := (0, \ldots, 0, \frac{v_{ij}}{2}, 0, \ldots, 0)$$

This strategy is motivated by two key observations: - We want to leverage the optimal matching structure. - Bidding half of their valuation $\frac{v_{ij}}{2}$ tends to perform well in first-price auctions.

Now, what can we say about $u_i(S_i^*, S_{-i})$?

Claim 1. *For any player i, and their utility under the strategy S_i^*, denoted $u_i(S_i^*, S_{-i})$, satisfies:*

$$u_i(S_i^*, S_{-i}) \geq \frac{v_{ij}}{2} - p_j(S)$$

where $p_j(S)$ is the price of item j in strategy profile S.

Proof. As before there are two cases.

- If player i wins item j he would pay his bid on item j in S_i^*, which is $v_{ij}/2$, and then they get their value minus the price, which is $v_{ij} - v_{ij}/2 = v_{ij}/2$, the right hand side is smaller than this by $p_j(S)$, so it satisfy the inequality.
- If if player i does not win item j it means $p_j(S) \geq v_{ij}/2$ and therefore $u_i(S_i^*, S_{-i})$ is greater than a negative number.

■

If i has no item in the optimal matching then

$$S_i^* := (0, 0, \dots, 0)$$

and clearly,

$$u_i(S_i^*, S_{-i}) \geq 0$$

Now, we do the same trick as always; we sum this inequality over all the players. Thus,

$$\sum_i u_i(S_i^*, S_{-i}) \geq \frac{1}{2} \sum_{(i,j \in M^*)} v_{ij} - \sum_{j \text{matched in} M^*} p_j(S)$$
$$\geq \frac{1}{2}\text{OPTSW} - Rev(S)$$

where M^* is the maximum value matching, OPTSW is the optimal social welfare, and Rev(S) is the revenue under strategy profile S. Now, consider a mixed Nash equilibrium σ. We have

$$\mathop{\mathbb{E}}_{S\sim\sigma}(u_i(S)) \geq \mathop{\mathbb{E}}_{S\sim\sigma}(u_i(S_i^*, S_{-i}))$$

This holds more generally for correlated equilibria (CCE) as well. If we sum over all the players

$$\sum_i \mathop{\mathbb{E}}_{S\sim\sigma}(u_i(S)) \geq \sum_i \mathop{\mathbb{E}}_{S\sim\sigma}(u_i(S_i^*, S_{-i}))$$
$$\geq \mathop{\mathbb{E}}_{S\sim\sigma}(\frac{1}{2}\text{OPTSW} - Rev(S))$$

Hence we derived

$$\mathop{\mathbb{E}}_{S\sim\sigma}(SW(S)) \geq \frac{1}{2}OPTSW$$

Can players learn how to bid?

To enable players to learn optimal bidding strategies, we need to discretize the bid space, similar to the previous lectures. Specifically, for each player i, we discretize their bids into the following set:

$$\{0, \varepsilon v_{ij}, 2\varepsilon v_{ij}, \dots, (1-\varepsilon)v_{ij}\}$$

where ε represents a small discretization parameter. This discretization introduces a tradeoff — we give up a small ε-error but simplify the strategy space significantly. For simplicity, we assume $1/\varepsilon$ is an even integer, so that $v_{ij}/2$ is included on the list as well.

The total number of possible strategies for each player depends on the number of items, m, and the level of discretization. If each player can choose from $1/\varepsilon$ bid values for each item, then the total number of strategies is $(1/\varepsilon)^m$.

This presents a problem: the strategy space grows exponentially with the number of items m, making it computationally infeasible for players to explore all possible strategies. To mitigate this, we restrict each player to bid on only one item, reducing their strategy space to $m \cdot (1/\varepsilon)$, where m is the number of items, and $1/\varepsilon$ is the number of discretized bid options per item.

Remember for this learning setting with T rounds of bidding, the regret bound for a player is given by:

$$\sqrt{2T \ln k}$$

where k is the number of available strategies. In this case, $k = m \cdot 1/\varepsilon$. Although the regret bound is not overly large even with $(1/\varepsilon)^m$ strategies (namely it would be $\sqrt{2Tm \ln 1/\varepsilon}$). The main challenge is that running a regret-minimizing algorithm is computationally expensive due to the large number of potential strategies.

Challenges in practical application

Although first-price auctions typically achieve a similar price of anarchy even with incomplete information about the other players (a more realistic scenario), we cannot use the same argument here for the multi-item case. The no-regret strategy we proposed earlier use the optimal matching to decide what to not regret. This is a different match in each iteration, rather than a fixed strategy throughout the time, which is what the no-regret algorithms guarantee. So there is a buck for Bayesian setting that what I should not regret depends on the values of the other bidders and we do not know how to do that because we do not know who else is around.