

Lecture 18: Strategic Uncertainty: Combinatorial Auctions

Lecturer: *Sahil Singla*

Last updated: March 12, 2022

In this lecture, we'll take a game-theoretic look at algorithmic problems and consider algorithmic problems under "strategic" uncertainty. Specifically, we'll want to consider optimization problems where the input itself has a stake in the output selected. Consider the following problem, referred to as *combinatorial auctions*. There are n bidders and m items. Each bidder has a *valuation function* for the m items $v_i(\cdot)$. That is, player i receives value $v_i(S)$ for receiving the set S of items. Your optimization problem is to find an allocation, that is a partition of $[m]$ into S_1, \dots, S_n so as to maximize the *social welfare*, $\sum_i v_i(S_i)$.

This is already an interesting algorithmic problem (in HW-1 we saw a 2-approximation greedy algorithm for this problem when the bidders have submodular valuations), and we'll see more of it this lecture. But we'll also want to focus on the following: what if you don't know each player's valuation function, but you rely on them to provide you information. Then every algorithm you design imposes a *game*: depending on the behavior (information provided) of the other players and your own behavior (information you choose to provide), you get some payoff (your value for the output of the algorithm). In general, we might need to use payments to address this problem. This adds a new angle to the problem that we must consider: each player is *strategic* and wants to maximize their own utility, which is the difference of their payoff and their payment.

Example 1 (Single-Item Auction). *Consider the case where $m = 1$. This is an easy algorithmic problem: everyone just has a value v_i for getting the item, and you need to find $\arg \max_i \{v_i\}$ and give them the item. However, if this is your algorithm, it induces a game among the bidders where the only equilibrium is to report the highest allowable value (reporting a higher value makes you more likely to get the item and doesn't cost anything). So even though you can easily find the maximum value amongst the provided input, the provided input is garbage and doesn't tell you anything about the actual maximum.*

*One alternative is to use instead the first-price auction. Now, you still give the item to the highest reported bid, but the winner must pay their bid. Now the equilibrium of the induced game isn't quite so bad, but still really hard to reason about (**lots** of work goes into characterizing equilibria of first-price auctions, and it's basically a complete mess - there's not enough information here on a model to see why, but check AGT book for more details). At minimum at least we can conclude that no one will ever bid above their value (instead they would rather bid exactly their value, no matter the other bids), but anything beyond this is hard to come by.*

A third alternative is to use the second-price auction. Now, you again give the item to the highest reported bid, but the winner pays the second-highest price. For this auction, there's an equilibrium that's actually quite easy to reason about. Observe that if the highest bid (aside from bidder i) is p_{-i} , then bidder i will win the item and pay p_{-i} if the bid $b_i \geq p_{-i}$, and lose the item and pay 0 otherwise. So they want to win if and only if their value exceeds p_{-i} , and this can be achieved always by submitting a bid of v_i . So it is a Nash equilibrium

for everyone to tell the truth in a second-price auction, and the maximum reported bid is selected (so the maximum actual value is also selected in equilibrium). Actually, telling the truth is a dominant strategy, discussed below.

So using the single-item auction as an example, we see that algorithms without payments might be really chaotic, but that algorithms with “the right” payments can induce simple equilibria and lead to good algorithms (even when we rely on the bidders themselves to tell their values). The specific notion we saw in the second-price auction is a dominant strategy.

Definition 1. *Strategy s is a dominant strategy for player i in a game if for all $s' \neq s$ that i could use, and all strategies \vec{t} that the other bidders might use, $P_i(s, \vec{t}) \geq P_i(s', \vec{t})$, and there exists a \vec{t}^* such that $P_i(s, \vec{t}^*) > P_i(s', \vec{t}^*)$. Here, $P_i(s, \vec{t})$ denotes the payoff enjoyed by player i when they use strategy s and the other players use strategies \vec{t} .*

It should be clear that if every bidder has a dominant strategy, then it is an *equilibrium*—no bidder gains by deviating. In a later lecture we’ll formalize several notions of equilibrium that have been studied in the literature; e.g., you might have heard of the Nash equilibrium.

1 The Vickrey-Clarke-Groves Auction

The first result we’ll discuss is seminal work of Vickrey, Clarke, and Groves (that contributed to a Nobel prize for Vickrey). It’s actually a combination of three separate single-author works, and was phrased very differently than the theorem statement below (they are all economists). To state the result, we’ll need to be clear about how bidders interact with payments.

Definition 2. *A bidder is quasi-linear if their utility for receiving value v and paying p is $v - p$. Bidders always want to maximize their utility.*

Theorem 1. *Let \mathcal{A} be an algorithm that takes as input $v_1(\cdot), \dots, v_n(\cdot)$ and outputs S_1, \dots, S_n maximizing $\sum_i v_i(S_i)$ over all partitions of $[m]$. Then a dominant-strategy truthful mechanism exists that requires only $n+1$ black-box calls to \mathcal{A} , and selects the partition $\mathcal{A}(v_1, \dots, v_n)$.*

Let’s parse this. First, we need to define dominant-strategy truthful. This simply means that telling the truth is a dominant strategy. That is, every bidder prefers to report their true $v_i(\cdot)$ than any other valuation function, no matter what valuation functions the other bidders report. This, together with the final guarantee, that the allocation is $\mathcal{A}(v_1, \dots, v_n)$ means that the actual welfare-maximizing allocation is chosen: because the mechanism is dominant strategy truthful, the reported valuations will be the actual valuations. Because the allocation maximizes welfare on the reported valuations, this means that the actual welfare-maximizing outcome is chosen. The fact that the mechanism only requires $n + 1$ calls to \mathcal{A} defines its runtime. If \mathcal{A} is poly-time, then the mechanism is poly-time too.

Proof. Consider the payment rule where each bidder is charged their *externality* on the other bidders. That is, we will charge every bidder the “harm” they cause the others by existing. More specifically, we look at the total happiness of all other bidders with you

in the picture versus the total happiness of all other bidders without you and charge the difference. That is, we do the following payments: below, $v = \langle v_1, \dots, v_n \rangle$, and $v_{-i}^* = \langle v_1, \dots, v_{i-1}, 0, v_{i+1}, \dots, v_n \rangle$. We'll also let $\mathcal{A}_i(v)$ denote the set awarded to bidder i on input v .

$$P_i(v) = \sum_{j \neq i} v_j(\mathcal{A}_j(v_{-i}^*)) - \sum_{j \neq i} v_j(\mathcal{A}_j(v)).$$

Again, this is exactly the difference between everyone else's utility without you in the picture versus with you in the picture. Let's now compute bidder i 's utility for submitting any bid v'_i (we'll use v' to denote $\langle v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n \rangle$).

$$\begin{aligned} U_i(v'_i, v_{-i}) &= v_i(\mathcal{A}_i(v')) - P_i(v') \\ &= v_i(\mathcal{A}_i(v')) - \sum_{j \neq i} v_j(\mathcal{A}_j(v_{-i}^*)) + \sum_{j \neq i} v_j(\mathcal{A}_j(v')) \\ &= \sum_j v_j(\mathcal{A}_j(v')) - \sum_{j \neq i} v_j(\mathcal{A}_j(v_{-i}^*)). \end{aligned}$$

Now let's look at these two terms. The first term is exactly the total welfare generated by the algorithm \mathcal{A} on bids v' but evaluated according to the real values v . The second term is completely out of bidder i 's control: it only depends on bids submitted by other bidders. So bidder i wants to maximize the first term to maximize their utility. The first term is clearly maximized when \mathcal{A} selects the true welfare-maximizing allocation, which happens when bidder i submits v_i . \square

At first this seems great! No matter the valuations, as long as we have an algorithm maximizing welfare, we can turn it into a truthful mechanism maximizing welfare. The drawback is that for basically any interesting class of valuations, maximizing welfare is NP-hard. Also unfortunately, the VCG reduction is incompatible with approximation: if \mathcal{A} does not exactly maximize welfare on every input, the output mechanism need not be truthful. (We'll prove this in HW-4.)

2 Truthful Approximation Algorithms: Single-Minded Agents

In this section, we'll consider a (very) special case of valuations and derive a truthful approximation algorithm. Here, each bidder's value will be *single-minded*. That is, for each i , there exists a special set S_i , and $v_i(S) = V_i$ if $S \supseteq S_i$, and 0 otherwise. Note that maximizing welfare is NP-hard by a reduction from independent set, even in this super special case.

Consider any graph $G = (V, E)$. We'll make a player for each node in V , and an item for each edge in E . Player i 's interest set S_i will be exactly the edges adjacent to them, and $V_i = 1$ for all i . Now it should be clear that we can simultaneously give a set P of players their interest sets if and only if they form an independent set in G . Therefore, maximizing welfare is equivalent to finding a large independent set. Note also that independent set is

NP-hard to approximate within $n^{1-\varepsilon}$ for any $\varepsilon > 0$, so welfare maximization for single-minded bidders is also NP-hard to approximate within $n^{1-\varepsilon}$ (or $m^{1/2-\varepsilon}$) for any $\varepsilon > 0$.

Now, we'll show a greedy mechanism that is truthful, and guarantees a $m^{1/2}$ -approximation. The mechanism is the following:

- Ask each bidder to report V_i, S_i .
- Sort the bidders so that $V_1/\sqrt{|S_1|} \geq V_2/\sqrt{|S_2|} \dots$
- Initialize $A = \emptyset$ (the set of awarded items). Starting from $i = 1$, visit bidder i and declare them a winner if and only if $S_i \cap A = \emptyset$. If so, update $A := A \cup S_i$.
- Award each winner their declared interest set S_i .
- Charge each winner the minimum V_i they could have reported and still been a winner.

For example, say that the bids are $(1, \{1\}), (1, \{1, 2, 3, 4\}), (4, \{1, 2\}), (4, \{3, 4\})$. Then the bids will be sorted so that the two 4s go first, followed by the two 1s. Both the 4s will win, the other two won't. $(4, \{1, 2\})$ will pay $\sqrt{2}$, because they will win if and only if they appear before $(1, \{1\})$ in the ordering. $(4, \{3, 4\})$ will pay 0, because they would win as long as they bid at least 0.

Theorem 2. *The greedy algorithm above guarantees a truthful \sqrt{m} -approximation.*

Proof. There are two things we want to prove: the mechanism is truthful and that it gets the desired approximation ratio. We'll do the approximation ratio first.

Let OPT denote the true optimal allocation. For each i that wins, let $OPT_i = \{j \in OPT, j \geq i, S_i \cap S_j \neq \emptyset\}$. That is, OPT_i is the players in OPT who are blocked by i (including itself). Clearly, $OPT = \cup_i OPT_i$, as everyone not blocked by any i would have been selected by Greedy. Now we'll prove that for all winners, $\sum_{j \in OPT_i} V_j \leq \sqrt{m} \cdot V_i$.

Note that every $j \in OPT_i$ appears in the greedy order after i , so we have $V_j \leq V_i \sqrt{|S_j|} / \sqrt{|S_i|}$. Summing over all $j \in OPT_i$, we have:

$$\sum_{j \in OPT_i} V_j \leq V_i / \sqrt{|S_i|} \cdot \sum_{j \in OPT_i} \sqrt{|S_j|}.$$

Using the Cauchy-Schwarz inequality ($x \cdot y \leq |x|_2 \cdot |y|_2$, for $x = \langle 1, \dots, 1 \rangle$ and $y = \langle \sqrt{|S_j|} \rangle_{j \in OPT_i}$), this is:

$$\sum_{j \in OPT_i} \sqrt{|S_j|} \leq \sqrt{|OPT_i|} \cdot \sqrt{\sum_{j \in OPT_i} |S_j|}.$$

Finally, observe that every $S_j \in OPT_i$ intersects S_i (by definition). Also, since OPT is an allocation, we must have $S_j \cap S_{j'} = \emptyset$ for all $j, j' \in OPT$. Therefore, we have $|OPT_i| \leq |S_i|$. Again since OPT is an allocation, we have $\sum_{j \in OPT_i} |S_j| \leq m$. Therefore,

the RHS above is upper bounded by $\sqrt{|S_i|} \cdot \sqrt{m}$. Plugging back into the first inequality yields that: $\sum_{j \in OPT_i} V_j \leq V_i \cdot \sqrt{m}$, as desired.

Finally, we want to claim that the mechanism is dominant strategy truthful. First, observe that, no matter what set you report, it is a dominant strategy to report your actual value for that set (either 0 or V_i). This is because based on the other bidders, the Greedy ordering induces a minimum bid you can submit (with that set) and still win it. You will win if and only if you bid above that value. If you report your true value for that set, you will always be on the correct side (just like second price).

Now, we want to argue that you should always report your true interest set. It's obvious you should never report an S that doesn't contain S_i , as this guarantees non-positive utility. If you report an S that strictly contains S_i , then this will only increase the bid you'd need to make in order to win (because you conflict with more, and because your bid gets divided by a bigger number). So you both want to bid your true value for the reported set, and report your true interest set.

To see this last claim more formally, consider all bids aside from your own. If you weren't in the picture, then there is some allocation that Greedy would select. Now imagine that you report the bid (V'_i, S'_i) . Then you will be allocated S'_i if and only if your bid winds up ahead of the first bidder who intersects with S'_i . Call this bidder j . Then you will be allocated if and only if $V'_i / \sqrt{|S'_i|} \geq V_j / \sqrt{|S_j|} \Leftrightarrow V'_i \geq V_j \sqrt{|S'_i|} / \sqrt{|S_j|}$. Note that if $S'_i \supseteq S_i$, the RHS can only go up. This is because first, $|S'_i| \geq |S_i|$, and second, the earliest bidder which intersects S'_i can only be earlier. So the price offered to you can only go up as you report a set strictly containing S_i , and doing so is dominated. \square

Beyond single-minded agents, there is a lot of work on designing truthful approximation algorithms for other agent valuations. In HW-4, we'll design a truthful $O(\log m)$ -approximation for submodular agents. Interestingly, it's conjectured that for submodular valuations we can get a truthful $O(1)$ -approximation, however, the current best bound is truthful $O((\log \log m)^2)$ -approximation [2].

Notes

This lecture is based on course notes from COS 521 at Princeton University, which are themselves based on Chapter 11 of the AGT Book [1].

References

- [1] Nisan, Roughgarden, and Tardos. "Algorithmic Game Theory". Vazirani (eds.), Cambridge University Press 2007.
- [2] Assadi, Kesselheim, and Singla. "Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier". Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA), 2021.