601.436/636 Algorithmic Game Theory **Lecturer:** Michael Dinitz
Topic: Combinatorial Auctions **Date:** 4/21/20
Scribe: Michael Dinitz

22.1 Introduction

Today we're going to talk more about combinatorial auctions, which we introduced last time.

22.2 Combinatorial Auctions and Single-Minded Bidders

Let's recall the setup for combinatorial auctions.

- Players $[n]$.
- m items M.
- Outcome: each bidder i gets a set (bundle) $S_i \subseteq M$, where $S_i \cap S_j = \emptyset$ for all $i \neq j$. Note that this means there are *many* outcomes: $(n+1)^m$.
- Valuations: each player i has a valuation $v_i(S)$ for each $S \subseteq M$, which is the value that player i gets from receiving bundle S. Note that this means that player i does not care about what any other player j gets: they just care about what they get. As we discussed last time, this means that each player has 2^m private parameters, which is a lot. We'll assume (for today) that $v_i(\emptyset) = 0$ and that $v_i(S) \leq v_i(T)$ for all $S \subseteq T \subseteq M$, since player i could always just throw away the items in $T \setminus S$.
- The social welfare of outcome $(S_1, S_2, \ldots, S_n)$ is $\sum_{i=1}^n v_i(S_i)$.
- The utility of a player in an outcome is its value minus the price it is charged.

What mechanism should we use for combinatorial auctions? The classical answer is VCG, but as talked about last time, this has a number of issues. VCG can be bad for revenue, bids are massive and hard to communicate, and finding the welfare-maximizing outcome/allocation can be an extremely difficult computational problem. As computer scientists, we're mostly concerned with this last issue. One obvious approach would be to use an approximation algorithm for social welfare, but even if a good approximation algorithm exists we have another problem: if we try to define prices as we did before (the externality of a player), then the fact that we're approximating social welfare instead of computing it may result in prices that are not incentive compatible. So, unfortunately, we (to some extent) need to start from scratch every time we encounter a new combinatorial auction setting.

22.2.1 Single-Minded Bidders

We're going to spend most of today on one of the simplest combinatorial auctions: single-minded bidders. This is an important setting in practice since it models some realistic behavior, and is also one of the simplest non-single parameter environments.

This setting is simple to describe. Each bidder i has a (private) set $T_i \subseteq M$, and there is a private parameter $v_i \in \mathbb{R}^+$ such that

$$v_i(S) = \begin{cases} v_i & \text{if } S \supseteq T_i \\ 0 & \text{otherwise} \end{cases}$$

So this is *almost* a single-parameter environment, since there are now private parameters v_i and T_i. A bid in this context is a pair (b_i, S_i), which is easy enough to write down and transmit from the player to the auctioneer. In the homework you need to think about a related setting where T_i is public, which is truly a single-parameter environment, but now we're going to think about this trickier case.

22.3 Greedy Mechanism

Like last time, we need to think about simultaneously designing the allocation rule and the prices – there's no general theorem which will let us derive prices from an allocation like Myerson or VCG.

22.3.1 The Mechanism

We'll use the following allocation rule/algorithm:

- Sort the bidders so
$$\frac{b_1}{\sqrt{|S_1|}} \geq \frac{b_2}{\sqrt{|S_2|}} \geq \cdots \geq \frac{b_n}{\sqrt{|S_n|}}.$$
We'll see why we use this specific ordering later, when we analyze the approximation of social welfare. But it's worth noting that for incentive compatibility we just need an ordering that is increasing in b_i and decreasing in $|S_i|$ – any such ordering will do. Being linear in b_i and $1/\sqrt{|S_i|}$ is a choice for social welfare maximization.
- $W = \emptyset$
- For i from 1 to n: if $S_i \cap (\cup_{j \in W} S_j) = \emptyset$, add i to W.
- Return the allocation which gives S_i to player i if and only if $i \in W$.

To define prices, we'll (obviously) set $p_i = 0$ if $i \notin W$. If $i \in W$, the price is a little more complicated. Fix a bidder $i \in W$. Let $\alpha(i)$ be the minimum index (bidder in the sorted order) such that $S_i \cap S_{\alpha(i)} \neq \emptyset$ and $S_k \cap S_{\alpha(i)} = \emptyset$ for all $k < \alpha(i)$ with $k \neq i$ and $k \in W$. In other words, $\alpha(i)$ is the first player who lost *because* of bidder i: if i had not been participating then $\alpha(i)$ would have been in W. If no such $\alpha(i)$ exists, then we set $p_i = 0$. Otherwise, we set

$$p_i = \frac{b_{\alpha(i)}}{\sqrt{|S_{\alpha(i)}|/|S_i|}} = b_{\alpha(i)}\sqrt{\frac{|S_i|}{|S_{\alpha(i)}|}}$$

22.3.2 Analysis

This mechanism is obviously polynomial time, and obviously outputs a valid allocation. So we need to show that it is polynomial time and approximately maximizes the social welfare.

22.3.2.1 Incentive Compatibility

Let's prove that this mechanism has two important properties: monotonicity and critical payment.

- Monotonicity: if i wins with bid (b_i, S_i), then it would win with any bid (b_i', S_i') such that $b_i' \geq b_i$ and $S_i' \subseteq S_i$. This is trivial to prove: increasing b_i or decreasing S_i can only move i earlier in the greedy ordering, so if S_i did not conflict with any earlier $j \in W$ then it still will not conflict after moving earlier.

- Critical Payment: if i wins then the price it pays is the smallest x such that i would still win if it had bid (x, S_i).

  To prove this, first consider the case when $\alpha(i)$ does not exist. Then there is no other bidder who fails to win because of i's participation, and thus if i was last in the ordering it would still win. So it is charged 0, which is the critical payment.

  Now suppose that $\alpha(i)$ does exist. Then i will still win as long as it appears before $\alpha(i)$ in the ordering, and if it appears after $\alpha(i)$ then $\alpha(i)$ will win and so i will not (since $S_i \cap S_{\alpha(i)} \neq \varnothing$). Thus the critical payment is the x such that $\frac{x}{\sqrt{|S_i|}} = \frac{b_{\alpha(i)}}{\sqrt{|S_{\alpha(i)}|}}$. Solving for x, we get that the critical payment is

  $$b_{\alpha(i)}\sqrt{\frac{|S_i|}{|S_{\alpha(i)}|}} = p_i$$

So the greedy mechanism has these two properties.

Theorem 22.3.1. *Any mechanism where losers pay* 0 *which has both the monotonicity and critical payment properties is incentive compatible.*

Proof. Fix player $i \in [n]$ and all bids other than i's. Let $u(b, S)$ be the utility that player i would get by bidding (b, S), so $u(b, S) = v_i(S) - p_i(b, S)$. By the critical payment property, we know that $p_i(b, S) = \inf\{x : i \text{ wins with bid } (x, S)\}$.

Let's first show that truthful bidding results in nonnegative utility. If player i bids truthfully then it gets either utility 0 (if it does not win) or utility $v_i(T_i) - p_i(b, T_i) = v_i - p_i(b, T_i) \geq 0$ (by the definition of $p_i(b, S)$).

Now let's show that telling the truth is a dominant strategy, i.e., that $u(v_i, T_i) \geq u(b, S)$ for all bids (b, S). This is obvious if (b, S) is a losing bid (since we just proved that telling the truth gives nonnegative utility). This is also obvious if $T_i \notin S$, since then even if player i won with bid (b, S) they would receive a bundle which has zero value to them. So without loss of generality, assume that $T_i \subseteq S$ and that (b, S) is a winning bid.

We'll prove that $u(v_i, T_i) \geq u(b, S)$ by a two-step argument: first we'll show that $u(b, T_i) \geq u(b, S)$, and then we'll show that $u(v_i, T_i) \geq u(b, T_i)$.

For the first step, we know that both (b, S) is a winning bid and thus (b, T_i) would also be a winning bid by the monotonicity property. So $u(b, T_i) = v_i - p_i(b, T_i)$ and $u(b, S) = v_i - p_i(b, S)$, and hence we just need to show that $p_i(b, T_i) \leq p_i(b, S)$. By the critical payment property, $p_i(b, S)$ is the minimum x such that (x, S) would win. By the monotonicity property, this means that (x, T_i) would also win, and thus the minimum y such that (y, T_i) would win is at most x, and so (again by critical payment) $p_i(b, T_i) = y \leq x = p_i(b, S)$.

For the second (and final) step, we need to show that $u(v_i, T_i) \geq u(b, T_i)$. As discussed, by monotonicity we have that (b, T_i) is a winning bid. If (v_i, T_i) is a winning bid, then $u(v_i, T_i) = u(b, T_i)$ since the value of the bundle is the same and the price is the same (by the critical payment property). If (v_i, T_i) is not a winning bid then $u(v_i, T_i) = 0$, but on the other hand $p_i(b, T_i) \geq v_i$ and thus $u(b, T_i) \leq 0$. □

22.3.2.2 Social Welfare

Let OPT be the winners in a welfare-maximizing allocation, and W be the winners from our greedy mechanism.

Theorem 22.3.2. $\sum_{i \in OPT} v_i \leq \sqrt{m} \cdot \sum_{i \in W} v_i$

Proof. Let $OPT_i = \{j \in OPT : j \geq i \wedge T_i \cap T_j \neq \varnothing\}$. Note that $OPT = \cup_{i \in W} OPT_i$, since every player $j \in OPT$ is either in W (in which case it is in OPT_j) or was not added to W because some player $i < j$ with $T_i \cap T_j \neq \varnothing$ was added to W, in which case j is in OPT_i.

Let $j \in OPT_i$. Then $j \geq i$, so

$$\frac{v_j}{\sqrt{|T_j|}} \leq \frac{v_i}{\sqrt{|T_i|}}$$
$$\implies v_j \leq \frac{v_i}{\sqrt{|T_i|}} \sqrt{|T_j|}$$

Since this was true for all $j \in OPT_i$, we can add them all up to get

$$\sum_{j \in OPT_i} v_j \leq \frac{v_i}{\sqrt{|T_i|}} \sum_{j \in OPT_i} \sqrt{|T_j|} \leq \frac{v_i}{\sqrt{|T_i|}} \sqrt{|OPT_i|} \sqrt{\sum_{j \in OPT_i} |T_j|}$$

This last inequality is due to the Cauchy-Schwarz inequality, a super important and useful inequality which sounds obvious: Cauchy-Schwarz says that $|\langle x, y \rangle| \leq \|x\| \cdot \|y\|$ for all vectors x and y (where $\langle \cdot \rangle$ denotes the inner product and $\| \cdot \|$ denotes the length in the usual Euclidean norm). If we set

x to be the all 1's vector and y to be the vector with $y_j = \sqrt{|T_j|}$, then Cauchy-Schwarz implies the final inequality above.

Now let's bound $|OPT_i|$. This is actually pretty easy to do: every $j \in OPT_i$ has $T_j \cap T_i \neq \varnothing$, but for any $j, j' \in OPT_i$ with $j \neq i$ we know that $T_j \cap T_{j'} = \varnothing$ (since they are both able to get their desired bundles in OPT). So every every player in OPT_i conflicts with i in a *different* item, and so $|OPT_i| \leq |T_i|$. So if we continue the above series of inequalities, we get

$$\sum_{j \in OPT_i} v_j \leq v_i \sqrt{\sum_{j \in OPT_i} |T_j|} \leq v_i \sqrt{m},$$

where the last inequality is again because players who win in OPT cannot have overlapping desired bundles. Thus

$$\sum_{i \in OPT} v_i \leq \sum_{i \in W} \sum_{j \in OPT_i} v_j \leq \sum_{i \in W} v_i \sqrt{m} = \sqrt{m} \sum_{i \in W} v_i,$$

as claimed. □

22.4 Spectrum Auctions

I'm not going to talk much about spectrum auctions, but they're a very interesting "case study" in real-life combinatorial auctions. Chapter 8 in Roughgarden's book talks all about them, and I suggest you read it.

At a super high level, spectrum auctions involve a government auctioning off electromagnetic spectrum to various companies. So the items are blocks of EM spectrum, and the bidders are companies interesting in owning such a block (e.g., a TV station, radio station, telecommunications company, etc.). These auctions really happen, and they involve large amounts of money and important public goods, so they really matter. How should they be run?

One answer would be VCG, but as we've discussed, this has a number of issues which make it impractical. An important one here is that the complexity of the auction (e.g., the size of the bids) should be polynomial in m (the number of blocks of spectrum being sold) rather than exponential like in VCG.

A natural approach would be to sell the m items separately. For example, we could do this sequentially, where we first sell one block using our normal auction (like a second-price auction), then another block, etc. It turns out that if you do this, even in simple cases like identical items, the fact that the players know that you're going to hold sequential auctions means that bidding truthfully is no longer a dominant strategy – these auctions are not incentive compatible. See Roughgarden's book for more details of an example which demonstrates this. And this is also a huge problem in real life: in 2000, Switzerland auctioned off three blocks of spectrum using three sequential second-price auctions. The first two blocks were basically identical (28 MHz blocks), while the third was for a better item (a 56 MHz block). The first block went for 121 million francs and the second went for 134 million, but the third (and best) went for 55 million francs. Clearly there were some non-optimal bids, so truthtelling was not a dominant strategy!

We could also sell the m items separately but simultaneously, i.e., just run m simultaneous but separate second-price auctions. But now it is highly unclear how a payer should bid. Suppose that all of the items are identical and you want any one of them – should you bid low on a lot of items, or high on just one? What does it even mean to be truthful here? Because it's not clear how to bid, these auctions can have very strange behavior, and tend to be very bad on revenue. In real life, in 1990 New Zealand used this format and made only $36 million rather than the projected $250 million. And the auction was kind of crazy: one block had a high bid of $100,000 while the second-highest bid (and thus the price) was $6, and another block had a high bid of $7 million and second-highest of $5,000. Clearly these bids were not the true valuations of the companies, and (even worse) everything was made public after the auction ended and so everyone could see how much money was "left on the table".

The most popular solution is something called a "simultaneous ascending auction" (SAA). These are a bit more complicated, so I'm not going to get into them, but are basically like running simultaneous classical "raise your paddle to bid" auctions. Note that if we're just running one auction, then as we've discussed this classical auction (sometimes called an *English auction*) is basically the second-price auction. But if we run them simultaneously then this is no longer true, since the fact that bids are public during the process changes things. There are a number of complications (SAAs usually have "activity rules" that sometimes forbid certain bids, for example), but SAAs have a number of nice properties. See Roughgarden Chapter 8 for much more discussion.