

General Combinatorial Auctions in Polynomial Time

Instructor: Thomas Kesselheim

Last time, we got to know the famous VCG mechanism. It almost gives us everything we ask for: It is reasonably easy to understand, dominant-strategy incentive compatible (truthful), and in the dominant-strategy equilibrium the social welfare is maximized. Unfortunately, there is one drawback, which is that it requires us to maximize the declared welfare exactly. Approximate solutions are not enough.

There are many ways to still design truthful mechanisms outside the VCG framework—we have seen examples for single-parameter settings, in which we only had to come up with monotone algorithms. Beyond single-parameter settings, things become more difficult. However, there are still many ideas to build truthful mechanisms. We will get to know one such approach for combinatorial auctions today.

1 Setting

Recall that in a *combinatorial auction*, we have a set N of n bidders and a set M of m items. Each bidder i has a private valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$, defining a non-negative value of each subset of items. We will not make any assumptions on v_i today except for normalization ($v_i(\emptyset) = 0$) and monotonicity: If $S \subseteq S'$, then $v_i(S) \leq v_i(S')$. The set of feasible allocations is given as $X = \{(S_1, \dots, S_n) \mid S_i \cap S_{i'} = \emptyset \text{ for } i \neq i'\}$.

Our goal is to come up with a direct mechanism $M = (f, p)$, consisting of an allocation function $f: V \rightarrow X$ and a payment rule $p: V \rightarrow \mathbb{R}^n$, where $V = V_1 \times \dots \times V_n$ and V_i is the set of all monotone normalized functions $2^M \rightarrow \mathbb{R}_{\geq 0}$. Today, our mechanism will be randomized. That is, both $f(b)$ and $p(b)$ are random variables.

The approximation guarantee will be with respect to the *expected social welfare*. We evaluate the allocation f by $\mathbf{E}[\sum_{i \in N} v_i(f_i(v))]$.

2 Configuration Linear Program

Our mechanism will be based on the configuration linear program. For $i \in N$, $S \subseteq M$, we introduce variables $x_{i,S}$ that indicate if bidder i gets exactly the items in set S . We can now write the following linear program

$$\begin{array}{ll}
 \text{maximize} & \sum_{i \in N} \sum_{S \subseteq M} b_i(S) x_{i,S} \\
 \text{subject to} & \sum_{i \in N} \sum_{S: j \in S} x_{i,S} \leq 1 & \text{for all } j \in M \\
 & \sum_{i \in N} x_{i,S} \leq 1 & \text{for all } i \in N \\
 & x_{i,S} \geq 0 & \text{for all } i \in N, S \subseteq M
 \end{array}$$

Note that we actually would like to find the optimal *integral* solution, that is $x_{i,S} \in \{0, 1\}$. This, of course, is hard. Instead, we relax this assumption and allow fractional assignments of variables. Then, this LP can be solved in polynomial time. Note that this is not quite obvious as we have exponentially many variables. However, as the number of constraints is polynomial, one can do some tricks given access to the right queries.

Given a set $N' \subseteq N$ of bidders, let $\text{opt}_{N'}^*$ be the value of the optimal fractional solution to this LP if we restrict it to only consider bidders from the set N' . Observe that $\text{opt}_{N'}^*$ is an upper

bound on the optimal declared welfare, i.e., $\text{opt}_N^* \geq \max_{(S_1, \dots, S_n) \in X} \sum_{i \in N} b_i(S_i)$. Therefore, opt_N^* will be our point of comparison.

3 The Mechanism

Our mechanism is a simplified version of one given by Dobzinski, Nisan, and Schapira (STOC 2006). It is, in fact, two mechanisms. Each of them is run with probability $\frac{1}{2}$. Mechanism (A) is simply a second-price auction. Mechanism (B) uses a random subset of the bidders that defines an item price for the other bidders. Given these item prices, each bidder then chooses the bundle that she prefers most.

With probability $\frac{1}{2}$ run Mechanism (A), otherwise run Mechanism (B).

Mechanism (A):

Let i^* be the bidder of maximum $b_i(M)$. Assign him all items and let him pay $\max_{i \neq i^*} b_i(M)$.

Mechanism (B):

- For each $i \in N$: Add i to *STAT* with probability $\frac{1}{2}$ independently. Set $FIXED := N \setminus STAT$.
- Compute opt_{STAT}^* . Do not assign any items to *STAT*.
- Approach bidders in *FIXED* in order of increasing index. Offer each item at price $p := \frac{1}{8m} \text{opt}_{STAT}^*$. That is, to bidder i , assign S that maximizes $b_i(S) - |S|p$, where S is a subset of the remaining items.

Why all these choices make sense to some extent will become visible in the analysis.

4 Truthfulness

We will show that this mechanism is truthful. Actually, in case of randomized mechanism, there are multiple ways to define truthfulness. Our mechanism fulfills the strongest of them: No matter which random bits are used in the course of the mechanism, no bidder is ever better off by lying about the valuation.

Theorem 15.1. *Fixing any outcome of the random coin flips, this mechanism is truthful.*

Proof. Having fixed the randomization, we either definitely run Mechanism (A) or definitely run Mechanism (B). Mechanism (A) is truthful because it is a second-price auction. In Mechanism (B), a bidder can belong to either *STAT* or *FIXED*. If it belongs to *STAT*, then the utility is always 0 because the bidder will neither get anything nor pay anything. If it belongs to *FIXED*, then it cannot influence the item prices because these are set only by *STAT*. The bidder is assigned whatever maximizes the utility given these prices, so also here there is no point in lying. \square

Note that in this proof, we made use of the strange design of this mechanism: For example, it would sound more reasonable to run both Mechanism (A) and (B) and to then return the better of the two allocations. However, we have to be very careful to keep the mechanism truthful: It might be that one of the bidders manipulates, for example $b_i(M)$, to get a different bundle. This is also the reason why bidders in *STAT* cannot get anything. Otherwise they might be tempted to misreport some value because this influences the price.

5 Approximation Guarantee

We will show the following guarantee.

Theorem 15.2. *Fix a bid vector b and let $S_1, \dots, S_n \subseteq M$ be random variables indicating the assignment by the mechanism. Then $\mathbf{E} [\sum_{i \in N} b_i(S_i)] = \Omega\left(\frac{1}{\sqrt{m}}\right) \text{opt}_N^*$.*

Note that by the hardness-of-approximation result, we cannot hope to get any algorithm whose guarantee is a lot better than \sqrt{m} . For the sake of readability, we do not care much about the constants involved here.

To prove the theorem, we make a case distinction, asking whether there is a bidder i of very high value. In this case, we use that with probability $\frac{1}{2}$, Mechanism (A) is run.

Observation 15.3. *If there is a bidder i for which $b_i(M) \geq \frac{\text{opt}_N^*}{\sqrt{m}}$, then Mechanism (A) returns a solution of declared value at least $\frac{\text{opt}_N^*}{\sqrt{m}}$.*

If there is no such bidder, then we rely on Mechanism (B). The idea of partitioning the bidders is that set $STAT$ should give us a reasonable estimate of the values in $FIXED$. Only then, p can work as a good price for the item. Fortunately, if there is no dominant bidder, meaning that $b_i(M) < \frac{\text{opt}_N^*}{\sqrt{m}}$, then this happens with good probability.

Lemma 15.4. *If there is a bidder i for which $b_i(M) \geq \frac{\text{opt}_N^*}{\sqrt{m}}$, with probability at least $1 - \frac{4}{\sqrt{m}}$, both opt_{STAT}^* and opt_{FIXED}^* are at least $\frac{1}{4}\text{opt}_N^*$.*

Proof. Let x^* indicate an optimal solution to the configuration LP including all bidders. That is, $\text{opt}_N^* = \sum_{i \in N} \sum_{S \subseteq M} b_i(S) x_{i,S}^*$. Now, observe that $\text{opt}_{STAT}^* \geq \sum_{i \in STAT} \sum_{S \subseteq M} b_i(S) x_{i,S}^*$ because we could simply take the allocation x^* , remove all allocations to bidders not in $STAT$, and get a feasible solution. opt_{STAT}^* in turn is the value of the best feasible solution. By the same reasoning $\text{opt}_{FIXED}^* \geq \sum_{i \in FIXED} \sum_{S \subseteq M} b_i(S) x_{i,S}^* = \sum_{i \in N} \sum_{S \subseteq M} b_i(S) x_{i,S}^* - \sum_{i \in STAT} \sum_{S \subseteq M} b_i(S) x_{i,S}^* = \text{opt}_N^* - \sum_{i \in STAT} \sum_{S \subseteq M} b_i(S) x_{i,S}^*$.

So, it suffices to understand the random variable $Y = \sum_{i \in STAT} \sum_{S \subseteq M} b_i(S) x_{i,S}^*$. Whenever $\frac{1}{4}\text{opt}_N^* \leq Y \leq \frac{3}{4}\text{opt}_N^*$, then both opt_{STAT}^* and opt_{FIXED}^* are at least $\frac{1}{4}\text{opt}_N^*$.

We will use Chebyshev's inequality that states $\Pr[|Y - \mathbf{E}[Y]| \geq \alpha] \leq \frac{\text{Var}(Y)}{\alpha^2}$ for all $\alpha > 0$. Let us define $Y_i = \sum_{S \subseteq M} b_i(S) x_{i,S}^*$ if $i \in STAT$ and 0 otherwise, so $Y = \sum_{i \in N} Y_i$. The expectation of Y can now be calculated using linearity of expectation

$$\mathbf{E}[Y] = \mathbf{E}\left[\sum_{i \in N} Y_i\right] = \sum_{i \in N} \mathbf{E}[Y_i] = \frac{1}{2} \sum_{i \in N} \sum_{S \subseteq M} b_i(S) x_{i,S}^* = \frac{1}{2} \text{opt}_N^*$$

The variance of a single Y_i is given as

$$\text{Var}(Y_i) = \frac{1}{4} \left(\sum_{S \subseteq M} b_i(S) x_{i,S}^* \right)^2 \leq \frac{1}{4} \left(\frac{\text{opt}_N^*}{\sqrt{m}} \right) \left(\sum_{S \subseteq M} b_i(S) x_{i,S}^* \right),$$

where we used that there is no dominant bidder. As the Y_i are independent, the variance of their sum is the sum of their variances. Therefore, we get for the variance of Y

$$\text{Var}(Y) = \text{Var}\left(\sum_{i \in N} Y_i\right) = \sum_{i \in N} \text{Var}(Y_i) \leq \frac{1}{4} \left(\frac{\text{opt}_N^*}{\sqrt{m}} \right) \left(\sum_{i \in N} \sum_{S \subseteq M} b_i(S) x_{i,S}^* \right) = \frac{1}{4\sqrt{m}} (\text{opt}_N^*)^2.$$

So, if we set $\alpha = \frac{1}{4}\text{opt}_N^*$, then $\Pr\left[|Y - \frac{1}{2}\text{opt}_N^*| \geq \frac{1}{4}\text{opt}_N^*\right] \leq \frac{\frac{1}{4\sqrt{m}} (\text{opt}_N^*)^2}{\left(\frac{1}{4}\text{opt}_N^*\right)^2} = \frac{4}{\sqrt{m}}$. \square

The last lemma is to show that if *STAT* and *FIXED* are such a “good” partition, then using the item price p will give us a reasonably good outcome.

Lemma 15.5. *If for all bidders i we have $b_i(M) < \frac{\text{opt}_N^*}{\sqrt{m}}$, and *STAT* and *FIXED* are chosen such that both opt_{STAT}^* and opt_{FIXED}^* are at least $\frac{1}{4}\text{opt}_N^*$, then Mechanism (B) returns a solution of declared welfare at least $\frac{\text{opt}_N^*}{512\sqrt{m}}$ irrespective of precise random outcomes.*

Proof. By our assumption we have $\text{opt}_{STAT}^*, \text{opt}_{FIXED}^* \geq \frac{1}{4}\text{opt}_N^*$. Also, $\text{opt}_{STAT}^*, \text{opt}_{FIXED}^* \leq \text{opt}_N^*$ because both are feasible choices also if all bidders are allowed.

Let A be the set of items that are assigned. For every item that is sold, some bidder pays us p . Importantly, the payments never exceed the bids. Therefore, the declared welfare of the final allocation is at least $p|A|$, because this is the money that the mechanism collects.

We will show that A has size at least $\frac{1}{16}\sqrt{m}$. This then gives us that the allocation has declared welfare $p|A| \geq \frac{\text{opt}_{STAT}^*}{8m} \cdot \frac{1}{16}\sqrt{m} \geq \frac{\text{opt}_N^*}{16 \cdot 8 \cdot 4\sqrt{m}} = \frac{\text{opt}_N^*}{512\sqrt{m}}$.

Let B denote the set of bidders that get at least one item. Naturally, we have $|A| \geq |B|$. Consider a pair (i, S) of a buyer $i \in \text{FIXED}$ and a bundle $S \subseteq M$. Observe that at least one of the following conditions hold:

- (1) $b_i(S) < |S|p$
- (2) one of the items in S gets allocated, so $S \cap A \neq \emptyset$
- (3) bidder i gets an item, so $i \in B$.

Let \mathcal{T}_k denote the pairs for which the k th condition holds.

Let x^{FIXED} be an optimal solution to the LP relaxation that only considers bidders in *FIXED*. That is, $\sum_{i \in \text{FIXED}} \sum_{S \subseteq M} b_i(S)x_{i,S}^{\text{FIXED}} = \text{opt}_{FIXED}^*$. As every pair (i, S) is contained in at least one of the three sets, we have

$$\text{opt}_{FIXED}^* \leq \sum_{(i,S) \in \mathcal{T}_1} b_i(S)x_{i,S}^{\text{FIXED}} + \sum_{(i,S) \in \mathcal{T}_2} b_i(S)x_{i,S}^{\text{FIXED}} + \sum_{(i,S) \in \mathcal{T}_3} b_i(S)x_{i,S}^{\text{FIXED}} .$$

First, let us observe that we can bound the contribution of pairs that are too expensive by

$$\begin{aligned} \sum_{(i,S) \in \mathcal{T}_1} b_i(S)x_{i,S}^{\text{FIXED}} &< \sum_{(i,S) \in \mathcal{T}_1} |S|px_{i,S}^{\text{FIXED}} \\ &\leq \sum_{i \in \text{FIXED}} \sum_{S \subseteq M} |S|px_{i,S}^{\text{FIXED}} \\ &= p \sum_{j \in M} \sum_{i \in \text{FIXED}} \sum_{S: j \in S} x_{i,S}^{\text{FIXED}} . \end{aligned}$$

The first inequality uses the definition of \mathcal{T}_1 whereas the second inequality simply extends the sum to all pairs (i, S) . Next, we use that $\sum_{i \in \text{FIXED}} \sum_{S: j \in S} x_{i,S}^{\text{FIXED}} \leq 1$ for all $j \in M$ as x^{FIXED} is a feasible LP solution. Therefore

$$\sum_{(i,S) \in \mathcal{T}_1} b_i(S)x_{i,S}^{\text{FIXED}} \leq pm = \frac{1}{8}\text{opt}_{STAT}^* .$$

By our assumption $\text{opt}_{STAT}^* \leq \text{opt}_N^* \leq 4\text{opt}_{FIXED}^*$. So,

$$\sum_{(i,S) \in \mathcal{T}_1} b_i(S)x_{i,S}^{\text{FIXED}} \leq \frac{1}{2}\text{opt}_{FIXED}^*$$

and therefore

$$\frac{1}{2}\text{opt}_{FIXED}^* \leq \sum_{(i,S) \in \mathcal{T}_2} b_i(S)x_{i,S}^{\text{FIXED}} + \sum_{(i,S) \in \mathcal{T}_3} b_i(S)x_{i,S}^{\text{FIXED}} .$$

Next, we bound the contribution to this sum by \mathcal{T}_2 in terms of the number of items that are sold

$$\begin{aligned} \sum_{(i,S) \in \mathcal{T}_2} b_i(S) x_{i,S}^{FIXED} &= \sum_{j \in A} \sum_{i \in FIXED} \sum_{S: j \in S} b_i(S) x_{i,S}^{FIXED} \\ &\leq \frac{\text{opt}_N^*}{\sqrt{m}} \sum_{j \in A} \sum_{i \in FIXED} \sum_{S: j \in S} x_{i,S}^{FIXED} \\ &\leq \frac{\text{opt}_N^*}{\sqrt{m}} |A| , \end{aligned}$$

where again we use that $\sum_{i \in FIXED} \sum_{S: j \in S} x_{i,S}^{FIXED} \leq 1$ for all $j \in M$.

And finally, we bound the contribution of \mathcal{T}_3 using the number of bidders that get at least one item

$$\sum_{(i,S) \in \mathcal{T}_2} b_i(S) x_{i,S}^{FIXED} = \sum_{i \in B} \sum_{S \subseteq M} b_i(S) x_{i,S}^{FIXED} \leq \sum_{i \in B} b_i(M) \sum_{S \subseteq M} x_{i,S}^{FIXED} \leq \sum_{i \in B} b_i(M) \leq \frac{\text{opt}_N^*}{\sqrt{m}} |B| \leq \frac{\text{opt}_N^*}{\sqrt{m}} |A| .$$

Here, we use that $\sum_{S \subseteq M} x_{i,S}^{FIXED} \leq 1$ for all i and $|A| \geq |B|$.

In combination this gives us

$$\frac{1}{2} \text{opt}_{FIXED}^* \leq 2 \frac{\text{opt}_N^*}{\sqrt{m}} |A| .$$

As we have $\text{opt}_{FIXED}^* \geq \frac{1}{4} \text{opt}_N^*$, this implies the desired bound on $|A|$. □

Now, we can plug together these lemmas to prove the theorem.

Proof of Theorem 15.2. Again, we follow the case distinction. If there is a bidder i for which $b_i(M) \geq \frac{\text{opt}_N^*}{\sqrt{m}}$, then we run Mechanism (A) with probability $\frac{1}{2}$. We also run Mechanism (B) with probability $\frac{1}{2}$ but we simply use non-negativity:

$$\mathbf{E} \left[\sum_{i \in N} b_i(S_i) \right] \geq \mathbf{Pr} [\text{run Mechanism (A)}] \frac{\text{opt}_N^*}{\sqrt{m}} + \mathbf{Pr} [\text{run Mechanism (B)}] 0 = \frac{\text{opt}_N^*}{2\sqrt{m}} .$$

We do the same if for all bidders i we have $b_i(M) < \frac{\text{opt}_N^*}{\sqrt{m}}$. This time we get

$$\begin{aligned} \mathbf{E} \left[\sum_{i \in N} b_i(S_i) \right] &\geq \mathbf{Pr} \left[\text{run Mechanism (B) and } \text{opt}_{STAT}^*, \text{opt}_{FIXED}^* \geq \frac{1}{4} \text{opt}_N^* \right] \frac{\text{opt}_N^*}{512\sqrt{m}} \\ &\quad + \mathbf{Pr} \left[\text{run Mechanism (A) or } \text{opt}_{STAT}^* < \frac{1}{4} \text{opt}_N^* \text{ or } \text{opt}_{FIXED}^* < \frac{1}{4} \text{opt}_N^* \right] \\ &= \frac{1}{2} \left(1 - \frac{4}{\sqrt{m}} \right) \frac{\text{opt}_N^*}{512\sqrt{m}} = \Omega \left(\frac{1}{\sqrt{m}} \right) \text{opt}_N^* . \end{aligned}$$

□

6 Outlook

This mechanism is not nearly as elegant and clean as VCG. However, it is the best that people have come up with when it comes to designing truthful mechanisms that run in polynomial time. For example, it is still an open question whether randomization is actually necessary. An alternative to such complicated truthful mechanisms is to analyze simpler, non-truthful ones. This will be our next topic in this course.