

Greedy Algorithms for Combinatorial Auctions

Thomas Kesselheim

Last Update: December 3, 2020

Today, we will consider a subclass of combinatorial auctions. Our goal will be to design mechanisms that are truthful, run in polynomial time, and yield good approximation guarantees.

The subclass will be a single-parameter settings, for which we already know by Myerson's Lemma that we have to confine ourselves to monotone outcome rules.

1 Combinatorial Auctions

Recall combinatorial auctions.

Definition 12.1 (Combinatorial Auction). *In a combinatorial auction, a set of m items M shall be allocated to a set of n bidders $\mathcal{N} = \{1, \dots, n\}$. The bidders have private values for bundles of items. The goal is to maximize social welfare.*

- *Feasible allocations:* $X = \{(S_1, \dots, S_n) \in (2^M)^n \mid S_i \cap S_j = \emptyset, i \neq j\}$
- *Valuation functions:* $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}, i \in \mathcal{N}$ (private)
- *Objective:* Maximize social welfare $\sum_{i=1}^n v_i(S_i)$

We will generally assume free disposal, i.e., $v_i(S) \geq v_i(T)$ for $T \subseteq S$, and that valuations are normalized, i.e., $v_i(\emptyset) = 0$.

We will focus on the case where each bidder is interested in a single bundle of items. We will call these bidders single minded.

Definition 12.2 (Single-Minded Bidders). *Bidders are called single minded if, for every bidder $i \in \mathcal{N}$, there exists a bundle $S_i^* \subseteq M$ and a value $v_i \in \mathbb{R}_{\geq 0}$ such that*

$$v_i(T) = \begin{cases} v_i & \text{if } T \supseteq S_i^*, \\ 0 & \text{otherwise.} \end{cases}$$

We call a bidder that is granted their bundle or a superset a winner, and we say that this bidder wins the bundle.

Note that we again overload notation and use v_i to both denote the function as well as the value that a player has for winning. This is in line with our notation for single-item auctions and single-parameter problems.

Our objective will be to maximize *social welfare* $\sum_{i \in \mathcal{N}} v_i(x)$. In the case of single-minded bidders, we can write it simply as $\sum_{i \in W} v_i$, where $W \subseteq \mathcal{N}$ is the set of winners.

Example 12.3 (Single-Minded Combinatorial Auction). *There are two items a and b and three bidders Red, Green, and Blue. Red has a value of 10 for $\{a\}$, Green has a value of 14 for the set $\{a, b\}$, and Blue has a value of 8 for $\{b\}$. Social welfare is maximized by allocating $\{a\}$ to Red and $\{b\}$ to Blue.*

Our goal will be to design a mechanism for combinatorial auction with single-minded bidders. We will assume that the bundle S_i^* that bidder i is interested in is *public* and only the valuation v_i is *private*. As we will see, this turns the problem into a single-parameter problem, to which our previous results apply.

This mechanism $\mathcal{M} = (f, p)$ consists of an outcome rule f and a payment rule p . Both take as their input a vector of bids $b = (b_i)_{i \in \mathcal{N}}$. The outcome rule determines a set of winners W and the payment rule assigns a payment to each bidder.

Recall our design goals:




Figure 1: Single-minded CA instance from Example 12.3. The items are shown as black rectangles and the bundles as color-coded ellipses.

1. The mechanism should be *dominant strategy incentive compatible* (DSIC).
2. At equilibrium, the mechanism should obtain high *social welfare*.
3. The mechanism should run in *polynomial time*.

2 Hardness

A first observation is that we cannot hope to get an exact solution because the allocation problem is NP-hard. Note that this is purely an optimization question, the incentives do not matter at all.

Theorem 12.4 (Lehmann, O’Callaghan, Shoham 1999). *The allocation problem among single-minded bidders is NP-hard.*

Proof sketch. We will prove the claim by reduction from independent set. Consider a graph $G = (V, E)$. Each node is represented by a bidder. Each edge is represented by an item. For bidder i , set $S_i^* = \{e \in E \mid i \in e\}$ and $v_i = 1$.

Note that a set of bidders W corresponds to an independent set if and only if their sets S_i^* are disjoint, that is, if and only if W is a feasible set of winners. This implies that there is an independent set of size y if and only if there is an allocation (i.e. a set of winners W) such that $\sum_{i \in W} v_i = y$. \square

Due to this hardness result, we will consider approximation algorithms. We call an algorithm an α -approximation, if for the solution x computed by the algorithm on input $(v_i)_{i \in \mathcal{N}}$, we have $\sum_{i \in \mathcal{N}} v_i(x) \geq \frac{1}{\alpha} \max_{x^*} \sum_{i \in \mathcal{N}} v_i(x^*)$.

3 Greedy Mechanism for Small Bundle Sizes

We will first consider the following algorithm. It is clearly a polynomial-time algorithm. We will show that it can be combined with payments to get a truthful mechanism and that it yields a good approximation with respect to the maximum bundle size $d = \max_{i \in \mathcal{N}} |S_i^*|$.

Greedy-by-Value

1. Re-order the bids such that $b_1 \geq b_2 \geq \dots \geq b_n$.
2. Initialize the set of winning bidders to $W := \emptyset$.
3. For $i = 1$ to n do: If $S_i^* \cap \bigcup_{j \in W} S_j^* = \emptyset$, then $W := W \cup \{i\}$.

Example 12.5. *Consider the instance from Example 12.3. The ranking computed by Greedy-by-Value is Green, Red, Blue. Green is considered first and accepted, which leads to the removal of both Red and Blue. Green’s threshold bid is 10.*

First, we will discuss the question of truthfulness. To this end, we will define threshold bids.

Definition 12.6. Let $W(b)$ denote the set of winners when the bids are b . Define the threshold bid $\tau_i(b_{-i})$ for player i against bids b_{-i} of the bidders other than i as the smallest bid such that $i \in W(b_i, b_{-i})$, that is

$$\tau_i(b_{-i}) = \inf\{b_i \mid i \in W(b_i, b_{-i})\} .$$

Theorem 12.7. Greedy-by-Value is a monotone algorithm and charging winners the respective threshold bids yields a truthful mechanism.

Proof. In the notation of single-parameter problems, we would write $f_i(b) = 1$ if $i \in W(b)$ and 0 otherwise.

To prove monotonicity, we have to observe that $i \in W(b)$ also implies $i \in W(b'_i, b_{-i})$ for all $b'_i \geq b_i$. This follows from the fact that bidder i is only moving further to the front of the sorted list of all bids. Therefore, if all items are still available when reaching this bidder when bidding b_i , the same will have to be true for b'_i .

Myerson's lemma tells us to charge payments

$$p_i(b) = b_i f_i(b) - \int_0^{b_i} f_i(t, b_{-i}) dt .$$




Figure 2: Allocation curve for monotone algorithm

As f_i will take only values 0 and 1, due to monotonicity, the allocation curve will be very simple as depicted in Figure 2. Also the integral expression tell us that if $i \notin W(b)$ then, due to monotonicity, $p_i(b) = 0$. Otherwise, if $i \in W(b)$

$$p_i(b) = b_i f_i(b) - \int_0^{b_i} f_i(t, b_{-i}) dt = b_i - \int_{\tau_i(b_{-i})}^{b_i} 1 dt = \tau_i(b_{-i}) . \quad \square$$

Finally, the approximation guarantee follows by a simple charging argument.

Theorem 12.8. Greedy-by-Value is a d -approximation.

Proof. Let W be the set of bidders selected by the algorithm and let OPT be the optimal solution. For $i \in W$, let

$$OPT_i = \{j \in OPT, j \geq i \mid S_i^* \cap S_j^* \neq \emptyset\} .$$

That is, OPT_i contains the indices of the bidders $j \geq i$ that are in OPT and get blocked if we accept i . Note that if $i \in OPT$ then $OPT_i = \{i\}$. Each $j \in OPT$ is included in at least one set OPT_i for $i \in W$ for the following reason: If $j \notin OPT_i$ for $i \in W$ with $i < j$, then all items in S_j^* are still available when reaching bidder j in the execution. So, j would be accepted by the algorithm and, hence, $j \in W$ and $j \in OPT_j$. Therefore, we can write

$$\sum_{j \in OPT} b_j \leq \sum_{i \in W} \sum_{j \in OPT_i} b_j .$$




Figure 3: Challenge instance for Greedy-by-Value

Next, we have $|OPT_i| \leq |S_i^*| \leq d$. This is due to the fact that the sets S_j^* for $j \in OPT_i$ are disjoint but each have a non-empty intersection with S_i^* . Furthermore, by the greedy ordering $b_j \leq b_i$ for $j \in OPT_i$. Therefore

$$\sum_{j \in OPT} b_j \leq \sum_{i \in W} \sum_{j \in OPT_i} b_j \leq \sum_{i \in W} d \cdot b_i = d \cdot \sum_{i \in W} b_i . \quad \square$$

That the approximation guarantee can be as bad as d can be seen from examples such as the one in Figure 3. Assume without loss of generality that m is a multiple of d . Every set of d items is wanted by a distinct “big” bidder, who has a value of $1 + \epsilon$ for it. Each of the d items this bidder is interested in is requested by a distinct “small” bidder, each of which has a value of 1. Greedy-by-Value will accept all the big bidders resulting in welfare $m/d \cdot (1 + \epsilon)$, while accepting all small bidders would have social welfare of m .

There are also hardness-of-approximation results showing that (under some assumptions) it is impossible to get approximation factors that a lot better than d in polynomial time. However, one get also get a \sqrt{m} -approximation in polynomial time, which is better for large bundle sizes.

4 Greedy Mechanism for Large Bundle Sizes

Our next algorithm avoids the trap in which our Greedy-by-Value algorithm stepped by normalizing bids with their bundle size. More specifically, it divides each bid by the square root of the bundle size.

Greedy-by-Sqrt-Value-Density

1. Re-order the bids such that $\frac{b_1}{\sqrt{|S_1^*|}} \geq \frac{b_2}{\sqrt{|S_2^*|}} \geq \dots \geq \frac{b_n}{\sqrt{|S_n^*|}}$.
2. Initialize the set of winning bidders to $W := \emptyset$.
3. For $i = 1$ to n do: If $S_i^* \cap \bigcup_{j \in W} S_j^* = \emptyset$, then $W := W \cup \{i\}$.

Example 12.9. Consider again the instance from Example 12.3. The ranking computed by Greedy-by-Sqrt-Value-Density is $10 \geq 14/\sqrt{2} \geq 8$. So Red is considered first and accepted. This leads to the removal of Green. Afterwards Blue is accepted. The threshold bid for Red is $14/\sqrt{2}$, for Blue it is zero.

Theorem 12.10 (Lehmann, O’Callaghan, Shoham 1999). Greedy-by-Sqrt-Value-Density is a \sqrt{m} -approximation. It is monotone, so charging threshold bids makes it a truthful mechanism.

Proof. That Greedy-by-Sqrt-Value-Density is monotone can be shown by essentially the same argument that showed that Greedy-by-Value is monotone. Holding a bidder and the bids of the other bidders fixed, the bidder faces a ranked list of bids. Its position in this sorted list determines whether he wins or not. A higher bid can only improve its position.

To establish an upper bound on the approximation guarantee we again write W and OPT for the set of winners selected by the algorithm and the optimal one. Again, we define

$$OPT_i = \{j \in OPT, j \geq i \mid S_i^* \cap S_j^* \neq \emptyset\} .$$




Figure 4: Challenge instance for Greedy-by-Sqrt-Value-Density

And we can write

$$\sum_{j \in OPT} b_j \leq \sum_{i \in W} \sum_{j \in OPT_i} b_j .$$

So, if we can show $\sum_{j \in OPT_i} b_j \leq \sqrt{m} \cdot b_i$, we are done.

As $b_j \leq \sqrt{|S_j^*|} \cdot b_i / \sqrt{|S_i^*|}$, for $j \in OPT_i$, we obtain

$$\sum_{j \in OPT_i} b_j \leq \frac{b_i}{\sqrt{|S_i^*|}} \cdot \sum_{j \in OPT_i} \sqrt{|S_j^*|}$$

Next we will show that $\sum_{j \in OPT_i} \sqrt{|S_j^*|} \leq \sqrt{m} \cdot \sqrt{|S_i^*|}$. To this end, we will use that the function $\sqrt{\cdot}$ is concave. Therefore, by Jensen's inequality, we have for all $y_1, \dots, y_\ell \geq 0$ that $\frac{1}{\ell} \sum_{k=1}^{\ell} \sqrt{y_k} \leq \sqrt{\frac{1}{\ell} \sum_{k=1}^{\ell} y_k}$ and therefore $\sum_{k=1}^{\ell} \sqrt{y_k} \leq \sqrt{\ell \sum_{k=1}^{\ell} y_k}$. So, we get

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

Now $|OPT_i| \leq |S_i^*|$ since every S_j^* , for $j \in OPT_i$, intersects S_i^* and these intersections are disjoint. Furthermore, $\sum_{j \in OPT_i} |S_j^*| \leq m$ since OPT_i is an allocation.

We obtain,

$$\sum_{j \in OPT_i} b_j \leq \frac{b_i}{\sqrt{|S_i^*|}} \cdot \sum_{j \in OPT_i} \sqrt{|S_j^*|} \leq \frac{b_i}{\sqrt{|S_i^*|}} \cdot \sqrt{|OPT_i|} \cdot \sqrt{\sum_{j \in OPT_i} |S_j^*|} \leq b_i \sqrt{m} .$$

□

We can also obtain a lower bound of \sqrt{m} on the approximation guarantee. Namely, we consider instances such as the one given in Figure 4. There is one “big” bidder with a bundle size of m and a value of $\sqrt{m} + \epsilon$ and m bidders, one for each item, with a bundle size and a value of 1. Greedy-by-Sqrt-Value-Density accepts the big bidder for a social welfare of $\sqrt{m} + \epsilon$, while accepting all small bidders would have led to a social welfare of m .

There is also a hardness-of-approximation result for this setting. Namely, our reduction to prove Theorem 12.4 in combination with hardness-of-approximation results for independent set show that one cannot hope to get approximation factors a lot better than \sqrt{m} in polynomial time.

5 Conclusion

As we see, building truthful mechanisms with good approximation guarantees is not necessarily difficult. Indeed, in the cases that we considered, even ignoring truthfulness one could not do better. This is certainly not true for every setting. For others, there is a separation between truthful and non-truthful approximation.

References

- D. Lehmann, L. I. O'Callaghan, Y. Shoham. Truthful Revelation in Approximately Efficient Combinatorial Auctions. STOC 1999 and JACM 2002. (Greedy mechanism for single-minded CAs)
- J. Håstad. Clique is hard to approximate withing $n^{1-\epsilon}$. Acta Math., 182:105–142, 1999. (Hardness of approximation for single-minded CAs when the parameter is number of items)
- E. Hazan, S. Safra, O. Schwartz. On the complexity of approximating k-set packing. Computational Complexity, 15(1):20–39, 2006. (Hardness for single minded CAs with bounded bundle size)