Incentives vs Computation

ETH Zürich | Paul Dütting

Myerson's Lemma gives a very precise and constructive answer as to which outcome rules can be implemented in a truthful mechanism. If we are after a truthful mechanism, then we have to confine ourselves to monotone outcome rules.

We have already seen an example of an objective, minimizing makespan in scheduling on related machines, which we cannot attain exactly. In these cases we need to resort to approximation if we want our mechanism to run in polynomial time.

So while we can use any approximation algorithm if we do not care about incentives, we have to devise a monotone outcome rule if we want to be able to turn it into a truthful mechanism. Does the monotonicity requirement limit our ability to achieve near-optimal outcomes in polynomial time?

1 Combinatorial Auctions

We will study the tradeoff between incentives and computation through one of the canonical problems in mechanism design.

Definition 8.1 (Combinatorial Auction). *In a* combinatorial auction *a set of m items M shall be allocated to a set of n bidders $\mathcal{N}$. The bidders have private values for bundles of items. The goal is to maximize social welfare.*

- *Feasible allocations:* $A = \{(S_1, \ldots, S_n) \subseteq M^n \mid S_i \cap S_j = \emptyset, i \neq j\}$
- *Valuation functions:* $v_i : 2^M \to \mathbb{R}_{\geq 0}$, $i \in \mathcal{N}$ *(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 8.2 (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 his bundle a winner, and we say that this bidder wins the bundle.

We will further assume that the bundle S_i^* that bidder i is interested in is *public* and only the valuation v_i^* is *private*. This turns the problem into a single parameter problem, to which our previous results apply.

Example 8.3 (Single-Minded CA). *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.*

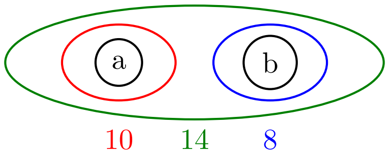

Figure 1: Single-minded CA instance from Example 8.3. The items are shown as black circles and the bundles as color-coded ellipses.

2 Hardness and Hardness of Approximation

A first observation is that the VCG mechanism, which maximizes social welfare and charges each bidder its externality, is not a viable solution. It is based on an allocation rule, which solves a NP-hard problem. We have already mentioned this result in the introductory lecture; so we only provide a proof sketch here.

Theorem 8.4 (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 a good.
- For bidder i, set $S_i^* = \{e \in E \mid i \in e\}$ and $v_i^* = 1$.

This way, winning bidders correspond to nodes in an independent set. □

The same reduction actually implies a hardness of approximation result in terms of the number of items m. A more recent results shows a lower bound in terms of the maximum bundle size of any bidder, $d = \max_i |S_i^*|$.

Theorem 8.5 (Lehmann, O'Callaghan, Shoham 1999; Håstad 1999)**.** *There is no polynomial-time algorithm for approximating the optimal allocation among single-minded bidders to within a factor of $m^{1/2-\epsilon}$, for any $\epsilon > 0$, unless* NP = ZPP*.*

Theorem 8.6 (Hazan et al. 2006)**.** *Approximating the optimal allocation among single-minded bidders to within a factor of $\Omega\left(\frac{d}{\log d}\right)$, is* NP*-hard.*

The class ZPP, for zero-error probabilistic polynomial time, is the subclass of NP consisting of those sets L for which there is some constant c and a probabilistic Turing machine $\mathcal{M}$ that on input x runs in expected time $O(|x|^c)$ and outputs 1 if and only if $x \in L$. More important for our purposes than the precise definition of the complexity class ZPP, is the fact that a conditional hardness result based on the assumption that ZPP $\neq$ NP is considered strong evidence of computational intractability.

3 Greedy Mechanisms for Single-Minded CAs

A natural question in light of the hardness results is whether we can find polynomial-time algorithms that match the lower bounds. In particular, is there a separation between the best algorithm subject to polynomial-time and the best monotone algorithm?

The answer to this question due to Lehmann, O'Callaghan, and Shoham is one of the foundational results of the field Algorithmic Game Theory: With respect to both parameters, the total number of items and the maximum bundle size, simple monotone greedy algorithms yield optimal approximation results.

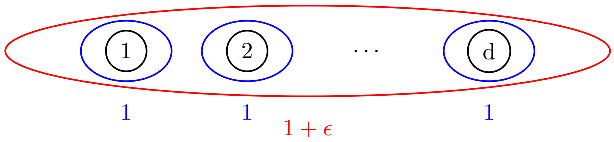

Figure 2: Challenge instance for Greedy-by-Value

Both algorithms use a carefully designed scoring function to rank the bidders. They then go through the bidders and greedily accept the next bidder in the ranked list, removing all future bidders that conflict with it.

Before we discuss the two algorithms let us first recall what truthful payments in a monotone algorithm for a setting like ours should look like.

Definition 8.7 (Threshold Payments)**.** *For an allocation rule for the single-minded CA problem denote by $W(b)$ the set of winners when the bids are b. If the allocation rule is monotone we define the* threshold bid b_i^* *for player i against bids b_{-i} of the bidders other than i as the smallest bid such that $i \in W(b_i^*, b_{-i})$.*

We first consider the algorithm that 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 $v_1^* \geq v_2^* \geq \cdots \geq v_n^*$.
2. Initialze 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 8.8. *Consider the instance from Example 8.3. 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*.*

Proposition 8.9 (Folklore)**.** *Greedy-by-Value is a $\Theta(d)$ approximation. It is monotone, so charging threshold bids yields a truthful mechansm.*

Proof. It is not difficult to see that the Greedy-by-Value algorithm is monotone. For every bidder i fixing the bids v_{-i}^* of the bidders other than i, player i's ouctome is determined by the position in the sorted list of bids of the other players. By increasing his bid v_i^*, bidder i can only move further to the front of the sorted list of all bids.

The approximation guarantee follows by a simple charging argument. Every bidder $i \in W(v)$ can block at most d bidders $j \in OPT(v)$, because both $W(v)$ and $OPT(v)$ are feasible allocations. Since we are ranking by non-increasing value each such bidder i must have a value v_i that is at least as high as the value v_j of the bidders $j \in OPT(v)$ that it blocks. We thus have the following guarantee

$$d \cdot \sum_{i \in W(v)} v_i^* \geq \sum_{i \in OPT(v)} v_i^* \Leftrightarrow \sum_{i \in W(v)} v_i^* \geq \frac{1}{d} \cdot \sum_{i \in OPT(v)} v_i^* \ .$$

That the approximation guarantee can be as bad as d can be seen from examples such as the one in Figure 2. Assume w.lo.g. 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 bid 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. □

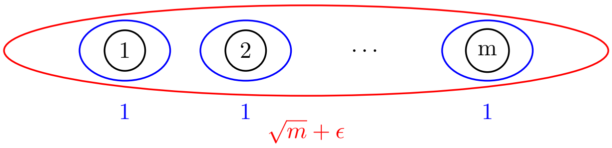

Figure 3: Challenge instance for Greedy-by-Value-Density

The same example that we used to establish a lower bound of d for Greedy-by-Value, also shows a lower bound of m. This is considerably worse than our lower bound of $\sqrt{m}$ on what can be achieved with a polynomial-time algorithm.

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-Value-Density

1. Re-order the bids such that $\frac{v_1^*}{\sqrt{|S_1^*|}} \geq \frac{v_2^*}{\sqrt{|S_2^*|}} \geq \cdots \geq \frac{v_n^*}{\sqrt{|S_n^*|}}$.
2. Initialze 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 8.10. *Consider again the instance from Example 8.3. The ranking computed by Greedy-by-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 8.11 (Lehmann, O'Callaghan, Shoham 1999)**.** *Greedy-by-Value-Density is a* $\Theta(\sqrt{m})$ *approximation. It is monotone, so charging threshold bids make it a truthful mechanism.*

Proof. That Greedy-by-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 proceed as follows. In the remainder we will simply write W for $W(v^*)$ and OPT for $OPT(v^*)$. For $i \in W$, let

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

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

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

Next we will show that $\sum_{j \in OPT_i} \sqrt{|S_j^*|} \leq \sqrt{m} \cdot \sqrt{|S_i^*|}$. By the Cauchy-Schwarz inequality,

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

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,

$$\sum_{j \in OPT} v_j^* \leq \sum_{i \in W} \sum_{j \in OPT_i} v_j^* \leq \sqrt{m} \cdot \sum_{i \in W} v_i^* \ .$$

To obtain a lower bound of $\sqrt{m}$ on the approximation guarantee we consider instances such as the one given in Figure 3. 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-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. □

We conclude that with respect to both quality measures, number of items m and maximum bundle size $d = \max_i |S_i^*|$, insisting on monotonicity did not lower our ability to obtain a near optimal outcome.

4 Case Study: Spectrum Auctions

Let us return to our initial motivation of studying combinatorial auctions, the upcoming FCC Incentive Auction (March 29, 2016). We will follow the literature, and make the simplifying assumption that bidders are single minded.

We have argued that the VCG mechanism is flawed as it requires exact optimization, which is computationally intractable. We have then reviewed two simple truthful mechanisms, which achieve the best possible performance subject to computability. It turns out that all three mechanisms are very vulnerable to collusion.

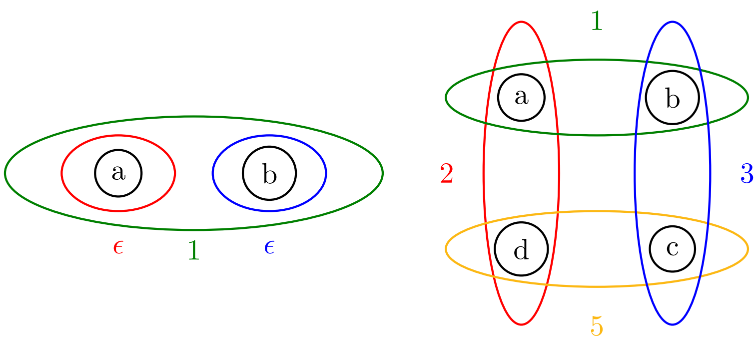

(a) VCG mechanism (a) Greedy mechanisms

Figure 4: Incentive Issues of VCG and Greedy

Example 8.12 (Incentive Issues of VCG)**.** *Consider the example in Figure 4 (a). There are two items* $\{a, b\}$ *and three bidders Red, Green, Blue. Red has a value of* ϵ *for* $\{a\}$*, Green has a value of* 1 *for* $\{a, b\}$*, and Blue has a value of* ϵ *for* $\{b\}$*.*

If all bidders bid truthfully, then the VCG mechanism gives both items to Green at a price of 2ϵ*; this is the welfare-maximizing outcome. If Red and Blue collude and they both report the same false valuation which is greater than 1, however, the VCG mechanism allocates each of these bidders the item that it wants, and for free!*

Example 8.13 (Incentive Issues of Greedy)**.** *The example in Figure 4 (b) shows an instance with four items* $\{a, b, c, d\}$ *and four bidders Red, Green, Blue, Yellow with values* $2, 1, 3, 5$*. If the bidders reported their valuations truthfully, both greedy algorithms would accept bids* 5 *and* 1*, and the corresponding bidders would be charged* 3 *and* 0*.*

If the two losing bidders were to instead both outbid the highest bid, they would both win, and pay nothing! In general, the impact of group deviations on social welfare can be arbitrarily large.

Definition 8.14. *A mechanism $M = (f, p)$ for a single-parameter setting is* (weakly) group-strategproof *(WGSP) if no group of bidders can jointly misreport their preferences, such that this misreport leads to a strictly higher utility for all members of the group.*

Clearly, any mechanism that is WGSP is DSIC; so WGSP is a strengthening of the DSIC requirement. We have seen in Example 8.12 and 8.13 that both, the VCG mechanism and the Greedy mechnisms, are *not* WGSP.

The proposal for the upcoming FCC auction foresees that a mechanism from the class of *deferred-acceptance auctions* (defined below) will be used. These can be thought of as (adaptive) reverse greedy algorithms. These mechanisms have several nice properties, amongst others these mechanisms are WGSP.

Definition 8.15 (Deferred-Acceptance Auction)**.** *A* deferred-acceptance auction *operates in stages $t \geq 1$. In each stage t a set of bidders $A_t \subseteq N$ is active; initially, $A_1 = \mathcal{N}$. The DA auction is fully defined by a collection of deterministic scoring functions $\sigma_i^{A_t}(b_i, b_{N \setminus A_t})$ that are non-decreasing in their first argument. Stage t proceeds as follows:*

- *If A_t is feasible, accept the bidders in A_t and charge each bidder $i \in A_t$ its threshold price $p_i(b_i) = \inf\{b'_i \mid i \in A(b'_i, b_{-i})\}$, where $A(b'_i, b_{-i})$ denotes the set of bidders that would have been accepted if the reported bids were (b'_i, b_{-i}) instead of (b_i, b_{-i}).*
- *Otherwise, set $A_{t+1} = A_t \setminus \{i\}$ where bidder $i \in \arg\min_{i \in A_t}\{\sigma_i^{A_t}(b_i, b_{\mathcal{N} \setminus A_t})\}$ is an active bidder with the lowest score.*

Theorem 8.16 (Milgrom and Segal 2014)**.** *Deferred-Acceptance Auctions are weakly group-strategyproof.*

The ranking functions used in the (forward) greedy algorithms of LOS satisfy the conditions on a scoring function. The greedy by value rule trivially satisfies it, the greedy by square root of the bundle size is a valid scoring function as the bundle sizes are part of the public information about the set of active bidders. However, when deployed in a DA auction, they can lead to arbitrarily bad social welfare.

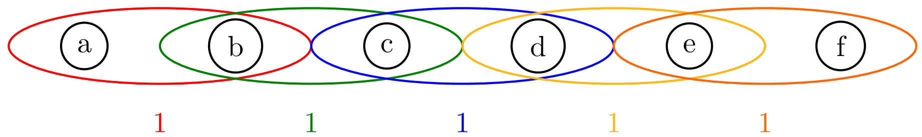

Figure 5: Counterexample for Greedy-by-Value and Greedy-by-Value-Density

Example 8.17 (Counterexample for LOS Algorithms)**.** *Consider the example in Figure 5. By adapting the weights slightly, we can ensure that the Greedy-by-Value and Greedy-by-Value-Density mechanisms, which are identical in this scenario, reject bundles from left to right. So the only bid that will be accepted is the orange bid, which a value of 1. It would have been optimal to accept the red, blue and orange bid for a social welfare of 3. Extending this example, we obtain a lower bound of $\Omega(m/2)$.*

These negative examples raise the question which performance can be achieved with a DA auction? It turns out that we can be as good as with a truthful mechanism.

Theorem 8.18 (Dütting, Gkatzelis, Roughgarden 2014)**.** *There is a deferred-acceptance auction for single-minded CAs, which guarantees an $O(d)$-approximation of the optimal social welfare.*

It is also possible to get an (almost tight) approximation with respect to the total number of items of $O(\sqrt{m \log m})$. This shows that insisting on stronger incentives, does not lead to worse performance. Quite contrarily, by considering a reverse greedy algorithm we can get the best of both worlds!

Recommended Literature

- 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)
- P. Milgrom and I. Segal. Deferred-Acceptance Auctions and Spectrum Re-Allocation. ACM EC 2014. (Definition of deferred-acceptance auctions and incentive properties)
- P. Dütting, V. Gkatzelis, T. Roughgarden. The Performance of Deferred-Acceptance Auctions. ACM EC 2014. (Deferred-acceptance auctions for single-minded CAs and Knapsack auctions)
- P. Dütting, I. Talgam-Cohen, T. Roughgarden. Modularity and Greed in Double Auctions. ACM EC 2014. (Extension of deferred-acceptance framework to double auctions)