Deferred-Acceptance Mechanisms

ETH Zürich Paolo Penna

In the previous lecture we have seen *obviously strategyproof* mechanisms. In this lecture we present a general construction of such mechanisms, called *deffered-acceptance* mechanisms. The construction resembles the cost-sharing mechanism, though in some details it is slightly more involved. We shall apply this construction to combinatorial auctions with single-minded bidders we already seen in a previous lecture.

1 Warm Up

Recall the main definition of the cost-sharing mechanism of the previous lecture, where we iteratively drop players:

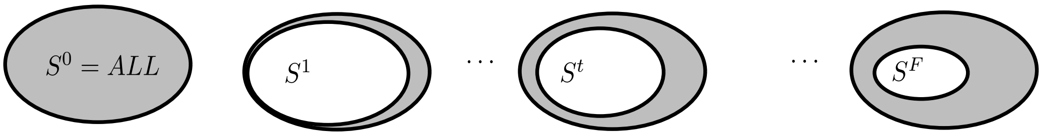

We want to apply this idea to **combinatorial auction** with single-minded bidders (seen in prior lectures):

- Each bidder i is interested in a (public) subset S_i of items.
- Bidder i is willing to pay some (private) amount v_i for getting the bundle S_i, and b_i is the bid (reported value) of i.

Recall that whenever two bidders i and j have intersecting bundles (they want the same item) it is not feasible to assign the desired bundle to all bidders.

We begin by considering a very simple and natural "greedy" algorithm which removes the bidder with **lowest bid** until obtaining a feasible solution. It will be convenient to write this algorithm is a slightly more general way towards the general algorithm.

Vanilla Deferred-Acceptance Mechanism

1. Start with the set S^0 of all players;
2. If S^t is not feasible then
   (a) Score bidders in S^t using the function
   $$\sigma_i(b_i) = b_i \tag{1}$$
   (b) Drop the lowest score bidder in S^t
3. Repeat Step 2 until getting a feasible set S^F (final set of winners).

Recall the one-parameter setting and the fact that truthful mechanisms are equivalent to monotone algorithms:

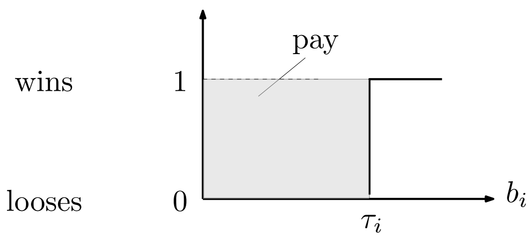

Exercise 1. *Show that the Vanilla Deferred-Acceptance Mechanism is monotone, and therefore the payments to make it into a truthful mechanism are the threshold values:*

$$\tau_i(b_{-i}) := \inf\{x \mid i \in S(x, b_{-i})\}$$

where $S(b)$ is the final set of winners in the mechanism.

Example 1. *Observe that on this instance of CA the Vanilla DA mechanism can be quite bad in terms of social welfare:*

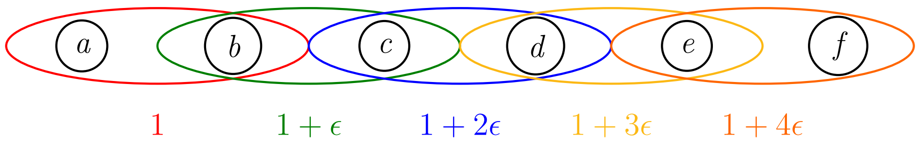

Indeed, Vanilla DA drops all bidders one by one, leaving only the last bidder.

Good: Vanilla DA is truthful (or strategyproof)
Bad: Vanilla DA has a very bad approximation

The previous example shows that the Vanilla DA is **not** the Greedy-by-Value mechanism which **adds** bidders one by one. Though both mechanisms score bidders in the same way (according to bids) they proceed in **opposite directions**:

- Greedy-by-value starts from the empty solution and adds the best feasible bidder at each step (Greedy);
- Vanilla DA starts from the unfeasible solution of all bidders and drops the worst one until a feasible solution is obtained (Reverse Greedy).

Recall that Greedy-by-Value is $O(d)$-approximate, where d is the largest bundle size. So in the instance above it is $O(1)$-approximate, while Vanilla DA is arbitrarily bad (just make the instance "longer" by extending the chain).

2 Deferred Acceptance Mechanisms

The bad example suggest to use more clever scoring function for this type of "reverse greedy" mechanisms. The resulting general scheme simply changes (2) into a more general scoring function:

Deferred-Acceptance Mechanism DA_σ

1. Start with the set $S^0 \leftarrow \mathcal{N}$ of all players;
2. If S^t is not feasible then
   (a) Score bidders in S^t using a scoring function
   $$\sigma_i^{S^t}(b_i, b_{\mathcal{N}\setminus S^t}) \tag{2}$$
   which is monotone non-decreasing in b_i.
   (b) Drop the lowest score bidder in S^t (break ties arbitrarily).
3. Repeat Step 2 until getting a feasible set S^F (final set of winners).

Note that the admissible scoring functions are all monotone functions that depend on (1) the current set of still active bidders S^t and (2) the bids of bidders who already left the auction.

We want both the following features:

- Good approximation;
- Obviously strategyproofness.

2.1 Good Approximation

We next derive a clever scoring function ("reverse greedy algorithm") which provides a good approximation when every bidder "conflicts" only with few others.

Conflict Graph:

- Each vertex is a bidder, and the **weight** of vertex i is
$$W(i) = b_i.$$
- Each edge represents a **conflict** between two bidders: edge (i, j) means that bidder i and j cannot both win ($S_i \cap S_j$ is not empty).

Given a conflict graph G, we let $c_G(i)$ be the degree of vertex i in this graph.

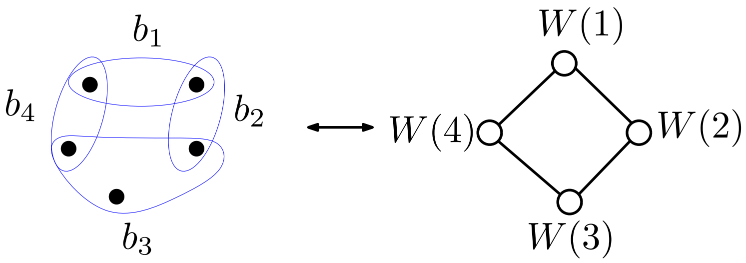

Note that we have simply reformulated our combinatorial auction problem:

Feasible Solutions $\leftrightarrow$ Independent Sets

The **social welfare** of a feasible solution is the **weight** of the independent set,

$$W(I) := \sum_{i \in I} W(i) = \sum_{i \in I} b_i \quad .$$

We shall use the following scoring function:

$$\sigma_i^{G_t}(b_i) := \frac{b_i}{c_{G_t}(i)(c_{G_t}(i)+1)} \tag{3}$$

where G_t is any subgraph (subset of vertices) of the conflict graph.[1] The resulting mechanism has a good approximation guarantee if each bidder "conflicts" with a few more. Consider the maximum degree in the conflict graph:

$$c_{\max} = \max_i c_G(i)$$

Theorem 2 (Sakai, Togasaki, Yamazaki 2003)**.** *The deferred-acceptance mechanism with the scoring function* (3) *is* $O(c_{\max})$*-approximate, where* $c_{\max}$ *is the maximum degree in the conflict graph.*

Proof. Given the sequence of graphs produced by iteratively removing one node until we obtain an independent set I:

$$G = G_0 \supset G_1 \supset \cdots \supset G_F = I \quad .$$

Claim 3. *For all* t *and for all* G_t *and* G_{t+1} *as above, it holds that*

$$\sum_{i \in G_{t+1}} \frac{W(i)}{c_{G_{t+1}}(i)+1} \geq \sum_{i \in G_t} \frac{W(i)}{c_{G_t}(i)+1} \tag{4}$$

We prove the claim below. Since the final graph G_F is an independent set, every node has degree 0, therefore

$$W(G_F) = \sum_{i \in G_F} W(i) = \sum_{i \in G_F} \frac{W(i)}{c_{G_F}(i)+1} \quad .$$

[1] We use G to denote both a subset of nodes as well as the induced subgraph.

Claim 3 implies the first of the following inequalities:

$$\sum_{i\in G_F} \frac{W(i)}{c_{G_F}(i)+1} \geq \sum_{i\in G} \frac{W(i)}{c_G(i)+1} \geq \sum_{i\in G} \frac{W(i)}{c_{\max}+1}$$

Since the optimum can at most include all nodes,

$$OPT \leq W(G) = \sum_{i\in G} W(i) \ ,$$

we obtain

$$W(G_F) \geq \frac{OPT}{c_{\max}+1}$$

which proves the theorem. □

Proof of Claim 3. Graph $G' = G_{t+1}$ is obtained from $G = G_t$ by removing a node r. Therefore, the degree of the neighbors of r decreases by 1, while all other nodes have the same degree.

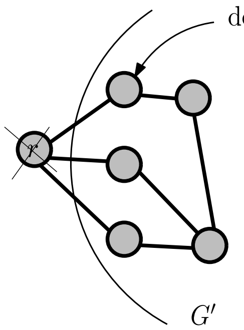

Call $N_G(r)$ the neighbors of r in G, and observe that the left summation in (4) can be otained by removing r's contribution and readjusting the contributions of $r' \in N_G(r)$:

$$\sum_{i\in G'} \frac{W(i)}{c_{G'}(i)+1} = \sum_{i\in G} \frac{W(i)}{c_G(i)+1} - \frac{W(r)}{c_G(r)+1} - \sum_{r'\in N_G(r)} \frac{W(r')}{c_G(r')+1} + \sum_{r'\in N_G(r)} \frac{W(r')}{c_{G'}(r')+1}$$

and because $c_{G'}(r') = c_G(r') - 1$

$$\begin{aligned} &= \sum_{i\in G} \frac{b_i}{c_G(i)+1} - \frac{b_r}{c_G(r)+1} + \left(\sum_{r'\in N_G(r)} \frac{b_{r'}}{c_G(r')} - \frac{b_{r'}}{c_G(r')+1} \right) \\ &= \sum_{i\in G} \frac{b_i}{c_G(i)+1} - \frac{b_r}{c_G(r)+1} + \left(\sum_{r'\in N_G(r)} \frac{b_{r'}}{c_G(r')(c_G(r')+1)} \right) \ . \end{aligned} \tag{5}$$

Since r minimizes the scoring function (3)

$$\frac{b_r}{c_G(r)(c_G(r)+1)} \leq \frac{b_{r'}}{c_G(r')(c_G(r')+1)} \ ,$$

and since $c_G(r) = |N_G(r)|$

$$\sum_{r' \in N_G(r)} \frac{b_{r'}}{c_G(r')(c_G(r')+1)} \geq \sum_{r' \in N_G(r)} \frac{b_r}{c_G(r)(c_G(r)+1)} = \frac{b_r}{c_G(r)+1} \ . \tag{6}$$

By putting together (6) and (5) the claim follows. □

2.2 Obvious Strategyproofness

Now we show that deferred-acceptance mechanisms are obviously strategyproof". Recall from the previous lecture that ascending-price auctions are obviously strategyproof.

Defferred-Acceptance Mechanism ⇒ Ascending Price Auction

That is, for any DA mechanism there is an equivalent ascending-price auction that "simulates" this mechanism. We shall only give the idea of this "simulation". Consider the following example:

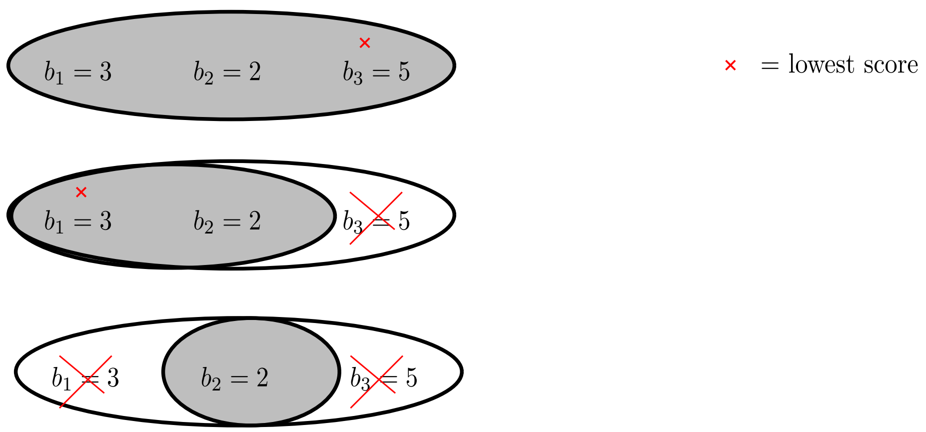

Suppose all possible valuations and bids are integers. We do the following

- Set prices equal 0,

$$p_1^1 = 0 \qquad p_2^1 = 0 \qquad p_3^1 = 0$$

- Compute the scoring function **according to these prices**

$$\sigma_1(0) \qquad \sigma_2(0) \qquad \sigma_3(0)$$

- Increase the price of the lowest score bidder (previous step) by 1. For instance, this was the second bidder, then the new prices are

$$p_1^2 = 0 \qquad p_2^2 = 1 \qquad p_3^2 = 0$$

We repeat these steps and drop players when the current price exceeds the bid. Note that in order to make the "simulation" correct, bidder 3 should be the first one to be dropped. Suppose we increase again the price of bidder 2:

$$p_1^3 = 0 \qquad p_2^3 = 2 \qquad p_3^3 = 0$$

Then it cannot happen that we further increase the price of bidder 2 before dropping bidder 3:

$$\sigma_2(2) > \sigma_3(5)$$

because the DA mechanism drops bidder 3 first. Since $\sigma_i()$ is monotone in b_i

$$\sigma_3(5) \geq \sigma_3(0)$$

and therefore for the prices p_i^3 bidder 2 is not the lowest score bidder (and thus we do not increase his/her price).

Final Remarks and Recommended Literature

The approximation algorithm in this lecture can be quite bad if the conflict graph has high degree. In the previous lectures we have seen polynomial-time strategyproof mechanism with approximation guarantee $O(d)$ and $O(\sqrt{m})$, where d is the largest bundle size and m is the number of items. These mechanisms are *not* obviously strategyproof.

Theorem 4 (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.*

This theorem says that $O(d)$-approximation is also possible with obviously strategyproof mechanisms. The same authors obtained an (almost tight) approximation with respect to the total number of items of $O(\sqrt{m \log m})$.

The best reference for the study of deferred-acceptance mechanisms with good approximation guarantee is the following one:

- P. Dütting, V. Gkatzelis, T. Roughgarden. The Performance of Deferred-Acceptance Auctions. ACM EC 2014.

  (from the application of Sakai et al. algorithm, and many other constructions)

The analysis of the reversed-greedy algorithm in this lecture is here:

- S. Sakai, M. Togasaki, K. Yamazaki. A note on greedy algorithms for maximum weighted independent set problem. Discrete Applied Mathematics 126: 2–3, 2003.

While the work introducing and proving properties of deferred-acceptance mechanisms is

- P. Milgrom and I. Segal. Deferred-Acceptance Auctions and Spectrum Re-Allocation. ACM EC 2014.

  (including the proof that DA mechanisms are obviously strategyproof)

Part of these notes is from last year's notes by Paul Dütting available here:

- http://www.cadmo.ethz.ch/education/lectures/HS15/agt_HS2015/