Topics in Algorithmic Game Theory and Economics

Pieter Kleer

Max Planck Institute for Informatics (D1)
Saarland Informatics Campus

January 13, 2020

Lecture 8
Some Mechanism Design

Mechanism Design

Mechanism design is a form of reversed game theory:

Given a (desired) outcome, how should we design the game to obtain that outcome as a result of strategic behaviour?

Examples:

- Auctions
  - Sponsored search auctions (e.g., Google)
  - Online selling platforms (e.g., eBay)
- (Stable) matching problems
  - Matching children to schools
  - Matching medical students to hospitals
- Kidney exchange markets

We focus mostly on (online) auctions.

Selling one item

Selling one item

Bidders:

- Set of bidders $\{1, \dots, n\}$ and one item.
- Bidder i has valuation v_i for the item.
  - Maximum amount she is willing to pay for it.
  - Private information: v_i not known to other players or seller.
- Bidder submits bid b_i.
  - Vector of all bids denoted by $b = (b_1, \dots, b_n)$.

Seller: *Collects (sealed) bids.*

- Gives item to some bidder (if any).
  - Allocation rule $x = x(b) = (x_1, \dots, x_n)$, with

$$x_i = \begin{cases} 1 & \text{if } i \text{ gets the item,} \\ 0 & \text{otherwise.} \end{cases}$$

- Charges price of p to bidder i^* receiving item.
  - Pricing rule $p = p(b)$.

Utility of bidder i:

$$u_i(b) = x_i(b)(v_i - p(b)) = \begin{cases} v_i - p(b) & \text{if } i \text{ gets the item,} \\ 0 & \text{otherwise.} \end{cases}$$

We have

- Bidders with valuations $v = (v_1, \ldots, v_n)$ and bids $b = (b_1, \ldots, b_n)$.
- Seller with allocation rule $x(b)$ and pricing rule $p(b)$.
- Utility of player given by $u_i(b) = x_i(b)(v_i - p(b))$.
- *Revenue of seller is p if item is sold.*

Definition

A (deterministic) mechanism (x, p) for selling an item to one of n bidders is given by an allocation rule $x : \mathbb{R}^n \to \{0, 1\}^n$ with $\sum_i x_i \leq 1$, and pricing rule $p : \mathbb{R}^n \to \mathbb{R}$.

Goal of bidder i is to maximize utility given mechanism (x, p).

- Bidders will try to bid strategically.
- How should we design auction to prevent undesirable outcomes?

First price auction

First price auction

Bidders report bids $b = (b_1, \ldots, b_n)$. Item is given to $i^* = \text{argmax}_i b_i$ and price $p = \max_i b_i$ is charged.

Example

Suppose there are three bidders

- Valuations $(v_1, v_2, v_3) = (10, 30, 25)$.
- Bids $(b_1, b_2, b_3) = (5, 22, 23)$.

Winner is bidder $i^* = 3$, with price $p = 23$. Utilities are $u = (0, 0, 2)$.

Is this a good auction format?

- Does not incentivize truthful bidding.
  - Bidders have incentive to lie (i.e., not report true valuation v_i).
- Bidder 2 values item the most, but does not get it.
  - Allocation rule does not maximize social welfare objective

$$\text{"Revenue for seller"} + \text{ "Player utilities"} = \sum_i v_i x_i(b) = v_{i^*}$$

Selling one item

Second price auction

Second price auction

Second price auction

Given bids $b = (b_1, \ldots, b_n)$:

- Item is allocated to highest bidder $i^* = \text{argmax}_i b_i$.
- Price charged is second-highest bid $p = \max_{j \neq i^*} b_j$.

- Ties are broken according to some fixed tie-breaking rule.

Example

Suppose we have three bidders.

- Valuations $(v_1, v_2, v_3) = (10, 30, 25)$.
- Bids $(b_1, b_2, b_3) = (10, 30, 22)$.

Winner is bidder $i^* = 2$ and pays $p = 22$. Utilities are $u = (0, 8, 0)$.

Second price auction has many desirable properties.

Desired properties

Bidders have incentive to be truthful: Reporting v_i is dominant strategy.

Definition (Strategyproof)

Mechanism (x, p) incentivizes truthful bidding if for every bidder i, alternative bid b_i', and bids $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, b_n)$ of other bidders, it holds that

$$u_i(b_{-i}, v_i) \geq u_i(b_{-i}, b_i'),$$

where $u_i(b) = x_i(b)(v_i - p(b))$.

Bidders have non-negative utility (when reporting truthfully).

Definition (Individually rational)

Mechanism (x, p) is individually rational if for every bidder i it holds

$$u_i(b) \geq 0$$

for every bid vector $b = (b_1, \ldots, b_{i-1}, v_i, b_{i+1}, \ldots, b_n)$.

Mechanism has good performance guarantee.

Definition (Welfare maximization)

Mechanism (x, p) is welfare maximizer if it maximizes

$$\sum_i v_i x_i(b) = \text{“Revenue for seller”} + \text{ “Player utilities”}$$

assuming that bidders are truthful.

- For now this just means we want to allocate item to a bidder with highest (true) valuation $v^* = \max_i v_i$.
- (*In online setting, we are content with approximation.)*

Definition (Computational efficiency)

Mechanism (x, p) should be implementable in polynomial time, i.e., compute allocation x and price p in polynomial time.

Proof strategyproofness (second price auction):

> Mechanism (x, p) incentivizes truthful bidding if for every i, alternative bid b'_i, and $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$, it holds that
> $$u_i(b_1, \ldots, v_i, \ldots, b_n) \geq u_i(b_1, \ldots, b'_i, \ldots, b_n).$$

Fix i and b_{-i}. Let $p(b) = s$ be second-highest bid.

- We compare v_i, b'_i and s (using case distinction).
    - Assume $v_i \neq s$ for simplicity.

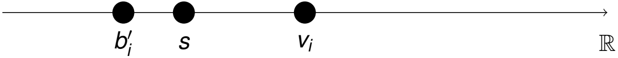

Case $s > v_i$:

- Bidder i would only win if $b'_i \geq b_{\max}$, but then $u_i = v_i - p < 0$.
- For any bid $b'_i < b_{\max}$ (then i does not get item), we have $u_i = 0$.

Case $s < v_i$:

- Bidder i wins. Charged price s same for all $b'_i > s$. For $b'_i < s$, we have $u_i = 0$. Hence, bidding v_i is an optimal choice. □

Myerson's lemma

Myerson's lemma holds in more general settings.

There exists a nice characterization, due to Myerson (1981), specifying what type of allocation rules yield strategyproof mechanisms.

- Pricing rule follows from allocation rule.

Myerson's lemma (very informal)

If there exists a monotone allocation rule x, then there is a unique pricing rule p so that the mechanism (x, p) is strategyproof (and vice versa).

Monotone allocation rule has the property that, if bidder i gets item when bidding b_i, she also gets item when bidding $b_i' \geq b_i$.

- That is, $\{0, 1\}$-variable $x_i = (b_i, b_{-i})$ is non-decreasing in bid b_i.

Exercise: Show second price auction has monotone allocation rule.

Selling multiple items

Unit-demand setting

Selling multiple items

Unit-demand setting:

- Set of items $M = \{1, \ldots, m\}$
- Set of bidders $N = \{1, \ldots, n\}$
- For every $i \in N$ a private valuation function $v_i : M \to \mathbb{R}_{\geq 0}$.
  - Value $v_{ij} = v_i(j)$ is value of bidder i for item j.
- For every $i \in N$ a bid function $b_i : M \to \mathbb{R}_{\geq 0}$.
  - Bid $b_{ij} = b_i(j)$ is maximum amount i is willing to pay for item j.

The goal is to assign (at most) one item to every bidder.

Example

Non-existing edges have $b_{ij} = 0$.

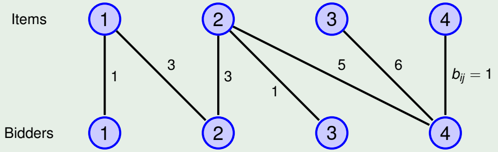

Example

Non-existing edges have $b_{ij} = v_{ij} = 0$ (i is not interested in item j)

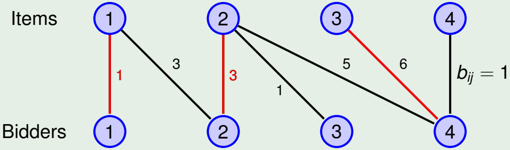

Definition (Mechanism)

A (deterministic) mechanism (x, p) is given by an allocation rule

$$x : \mathbb{R}_{\geq 0}^{n \times m} \to \{0, 1\}^{n \times m},$$

with $\sum_i x_{ij} \leq 1$ and $\sum_j x_{ij} \leq 1$, and pricing rule $p : \mathbb{R}_{\geq 0}^{n \times m} \to \mathbb{R}_{\geq 0}^{m}$.

- For bidder i, we have bid vector $b_i = (b_{i1}, \ldots, b_{im})$.
  - With $b = (b_1, \ldots, b_n)$, we have $x = x(b)$ and $p = p(b)$.
- Utility of bidder i is

$$u_i(b) = \begin{cases} v_{ij} - p_j(b) & \text{if j is the item } i \text{ receives,} \\ 0 & \text{if i does not get an item.} \end{cases}$$

Desired properties

- **Strategyproof:** For every $i \in N$, bidding true valuations $v_i = (v_{i1}, \ldots, v_{im})$ is dominant strategy.
  - It should hold that

$$u_i(b_{-i}, v_i) \geq u_i(b_{-i}, b_i').$$

  for all $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$ and other bid vector b_i'.
- **Individual rationality:** Non-negative utility when bidding truthfully.
- **Welfare maximization:** The allocation x maximizes

$$\sum_{i,j} x_{ij} v_{ij}$$

  with $x_{ij} = 1$ if bidder i gets item j, and zero otherwise.
  - Bipartite maximum weight matching in unit-demand setting.
- **Computationally tractable:** Allocation and pricing rules should be computable in polynomial time.

Vickrey-Clarke-Groves (VCG) mechanism

VCG mechanism works in more general settings than unit-demand.

Notation:

- Bipartite graph $B = (X \cup Y, E)$ with edge-weights $w : E \to \mathbb{R}_{\geq 0}$.
  - OPT(X', Y') is sum of edge weights of max. weight bipartite matching on induced subgraph $B' = (X' \cup Y', E)$ where $X' \subseteq X, Y' \subseteq Y$.

VCG mechanism

- Collect bid vectors $b_1, \ldots, b_n$ from bidders.
- Compute maximum weight bipartite matching L^* (the allocation x)
- If bidder i gets item j, i.e., $\{i, j\} \in L^*(N, M)$, then charge her

$$p_{ij}(b) = \text{OPT}(N \setminus \{i\}, M) - \text{OPT}(N \setminus \{i\}, M \setminus \{j\}),$$

  and otherwise nothing.

$\text{OPT}(N \setminus \{i\}, M) - \text{OPT}(N \setminus \{i\}, M \setminus \{j\})$ is welfare loss for other players by assigning j to i.

We use shorthand notation $a_i b_j$ for edge $\{a_i, b_j\}$.

Example

Suppose these are reported bids (non-existing edges have $b_{ij} = 0$)

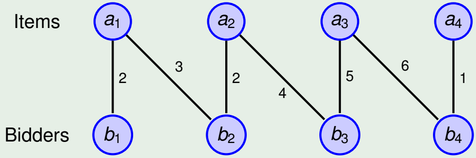

- $L^*(N, M) = \{a_1 b_2, a_2 b_3, a_3 b_4\}$.
  - $\text{OPT}(N, M) = 3 + 4 + 6 = 13$.
- $L^*(N \setminus \{b_4\}, M) = \{a_1 b_1, a_2 b_2, a_3 b_3\}$.
  - $\text{OPT}(N \setminus \{b_4\}, M) = 2 + 2 + 5 = 9$.
- $L^*(N \setminus \{b_4\}, M \setminus \{a_3\}) = \{a_1 b_2, a_2 b_3\}$.
  - $\text{OPT}(N \setminus \{b_4\}, M \setminus \{a_3\}) = 3 + 4 = 7$.
- Price charged to bidder b_4 for item a_3 is

$$p_{43}(b) = 9 - 7 = 2.$$

VCG mechanism satisfies all desired properties:

- Strategyproofness (bidding truthfully is optimal).
  - Exercise: Prove this.
- Individually rational.
  - Bidding truthfully gives non-negative utility.
- Social welfare maximizer.
  - It computes max. weight bipartite matching (where the weights are the true valuations).
- Computationally tractable.
  - Computing max. weight bipartite matching solvable in poly-time.

Online mechanism design

Selling one item

Selling one item online

Setting:

- Bidders have private valuation $v_i \geq 0$ for item.
- Whenever bidder arrives online, it submits bid b_i.

Bidders arrive one by one in unknown order $\sigma = (\sigma(1), \dots, \sigma(n))$.

Online mechanism (informal)

For $k = 1, \dots, n$, upon arrival of bidder $\sigma(k)$:

- Bid b_k is revealed.
- Decide (irrevocably) whether to allocate item to $\sigma(k)$.
  - If yes, charge price $p(b_{\sigma(1)}, \dots, b_{\sigma(k)})$ and STOP.

Goal: Allocate item to bidder with highest valuation $v^* = \max_i v_i$.

Utility of bidder i, when $\sigma(k) = i$, is given by

$$u_{i,k}(b_{\sigma(1)}, \dots, b_{\sigma(k)}) = \begin{cases} v_i - p(b_{\sigma(1)}, \dots, b_{\sigma(k)}) & \text{if } i \text{ gets item,} \\ 0 & \text{otherwise.} \end{cases}$$

Requirements for (online) deterministic mechanism (x, p):
Takes as input deterministic ordering $(y_1, \ldots, y_n)$ and bids $b_1, \ldots, b_n$ for the item.

- Specifies for every $k = 1, \ldots, n$ whether to allocate to y_k.
- This $\{0, 1\}$-variable x_k (and price p) for k is function of:
  - Total number of bidders n.
  - Bidders $y_1, \ldots, y_k$.
  - Bids $b_1, \ldots, b_k$.
  - The order $(y_1, \ldots, y_k)$.
    - Last aspect is usually irrelevant.

The variables x_i induce the allocation rule x.

Desired properties

Bidding truthfully should be dominant strategy for every arrival order σ and every arrival time of bidder i.

Definition (Strategyproof)

Consider arbitrary i, ordering $(\sigma(1), \ldots, \sigma(n))$, and k with $i = \sigma(k)$. We say an online mechanism $\mathcal{M} = (x, p)$ is strategyproof if for every alternative bid b_i' and every $b_{\sigma(1)}, \ldots, b_{\sigma(k-1)}$, it holds that

$$u_{i,k}(b_{\sigma(1)}, \ldots, b_{\sigma(k-1)}, v_i) \geq u_{i,k}(b_{\sigma(1)}, \ldots, b_{\sigma(k-1)}, b_i').$$

- **Individually rational:** Non-negative utility for bidder i when bidding truthful.
- **Constant factor α-approximation for welfare maximization**
  - For uniform random arrival model:

    $$\mathbb{E}_\sigma[v(\mathcal{M}(\sigma))] \geq \alpha \cdot \max_i v_i$$

    - With $v(\mathcal{M}(\sigma))$ valuation of bidder that gets item.
- **Computationally tractable:** Decision on who to allocate item to, and computation of charged price, in poly-time.