Introduction to Mechanism Design

Piotr Sankowski
Warsaw University
Algorithmic Trends 22.04.2014

Outline

- Introduction
- Social Choice
  - Condorcet's Paradox
  - Arrow's Theorem
- Mechanisms with Money
- Incentive Compatible Mechanism's
- Vickrey-Clarke-Groves Auction
  - Clarke's Payment Rules
  - Examples
- Combinatorial Auctions
  - Approximate Single-Minded Case

Introduction

Today lecture topic is designing games in such a way that players are forced to play according to the rules.

If players play according to the rules, then one can guarantee that the outcome of the game has some additional nice properties.

Introduction

This lecture is based on the book:

Algorithmic Game Theory

Edited by: Noam Nisan
Tim Roughgarden
Éva Tardos
Vijay V. Vazirani

Introduction

We are going to introduce the basic concepts of mechanism design.

Throughout the lecture we are going to talk about *social choice* problems.

In these problems we are asked to gather the preferences of individuals, to combine them, and to propose a single solution.

In *mechanism design* we assume that the players are strategic, i.e.,

- we assume that each player acts *rationally* and maximizes his own utility.

Introduction

The social choice abstraction can be perceived as an scenario for the following examples:

election every participant has his own preferences, the outcome is the joint choice;

markets everyone has preferences what and where to buy, the outcome is a allocation of goods and money;

auctions a market with single seller, who sets the rules of selecting the winners;

politics every citizen has his own opinion of what the government should do, and the government needs to make a decision that influences everyone.

Introduction

Observe that rationality of agents can greatly influence the intended outcome.

In routing games it causes the solutions to be very inefficient.

Many market interactions are nowadays implemented with usage of internet, e.g., Goodle Ads.

Social Choice

Consider elections with two candidates, where every elector prefers one of them.

When we want to group these preferences it seems that intuitively the majority vote is a good solution.

What shall we do when there are three candidates?

In 1785 Marquis de Condorcet observed that using majority voting for three candidates can be problematic.

Social Choice

Consider three candidates a,b and c, and three electors with the following preferences:

(1) $a \succ_1 b \succ_1 c$

(2) $b \succ_2 c \succ_2 a$

(3) $c \succ_3 a \succ_3 b$

In this case, the majority prefers $a \succ b \succ c \succ a$, i.e., it is cyclic.

In particular, who ever we elect the majority will want to change him.

Social Choice

There are many *election rules* — methods to chose one of many alternatives.

One should consider strategic players as well.

Assume that one elector has preferences $a \succ_i b \succ_i c$, but knows that a will not win, because he is hated by everybody.

Such elector can strategically vote on b instead of a, so that b is elected instead of c.

Social Choice

In strategic scenario it is hard to run the election, because the electors do not reveal their true prefereces.

This brings us to the question whether one can design reasonable elections that would be waterproofed against such manipulations?

The answer is negative and was given by Arrow in 1950.

Arrow's Theorem

We will consider a set A of possible alternatives and a set I of n players.

Let L be the set of linear orders on A.

In particular, every $\prec \in L$ is antisymmetric and transitive..

The preferences of every elector are given by $\prec_i \in L$, where $a \succ_i b$ means that i prefers a over b.

Arrow's Theorem

Definition 1

- *The function* $F : L^n \rightarrow L$ *is called social preference function.*
- *The function* $f : L^n \rightarrow A$ *is called social choice function.*

In other words, the social preference function aggregates all preferences into one joint preference.

The social choice function based on preferences chooses one of the alternatives.

Arrow's Theorem

Definition 2 *The social preference function* F *is* unanimous *if for every* $\prec \in L$ *we have* $F(\prec, \ldots, \prec) = \prec$.

In other words, if all electors have equal preferences then the social preference is the same.

Arrow's Theorem

Definition 3 *The elector i is called a* dictator *for the function F when for all $\prec_1, \ldots, \prec_n \in L$ we have $F(\prec_1, \ldots, \prec_n) = \prec_i$.*

The social preference is equal to the preferences of the dictator.

Definition 4 *The function F is* a dictatorship *if there exist a dictators for F.*

Arrow's Theorem

Definition 5 *The social preference function is* independent of irrelevant alternatives *when for every* $a, b \in A$ *and for every* $\prec_1, \ldots, \prec_n, \prec'_1, \ldots, \prec'_n \in L$, $a \prec_i b \Leftrightarrow a \prec'_i b$ *for every* i *implies that* $a \prec b \Leftrightarrow a \prec' b$, *where* $\prec = F(\prec_1, \ldots, \prec_n)$ *and* $\prec' = F(\prec'_1, \ldots, \prec'_n)$.

In other words, social preference between a and b depends only on preferences of electors between a and b, and not on preferences on c.

Twierdzenie Arrow'a

Theorem 6 (Arrow) *Every social preference function over a set of more than two candidates, which is unanimous and independent of irreverent alternatives is a dictatorship.*

Gibbard-Satterthwaite Theorem

Definition 7 *The social choice function* f *can be strategically manipulated by player* i *when for* $\prec_1, \ldots, \prec_n \in L$ *and* $\prec'_i \in L$ *such that* $a \prec_i a'$ *we have* $a = f(\prec_1, \ldots, \prec_i, \ldots, \prec_n)$*, but* $a' = f(\prec_1, \ldots, \prec'_i, \ldots, \prec_n)$*.*

In other words, when i prefers a' over a and can misreport his preferences in such a way that a' is elected instead of a.

Definition 8 *A function is* incentive compatible *when it cannot be strategically manipulated.*

Gibbard-Satterthwaite Theorem

Definition 9 *The elector i is a dictator for the social choice function f when for every $\prec_1, \ldots, \prec_n \in L$, for every $a \neq b$, we have $a \succ_i b \Rightarrow f(\prec_1, \ldots, \prec_n) = a$.*

Definition 10 *The function f is a* dictatorship *if it has a dictator.*

Theorem 11 (Gibbard-Satterthwaite) *Let f be a incentive compatible function on A, where $|A| \geq 3$, then f is a dictatorship, or some alternative is never choosen.*

Mechanisms with Money

In the social choice/preference function we were modeling preferences of the electors using linear orders.

We were not modeling who much they prefer different outcomes.

The money can be used to quantitatively measure such preferences.

Moreover, the money can be transferred between the players.

This allows to overcome the impossibility results given in these two theorems.

Mechanisms with Money

Consider a set of alternatives A and a set I of n players.

The preferences of player i are given by a *valuation function* $v_i : A \to \mathcal{R}$, where $v_i(a)$ denotes the value given by i to the outcome a.

This value is given in some currency, and when i gets m units of this currency, then his *utility* is equal to $u_i = v_i(a) + m$.

The utility is the abstraction of want the player expects and wants to maximize.

Incentive Compatibility

In the world with money, the mechanism not only should choose the single alternative, but should set the payments for all player.

The preference of player i is denoted by $v_i \in V_i$, where $V_i \subseteq \mathcal{R}^A$ is the publicly known set of alternatives for player i.

We will use the following standard notation

$$v_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$$

$$V_{-i} = V_1 \times \cdots \times V_{i-1} \times V_{i+1} \times \cdots \times V_n.$$

Incentive Compatibility

Definition 12 *The* mechanism *is a social choice function* $f : V_1 \times \cdot \times V_n \to A$ *together with payment rules* $p_1, \ldots, p_n$*, where* $p_i : V_1 \times \cdot \times V_n \to \mathcal{R}$*.*

Definition 13 *We say that the mechanism* $(f, p_1, \ldots, p_n)$ *is* incentive compatible *if for every player i, for every* $v_1 \in V_1, \ldots, v_n \in V_n$*, for every* $v'_i \in V_i$*, when* $a = f(v_i, v_{-i})$ *and* $a' = f(v'_i, v_{-i})$ *we have*

$$v_i(a) - p_i(v_i, v_{-i}) \geq v_i(a') - p_i(v'_i, v_{-i}).$$

Alternatively *truthful*, or *strategyproof*.

VCG Mechanism

The social welfare of alternative $a \in A$ is defined to be the sum of values given by all players to a, i.e., $U(a) = \sum_i v_i(a)$.

Definition 14 *The mechanism* $(f, p_1, \ldots, p_n)$ *is called Vickrey-Clarke-Groves (VCG) mechanism when*

- $f(v_1, \ldots, v_n) \in \text{argmax}_{a \in A} U(a)$,
- *for some functions* $h_1, \ldots, h_n$ *where* $h_i : V_{-i} \to \mathcal{R}$ *and for all* $v_1 \in V_1, \ldots, v_n \in V_n$ *we have:*

$$p_i(v_1, \ldots, v_n) = h_i(v_{-i}) - \sum_{j \neq i} v_j(f(v_1, \ldots, v_n)).$$

VCG Auction

The main idea of VCG is to pay to i'th player $-\sum_{j\neq i} v_j(f(v_1,\ldots,v_n))$, i.e., equal to how much other players value the outcome.

After adding this term to i'th player's valuation $v_i(f(v_1,\ldots,v_n))$ we are obtaining the social welfare.

In other words, this mechanism identifies the goals of all players with the maximization of social welfare.

The function h_i is independent from the bids of player i – from his point of view it is a constant.

VCG Mechanism

Theorem 15 (Vickrey-Clarke-Groves) *Every VCG mechanism is incentive compatible.*

VCG Mechanism

Let us fix i, v_{-i}, v_i and v_i'.

Let $a = f(v_i, v_{-i})$ and $a' = f(v_i', v_{-i})$.

The utility of i for declaring v_i is equal to:

$$v_i(a) + \sum_{j \neq i} v_j(a) - h_i(v_{-i}),$$

whereas for declaring v_i' is equal to:

$$v_i(a') + \sum_{j \neq i} v_j(a') - h_i(v_{-i}).$$

VCG Mechanism

$a = f(v_i, v_{-i})$ is the alternative that maximizes social welfare, so:

$$v_i(a) + \sum_{j \neq i} v_j(a) \geq v_i(a') + \sum_{j \neq i} v_j(a').$$

By subtracting $h_i(v_{-i})$ from both sides we get:

$$v_i(a) + \sum_{j \neq i} v_j(a) - h_i(v_{-i}) \geq$$

$$\geq v_i(a') + \sum_{j \neq i} v_j(a') - h_i(v_{-i}).$$

Clarke’a Payment Rules

We are left to chose the right functions h_i.

The simple solution is to take $h_i = 0$, but in such case we actually pay players a lot.

We would like to guarantee that players do not lose by participating, i.e., $u_i \geq 0$.

Moreover, we would not like to pay agents to take part in the auction.

Clarke'a Payment Rules

Definition 16 *The mechanism is called* individually rational *when it guarantees that each player has nonnegative utility:*

$$u_i = v_i(f(v_1, \ldots, v_n)) - p_i(v_1, \ldots, v_n) \geq 0.$$

Definition 17 *The mechanism make* no positive transfers *when it does not pay the players:*

$$p_i(v_1, \ldots, v_n) \geq 0.$$

Clarke'a Payment Rules

Definition 18 (Clarke'a Payment Rules) *Functions defined as:*

$$h_i(v_{-i}) = \max_{b \in A} \sum_{j \neq i} v_j(b),$$

are called Clarke's payment rules.

Theorem 19 *The VCG auction with Clarke's payment rules:*

- *does not make positive transfers,*
- *if $v_i(a) \geq 0$ for every $v_i \in V_i$ and $a \in A$, then it is individually rational.*

Clarke'a Payment Rules

Let $a = f(v_1, \ldots, v_n)$ maximize $U(a) = \sum_j v_j(a)$ and let b maximize $\sum_{j \neq i} v_j(b)$, then:

PT: $$p_i(v_1, \ldots, v_n) = \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a) \geq 0.$$

IR: $$u_i = v_i(a) + \sum_{j \neq i} v_j(a) - \sum_{j \neq i} v_j(b) \geq$$

$$\geq v_i(a) + \sum_{j \neq i} v_j(a) - v_i(b) - \sum_{j \neq i} v_j(b) =$$

$$\sum_j v_j(a) - \sum_j v_j(b) \geq 0.$$

Combinatorial Auctions

So far we have been considering very abstract auctions, where the list of alternatives is given explicitly.

In *combinatorial auctions* we are considering the case when we are giving a set of resources with some constraints.

The description of the constraints can be complex.

We should propose an efficient way to handle this complexity.

Combinatorial Auctions

We are given a set of m indivisible goods, which are being sold to n players.

Every player has preferences for each subset of items.

Definition 20 *The valuation v_i is a function which for each subset S of items gives the value $v_i(S) \in \mathcal{R}$ the player i would have for getting this set. It has to satisfy:*

- *free disposal – monotonicity, i.e., for $S \subseteq T$ we have $v_i(S) \leq v_i(T)$,*
- *be normalized, i.e., $v_i(\emptyset) = 0$.*

Combinatorial Auctions

We implicitly assume that utilities of the players are:

- *quasi-linear* in money, i.e., when player i gets set S and pays price p then his utility is $v_i(S) - p$,
- there are no *externalities*, i.e., utility of a player does not depend on what other players get.

Allocation of the items is denoted as $S_1, \ldots, S_n$ where $S_i \cap S_j = \emptyset$ for $i \neq j$.

Social welfare of the allocation is equals $\sum_i v_i(S_i)$.

Combinatorial Auctions

We assume that the valuation function v_i is private knowledge of player i – it is unknown to the seller and other players.

Our goal is to design a mechanism that will constrict an allocation that maximizes social welfare.

We want the mechanism to be incentive compatible.

Combinatorial Auctions

In such auction we often face the following difficulties:

computational complexity often determining the best allocation is NP-hard,

representation and communication the description of the valuation function requires exponential space,

strategic players how shall we analyze the behavior of the players.

Single-Minded Case

We restrict our attention to the case when players have just one goal.

This way we ignore communication problems and concentrate on complexity and strategy.

Definition 21 *The valuation function v is* single-minded *when there exists set S^* and value v^* such that*

- $v(S) = v^*$ *for* $S \supseteq S^*$*,*
- $v(S) = 0$ *for all other* S*.*

Single-minded bid is a pair (S^*, v^*).

Single-Minded Case

Let us now consider the complexity of computing the allocation that maximizes social welfare, when all players are single-minded.

This problem is given as:
INPUT: (S^*, v^*) for each player $i = 1, \ldots, n$.
OUTPUT: the set of winning bids $W \subseteq \{1, \ldots, n\}$ that maximizes social welfare $\sum_{i \in W} v_i^*$ such that for all $i \neq j \in W$ we have $S_i^* \cap S_j^* = \emptyset$.

Single-Minded Case

Lemma 22 *The problem of determining the allocation that maximizes social welfare is NPO-complete.*

Lemma 23 *The problem of determining the allocation that maximizes social welfare, even with approximation better then* $m^{1/2-\epsilon}$ *is NP-complete.*

Single-Minded Case

We will prove the existence of an incentive compatible mechanism that achieves approximation factor of $m^{1/2}$.

Moreover, in the case when $|S_i| = 2$ for each i it is possible to find the best allocation in polynomial time.

Similarly, when $S_i = \{j^i, j^i + 1, \ldots, k^i\}$ is a continous segment of items it is possible to find the optimal allocation as well.

There are more cases when it is possible to find optimum.

Single-Minded IC

Let V_{jc} be the set of possible single-minded bids for m items, and let A be the set of all possible allocation of these items to n players.

Definition 24 *The mechanism in the single minded case is composed out of:*

- *the allocation function* $f : (V_{jc})^n \to A$,
- *the payment functions* $p_i : (V_{jc})^n \to \mathcal{R}$ *for* $i = 1, \ldots, n$.

The mechanism is *efficient* if f and p_i can be computed in polynomial time.

Single-Minded IC

The main hardness is to get an incentive compatible mechanism that is efficient.

If efficiency was not an issue then we could:

- compute optimum allocation,
- set the prices according to VCG.

VCG works only when we can find optimum and no approximation is possible.

Single-Minded IC

- sort the bids $\frac{v_1^*}{\sqrt{|S_1^*|}} \geq \frac{v_2^*}{\sqrt{|S_2^*|}} \geq \ldots \geq \frac{v_n^*}{\sqrt{|S_n^*|}}$.
- for $i = 1 \ldots n$ do
  - if $S_i^* \cap (\bigcup_{j \in W} S_j^*) = \emptyset$ then $W \leftarrow W \cup \{i\}$.
- **Allocation:** W is the set of winners,
- **Payments:** for $i \in W$, $p_i = \frac{v_j^*}{\sqrt{|S_j^*|/|S_i^*|}}$, where j is the smallest index such that $S_i^* \cap S_j^* \neq \emptyset$, and for all $k < j$, $k \neq i$, $S_k^* \cap S_j^* = \emptyset$
  - if such j does not exist then $p_i = 0$.

Single-Minded IC

Lemma 25 *A single-minded mechanism, where the losers pay 0, is incentive compatible if and only if when it satisfies the following two conditions:*

monotonicity: *Every winner with bid* (S_i^*, v_i^*) *will still be a winner for every* $v_i' > v_i^*$ *and* $S_i' \subset S_i^*$ *(when other bids are fixed).*

critical payment: *Each winner pays the minimum amount that is needed for him to win, i.e, the infimum over valuations* v_i' *such that* (S_i^*, v_i') *is still a winning bid.*

Single-Minded IC

Before we prove this lemma we are going to show that the proposed mechanism satisfies these conditions:

monotonicity: By increasing v_i^* or decreasing S_i^* we move the bidder forward in the greedy order.

critical payments: Observe that i wins with j when i is before j in the greedy order. The payment rule gives exactly the moment when i is before j in greedy order.

Single-Minded IC

Proof of lemma: Observe that the player that reports truth will never have negative utility:

- when he looses his utility is 0, because losers pay 0,
- when he wins his utility is not smaller then the critical price.

Let us now show that the player cannot increase his utility by reporting (S', v') instead of his true bid $(S, v) = (S^*, v^*)$.

Single-Minded IC

When (S', v') is a loosing bid or S' does not contain S, then it is clear that it is not worse to report (S, v) .

Hence, let us assume that (S', v') is a winning bid and $S' \supseteq S$.

First, we will show that the player will be no worse when he reports (S, v') instead of (S', v').

Let us denote by p' the payment of the player for (S', v') and by p for (S, v').

Single-Minded IC

For every $x < p$, if player report (S, x) he will lose because p is critical.

By monotonicity (S', x) will be a loosing bid for every $x < p$.

This why the critical payment p' cannot be smaller then p.

Hence when the player reports (S, v') instead of (S', v') he still wins and pays no more.

Single-Minded IC

We are left to consider to show that reporting (S, v) is no worse then reporting (S, v').

Let us assume that (S, v) is a winning bid with payment $\tilde{p}$.

As long as v' is bigger then $\tilde{p}$, the player still wins and pays exactly the same amount.

If $v' < \tilde{p}$ the player loses and has zero utility.

Single-Minded IC

If (S, v) is a loosing bid, then v has to be smaller then the critical payment.

Hence, the payment for any winning bid (S, v') will be bigger then v.

Lemma 26 *Let OPT be the allocation that maximizes social welfare $\sum_{i \in OPT} v_i^*$, and let W be the outcome of the mechanism, then $\sum_{i \in OPT} v_i^* \leq \sqrt{m} \sum_{i \in OPT} v_i^*$.*

Single-Minded IC

For every $i \in W$ let $OPT_i = \{j \in OPT, j \geq i | S_i^* \cap S_j^* \neq \emptyset\}$ be the set of elements of OPT, that were not taken into W because of i (including i).

We have $OPT \subseteq \bigcup_{i \in W} OPT_i$, and we will prove that for all $i \in W$, we have $\sum_{j \in OPT_i} v_j^* \leq \sqrt{m} v_i^*$.

Observe that every $j \in OPT_i$ was after i in the greedy order and so $v_j^* \leq \frac{v_i^* \sqrt{|S_j^*|}}{\sqrt{|S_i^*|}}$.

Single-Minded IC

By summing up this inequality over all $j \in OPT_i$ we get:

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

By applying Cauchy-Schwarz inequality we obtain:

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

Single-Minded IC

Every S_j^* for $j \in OPT_i$ intersects S_i^*.

Because OPT is an allocation, these intersections need to be disjoint, and $|OPT_i| \leq |S_i^*|$.

Similarly $\sum_{j \in OPT_i} |S_j^*| \leq m$.

This way we obtain $\sum_{j \in OPT_i} \sqrt{|S_j^*|} \leq \sqrt{|S_i^*|}\sqrt{m}$.

By combining all inequalities this gives $\sum_{j \in OPT_i} v_j^* \leq \sqrt{m} v_i^*$.