Mechanism Design Without Money

Christos Tzamos

School of Electrical and Computer Engineering
National Technical University of Athens, Greece

Motivation

Mechanism Design and Social Choice

Design rules in order to make decisions based on people preferences when their interests are conflicting.

Game Theory

Game Theory

Studies strategic situations where players choose different actions in an attempt to maximize their returns.

Outcome Prediction - Solution Concepts

- Nash Equilibrium
- Pure Nash Equilibrium
- Dominant Strategy

Mechanism Design

Mechanism Design

Mechanism design is the art of designing rules of a game to achieve a specific outcome under a certain solution concept.

Social Choice as a Game

- A set A of different alternatives
- A set of n voters (the agents) N
- Each agent i has a linear order $\succ_i \in L$ over the set A

A function (mechanism) $f : L^n \to A$ that maps the agents' preferences to a single alternative is called social choice function.

Properties

Onto

$\forall a \in A, \exists x \in L^n$ such that $f(x) = a$

Unanimous

if $\exists a \in A$ such that $\forall b \in A$ and $i \in N$, $a \succ_i b$ then $f(\succ_1, \ldots, \succ_n) = a$

Pareto Optimal

if $f(\succ_1, \ldots, \succ_n) = a$, then $\nexists b \in A$ such that $b \succ_i a, \forall i \in N$

Properties - Incentive Compatibility

Strategic Manipulation by agent i

$\exists \succ_1, \ldots, \succ_n, \succ_i' \in L$ such that $b \succ_i a$ where $a = f(\succ_1, \ldots, \succ_i, \ldots, \succ_n)$ and $b = f(\succ_1, \ldots, \succ_i', \ldots, \succ_n)$.

Strategyproofness

A social choice function is called **incentive compatible** or **strategyproof** or **truthful** if no agent can strategically manipulate it.

Impossibility Result

Gibbard-Satterthwaite

Let f be an incentive compatible social choice function onto A, where $|A| \geq 3$, then f is a dictatorship.

Escape Routes

- Money

Impossibility Result

Gibbard-Satterthwaite

Let f be an incentive compatible social choice function onto A, where $|A| \geq 3$, then f is a dictatorship.

Escape Routes

- Money
- Randomization

Impossibility Result

Gibbard-Satterthwaite

Let f be an incentive compatible social choice function onto A, where $|A| \geq 3$, then f is a dictatorship.

Escape Routes

- Money
- Randomization
- Restricted domain of preferences

Single peaked preferences

Single peaked preferences

- One dimensional ordering of alternatives ($A = [0, 1]$)
- $\forall i \in N, \exists p_i$ such that $\forall x \in A - \{p_i\}$ and $\forall \lambda \in [0, 1)$, $p_i \succ_i \lambda x + (1 - \lambda) p_i$

The set of single peaked preferences is denoted by R

Properties

Onto $\equiv$ Unanimous $\equiv$ Pareto Optimal

Single peaked preferences

Generalized Median

A rule f is strategy-proof, onto and anonymous if and only if there exist $y_1, y_2, \ldots, y_{n-1} \in [0, 1]$ such that for all $\succeq \in R^n$,

$$f(\succeq) = median(p_1, p_2, \ldots, p_n, y_1, y_2, \ldots, y_{n-1})$$

Proof Idea

- Consider an agent i and fix all the other preferences. We get a function $f_i : R \to [0, 1]$
- The imageset $f_i(R)$ is closed
- The imageset $f_i(R)$ is an interval

Single peaked preferences

Generalized Median Voter Scheme

A rule f is strategy-proof and onto **iff** and only if there exist points $a_S \in [0,1], \forall S \subseteq N$ such that:

- $S \subseteq T$ implies $a_S \leq a_T$
- $a_\emptyset = 0$ and $a_N = 1$
- $\forall \succeq \in R^n$, $f(\succeq) = \max_{S \subseteq N} \min\{a_S, p_i : i \in S\}$

Setting and Outline

Facility Location Game

A number of facilities are to be placed in a metric space based on the preferences of strategic agents.

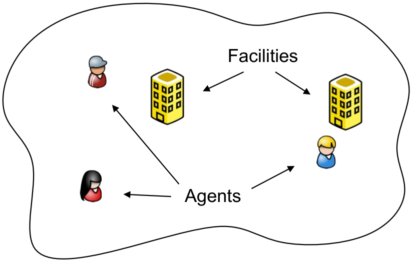

Agents

Agents

- $N = \{1, \ldots, n\}$ is the **set** of agents.

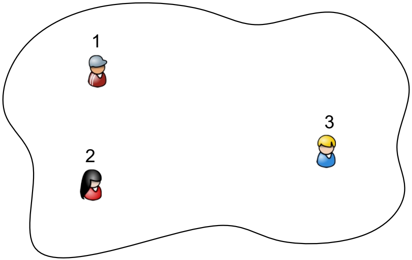

Agents

Agents

- $N = \{1, \ldots, n\}$ is the **set** of agents.
- Each agent *i* **has** a location x_i, which is *i*'s private information.

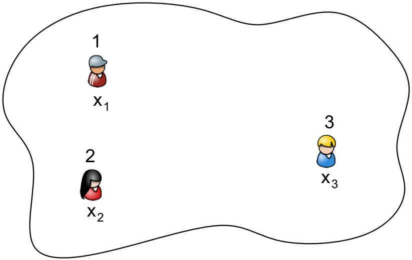

Agents

Agents

- $N = \{1, \ldots, n\}$ is the **set** of agents.
- Each agent i **has** a location x_i, which is i's private information.
- Each agent i **reports** a location y_i (y_i may not be equal to x_i).

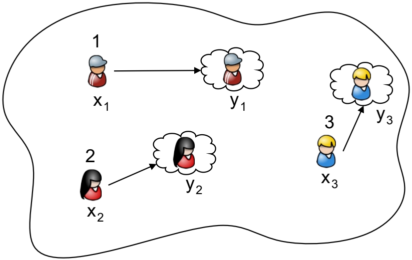

Agents

Agents

- $N = \{1, \ldots, n\}$ is the **set** of agents.
- Each agent i **has** a location x_i, which is i's private information.
- Each agent i **reports** a location y_i (y_i may not be equal to x_i).

The tuple $\vec{x} = (x_1, \ldots, x_n)$ is the **location profile**.

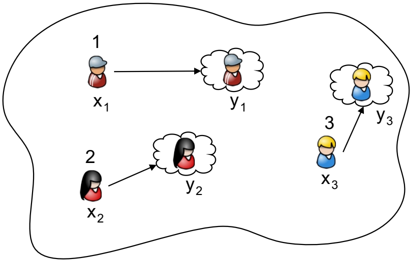

Mechanisms

Deterministic Mechanism

A function F that maps a location profile $\vec{x}$ to a non-empty set of facilities.

Randomized Mechanism

A function F that maps a location profile $\vec{x}$ to a probability distribution over non-empty sets of facilities.

Connection Cost

Connection Cost

The (expected) distance of agent i to her closest facility is called **connection cost**: $\text{cost}[x_i, F(\vec{x})] = d(x_i, F(\vec{x}))$

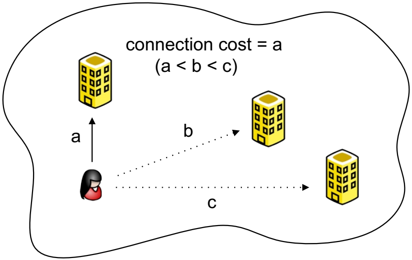

Objectives

Strategyproofness

For any location profile $\vec{x}$, any agent i, and any location y:
$\text{cost}[x_i, F(\vec{x})] \leq \text{cost}[x_i, F(y, \vec{x}_{-i})]$

Group-Strategyproofness

For any location profile $\vec{x}$, any set of agents S, and any location profile $\vec{y}_S$ for them:
$\exists i \in S : \text{cost}[x_i, F(\vec{x})] \leq \text{cost}[x_i, F(\vec{y}_S, \vec{x}_{-S})]$

Objectives II

Efficiency

Place the facilities in the metric space so as to minimize a given objective function.

Objectives II

Efficiency

Place the facilities in the metric space so as to minimize a given objective function.

Social Cost

$\sum_{i \in N} \mathrm{cost}[x_i, F(\vec{x})]$

Objectives II

Efficiency

Place the facilities in the metric space so as to minimize a given objective function.

Social Cost

$\sum_{i \in N} \mathrm{cost}[x_i, F(\vec{x})]$

Maximum Cost

$\max_{i \in N} \mathrm{cost}[x_i, F(\vec{x})]$

Single-Facility on the line - Social Cost

Single Facility on a line [Moul 80] [BarbBev 94] [Sprum 95]

The **generalized median** is the only strategyproof mechanism.
Median = Optimal Solution

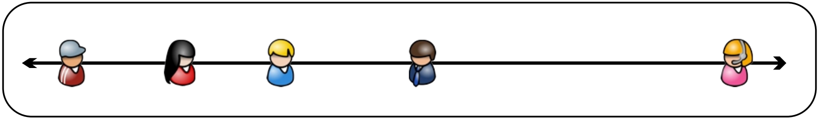

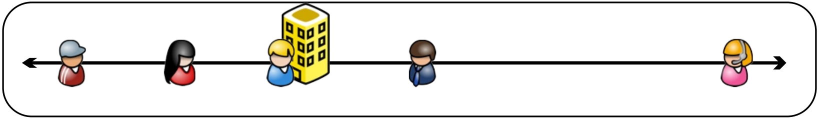

Single-Facility on the line - Social Cost

Single Facility on a line [Moul 80] [BarbBev 94] [Sprum 95]

The **generalized median** is the only strategyproof mechanism.
Median = Optimal Solution

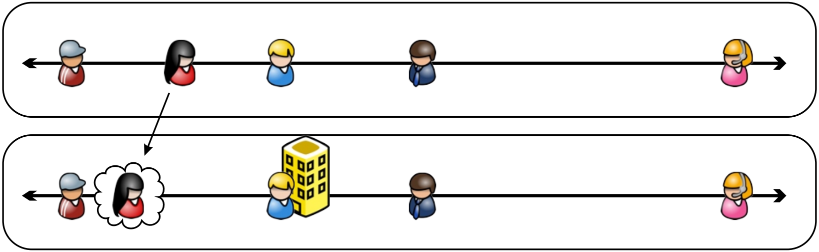

Single-Facility on the line - Social Cost

Single Facility on a line [Moul 80] [BarbBev 94] [Sprum 95]

The **generalized median** is the only strategyproof mechanism.
Median = Optimal Solution

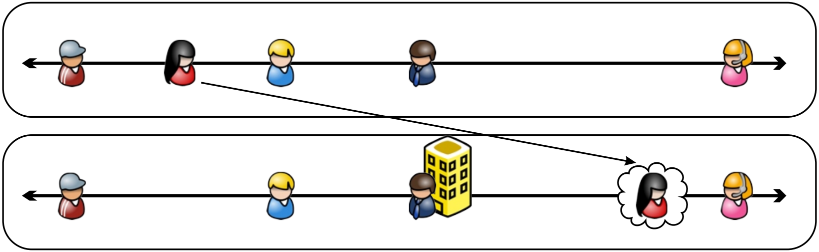

Single-Facility on the line - Social Cost

Proof of Optimality

If n is odd, $n = 2k + 1$. Any point that is to the left of the median has higher social cost since it is further away from at least $k + 1$ locations and closer to at most k locations, and the same holds for any point to the right of the median.

If n is even, $n = 2k$, and without loss of generality $x_1 \leq \cdots \leq x_n$, then any point in the interval $[x_k, x_k + 1]$ is an optimal facility location. In this case the median is considered to be the leftmost point of the optimal interval.

Single-Facility on general metric spaces - Social Cost

Single Facility on a tree [SchumVohr 02]

The **extended median** is the only strategyproof mechanism.
The optimal solution is strategyproof.

Single-Facility on general metric spaces - Social Cost

Single Facility on a tree [SchumVohr 02]

The **extended median** is the only strategyproof mechanism.
The optimal solution is strategyproof.

Impossibility Result [SchumVohr 02]

For non-tree metrics, **only dictatorial rules** can be both strategyproof and onto.
The optimal solution is **not** strategyproof.

Picking an agent deterministically $\rightarrow$ cost = $(n-1)OPT$.
Picking an agent at random $\rightarrow$ cost = $2OPT$.

Single-Facility on the line - Maximum Cost

Deterninistic Mechanisms

Upper Bound: 2
Any k-th order statistic of the reported locations.
Lower Bound: 2

Randomized Mechanisms

Upper Bound: 3/2
P(leftmost location) = 1/4
P(rightmost location) = 1/4
P(average of the leftmost and rightmost location) = 1/2
Lower Bound: 3/2

2-facility Location on the line

Deterministic Upper bound [ProcTenn 09]

The best known deterministic mechanism is **(n-2)**-approximate for the social cost by selecting the placing facilities at both the leftmost and rightmost location.

Deterministic Lower bound [LSWZ 10]

Any deterministic strategyproof mechanism has an approximation ratio of at least **(n-1)/2** for the social cost.

2-facility Location - randomized

Randomized Upper bound [LSWZ 10]

The best known randomized mechanism is **4**-approximate for the social cost.

Randomized Lower bound [LSW 09]

Any randomized strategyproof mechanism has an approximation ratio of at least **1.045** for the social cost.

Proportional Mechanism

Proportional Mechanism - Description

Facilities are placed at the locations of selected agents.

- **1st Round:** Agent i_1 is selected u.a.r.
- **2nd Round:** Agent i_2 is selected with probability $\frac{d(x_{i_2}, x_{i_1})}{\sum_{i \in N} d(x_i, x_{i_1})}$.

Proportional Mechanism

Proportional Mechanism - Description

Facilities are placed at the locations of selected agents.

- **1st Round:** Agent i_1 is selected u.a.r.
- **2nd Round:** Agent i_2 is selected with probability $\frac{d(x_{i_2}, x_{i_1})}{\sum_{i \in N} d(x_i, x_{i_1})}$.

Proportional Mechanism

4-approximate strategyproof mechanism on any metric space.
Not strategyproof for more than two facilities.

Proportional Mechanism

Example

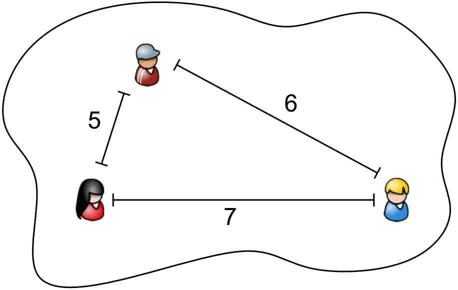

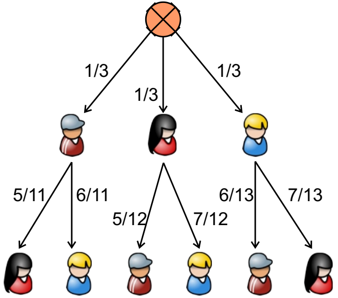

Proportional Mechanism

Example

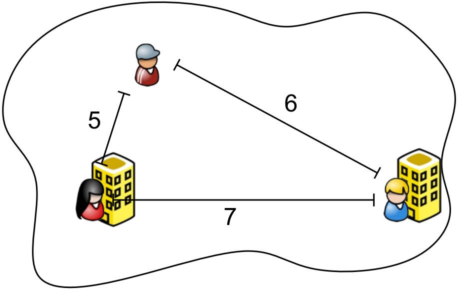

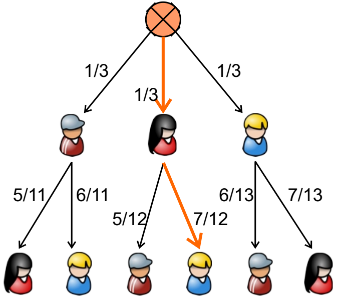

More than 2 facilities

No mechanism with bounded approximation ratio is known for the case of 3 or more facilities.
No deterministic mechanism with bounded approximation ratio is known for the case of 2 facilities in general metric spaces.

Impossibility result

Any deterministic mechanism that places facilities in the interval of the leftmost and rightmost location has unbounded approximation ratio.

Imposing Mechanisms

Imposing mechanisms may forbid an agent for connecting to certain facilities thus increasing her connection cost when she lies.

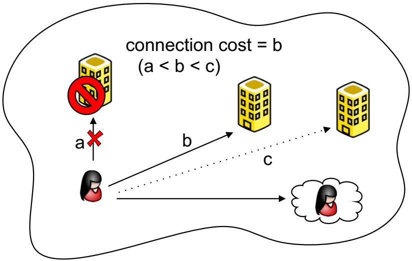

Imposing Mechanisms

Combining **differentially private** and **imposing** mechanisms [NisSmoTen 10] developed a general framework for strategyproof approximate mechanisms without money.

k-Facility Location on the [0,1] interval

Additive approximation roughly $n^{2/3}$

- Running time exponential in k.
- Does not imply constant approximation ratio.

Winner-Imposing Mechanisms

Winner-Imposing Mechanisms allow agents to connect only to the facility positioned at their reported location if any, otherwise there is no restriction.

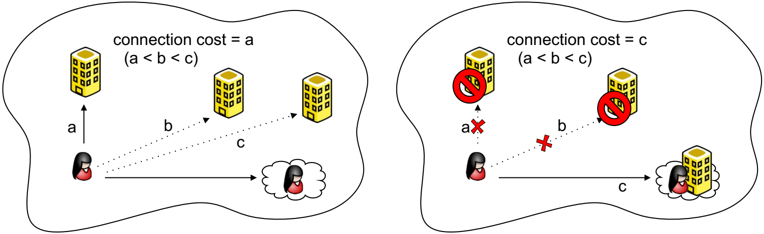

3-facility Location

Description

- Place two facilities in the leftmost agent L and the rightmost agent R. Assume L=0, R=1.
- Find $x_L = \max_{x_i \le 0.5} x_i$ and $x_R = \min_{x_i \ge 0.5} x_i$. Assume $d(x_L, 0.5) < d(x_R, 0.5)$.
- Place the third facility in $\min[\max\{x_L, 2(1 - x_R)\}, 0.5]$.

3-facility Location

Description

- Place two facilities in the leftmost agent L and the rightmost agent R. Assume L=0, R=1.
- Find $x_L = \max_{x_i \leq 0.5} x_i$ and $x_R = \min_{x_i \geq 0.5} x_i$. Assume $d(x_L, 0.5) < d(x_R, 0.5)$.
- Place the third facility in $\min[\max\{x_L, 2(1 - x_R)\}, 0.5]$.

Theorem

Strategyproof $(n-1)$-approximate mechanism.

WIProp

Winner-imposing version of the proportional mechanism of [LSWZ 10]

WIProp - Description

Facilities are placed at the locations of selected agents.

- **1st Round:** Agent i_1 is selected u.a.r.
- **ℓ-th Round:** Agent i_ℓ is selected with probability $\frac{d(x_{i_\ell}, C_{\ell-1})}{\sum_{i \in N} d(x_i, C_{\ell-1})}$.

WIProp

Winner-imposing version of the proportional mechanism of [LSWZ 10]

WIProp - Description

Facilities are placed at the locations of selected agents.

- **1st Round:** Agent i_1 is selected u.a.r.
- **ℓ-th Round:** Agent i_ℓ is selected with probability $\frac{d(x_{i_\ell}, C_{\ell-1})}{\sum_{i \in N} d(x_i, C_{\ell-1})}$.

Theorem

WIProp is a strategyproof $4k$-approximate mechanism.

General Facility Location

General Facility Location Game

Same as before but not a fixed number of facilities. Fixed cost f of opening a facility.

Objective

Minimize $\sum_{i=1}^{n} \mathrm{cost}[x_i, F(\vec{x})] + f\,|F(\vec{x})|$

Deterministic Mechanism for FL on the Line

Line Partitioning - Description

- **1st Round:** Partition the line in disjoint intervals of length 1. Every interval with more than 1 agent is considered active.
- **ℓ-th Round:** Partition every active interval of round $\ell - 1$ in half. Every interval with more than 2^ℓ agents is considered active.

Place a facility at the midpoint of every active interval.
(Active intervals of the 1st Round also get facilities at their endpoints)

Deterministic Mechanism for FL on the Line

Line Partitioning - Description

- **1st Round:** Partition the line in disjoint intervals of length 1. Every interval with more than 1 agent is considered active.
- **ℓ-th Round:** Partition every active interval of round $\ell - 1$ in half. Every interval with more than 2^ℓ agents is considered active.

Place a facility at the midpoint of every active interval.
(Active intervals of the 1st Round also get facilities at their endpoints)

Theorem

Line Partitioning is a group-strategyproof $O(\log n)$-approximate mechanism.

WI-OFL

Winner-imposing version of the online facility location algorithm of [Meyerson 01]

WI-OFL - Description

Consider any permutation of the agents.

- **1st Round:** Agent 1 is assigned a facility.
- **ℓ-th Round:** Agent ℓ is assigned a facility with probability $d(x_\ell, C_{\ell-1})$.

WI-OFL

Winner-imposing version of the online facility location algorithm of [Meyerson 01]

WI-OFL - Description

Consider any permutation of the agents.

- **1st Round:** Agent 1 is assigned a facility.
- **ℓ-th Round:** Agent ℓ is assigned a facility with probability $d(x_\ell, C_{\ell-1})$.

Theorem

WI-OFL is a strategyproof 8-approximate mechanism.

WI-OFL

Non-imposing OFL is not strategyproof even on the line.

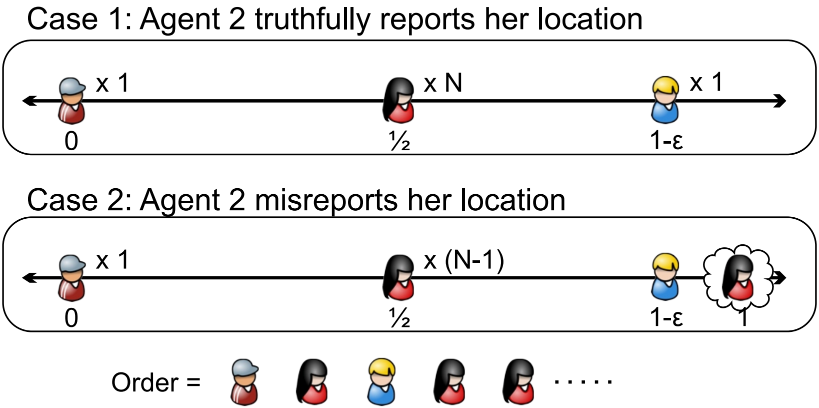

The End

Thank you!