Computationally Efficient Truthful Mechanisms

Markos Epitropou

Mechanism Design
@
NTUA

May 23, 2013

Outline

1. Dominant-Strategy Mechanisms
2. Characterizations of Truthful Mechanisms
3. Combinatorial Auctions
4. Job Scheduling

Outline

1. Dominant-Strategy Mechanisms

2. Characterizations of Truthful Mechanisms

3. Combinatorial Auctions

4. Job Scheduling

Dominant-Strategy Implementable

Truthful Mechanism

A mechanism $(f, p_1, \ldots, p_n)$ is called truthful if for every player i, every $v_1 \in V_1, \ldots, v_n \in V_n$ and every $v_i' \in V_i$, if we denote $a = f(v_i, v_{-i})$ and $b = f(v_i', v_{-i})$, then $v_i(a) - p_i(v_i, v_{-i}) \geq v_i(b) - p_i(v_i', v_{-i})$.

Dominant-Strategy Implementable

Revelation Principle

If there exists an arbitrary mechanism that implements f in dominant strategies, then there exists an incentive compatible mechanism that implements f. The payments of the players in the incentive compatible mechanism are identical to those, obtained at equilibrium, of the original mechanism.

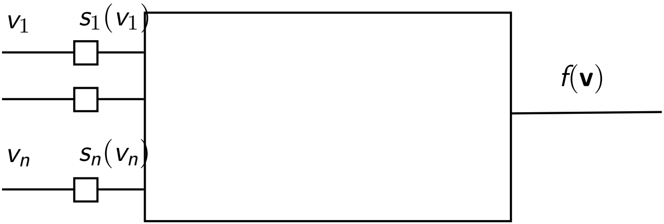

Dominant-Strategy Implementable

Revelation Principle

If there exists an arbitrary mechanism that implements f in dominant strategies, then there exists an incentive compatible mechanism that implements f. The payments of the players in the incentive compatible mechanism are identical to those, obtained at equilibrium, of the original mechanism.

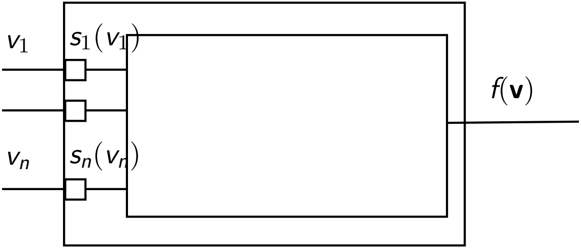

Payment Functions

The payment does not depend on v_i, but only on the alternative chosen $f(v_i, v_{-i})$.

$$\boxed{p_i = p_i(v_{-i}, a)}$$

Proof.

$$f(v_i, v_{-i}) = f(v_i', v_{-i})$$

if $p_i(v_i, v_{-i}) > p_i(v_i', v_{-i})$ then a player with type v_i declares v_i'. □

Weighted VCG

Affine Maximizers

A social choice function f is called an affine maximizer if for some subrange $A' \subset A$, for some player weights $w_1, \dots, w_n \in \mathbb{R}^+$ and for some outcome weights $c_a \in \mathbb{R}$ for every $a \in A'$, we have that $f(v_1, \dots, v_n) \in \arg\max_{a \in A'} (c_a + \sum_i w_i v_i(a))$.

VCG

Let f be an affine maximizer. Define for every i,
$p_i(v_1, \dots, v_n) = h_i(v_{-i}) - \sum_{j \neq i} \frac{w_j}{w_i} v_j(a) - \frac{c_a}{w_i}$, where h_i is an arbitrary function.

Outline

A Local Characterization

WMON

A social choice function f satisfies Weak Monotonicity (WMON) if for all i, all v_{-i} we have that $f(v_i, v_{-i}) = a \neq b = f(v_i', v_{-i})$ implies that

$$v_i'(b) - v_i(b) \geq v_i'(a) - v_i(a)$$

$IC \subseteq WMON$

If a mechanism $(f, p_1, \ldots, p_n)$ is incentive compatible, then f satisfies WMON.

A Local Characterization

$f(v_i, v_{-i}) = a$
$f(v'_i, v_{-i}) = b$

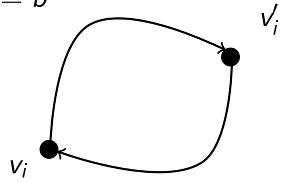

- v_i: $v_i(a) - p_a \geq v_i(b) - p_b \Rightarrow v_i(a) - v_i(b) \geq p_a - p_b$
- v'_i: $v'_i(a) - p_a \leq v'_i(b) - p_b \Rightarrow v'_i(a) - v'_i(b) \leq p_a - p_b$

A Local Characterization

Convex Domains - $WMON \subseteq IC$

If all domains of preferences V_i are convex sets then for every social choice function that satisfies WMON there exists payment functions $p_1, \ldots, p_n$ such that $(f, p_1, \ldots, p_n)$ is incentive compatible.

Setting (Job Scheduling)

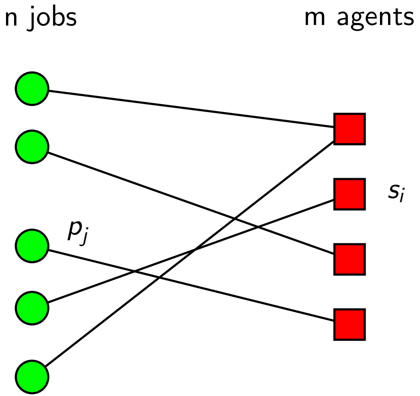

Single-Dimensional Domains

Setting (Job Scheduling)

- Agents' Types: $\{t_i = \frac{1}{s_i}\}$
- Mechanism's Input: $\{b_i\}$
- Mechanism's Output: $o(\mathbf{b})$
- Agent's Profit: $profit_i(t_i, \mathbf{b}) = p_i(\mathbf{b}) - cost_i(t_i, o(\mathbf{b}))$
- Agent's Cost: $cost_i(t_i, o) = t_i w_i(o) = \frac{l_i(o)}{s_i}$

Truthful Mechanism $(o, \mathbf{p})$

For every agent i,

$$profit_i(t_i, (\mathbf{b}_{-i}, t_i)) \geq profit_i(t_i, (\mathbf{b}_{-i}, b_i))$$

- $WMON \Rightarrow w'(b_i) \leq 0$ (decreasing load curve)
- $p_i(b_i) = h_i(b_{-i}) + b_i w_i(b_i) - \int_0^{b_i} w_i(u) du$

  $p_i(b_i) = b_i w_i(b_i) + \int_{b_i}^{\infty} w_i(u) du$

Single-Dimensional Domains

$$p_i(b_i) = h_i(b_{-i}) + b_i w_i(b_i) - \int_0^{b_i} w_i(u)du$$

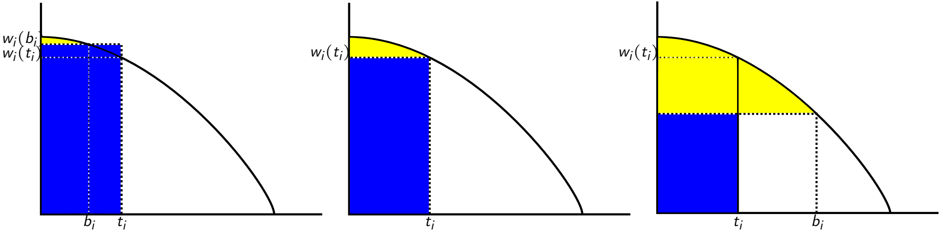

Arbitrary Valuations

Roberts Theorem

If $|A| \geq 3$, f is onto A, $V_i = \mathbb{R}^A$ for every i, and $(f, p_1, \ldots, p_n)$ is incentive compatible then f is an affine maximizer.

Outline

Notation

- n players
- m items
- Outcomes: All possible $S_1, \dots, S_n$
- $v_i(\emptyset) = 0$
- $S \subseteq T \Rightarrow v_i(S) \leq v_i(T)$

Notation

Linear Program Relaxation

$$\max \sum_{i,S\neq\emptyset} v_i(S)x_{i,S}$$

$$\text{subject to} \sum_{S\neq\emptyset} x_{i,S} \leq 1 \qquad \text{, for each player i}$$

$$\sum_i \sum_{S:j\in S} x_{i,S} \leq 1, \text{ for each item j}$$

$$x_{i,S} \geq 0 \qquad \text{, for each i,S}$$

Bidding Languages, Queries

Value Query: The auctioneer presents a bundle S, the bidder reports his value $v(S)$ for this bundle.
Demand Query: The auctioneer presents a vector of item prices $p_1, \ldots, p_m$; the bidder reports a demand bundle under these prices, i.e. some set S that maximizes $v(S) - \sum_{i \in S} p_i$.

Theorem

The Linear Program Relaxation can be solved in polynomial time using only demand queries with item prices.

Dual Linear Program

$$\min \sum_i u_i + \sum_j p_j$$

$$\text{subject to } u_i + \sum_{j \in S} p_j \geq v_i(S), \text{ for i, S}$$

$$u_i \geq 0 \quad \text{, for each player i}$$

$$p_j \geq 0 \quad \text{, for each item j}$$

Single-Minded Bidders

Valuations

A valuation v is called single minded if there exists a bundle of items S^* and a value $v^* \in \mathbb{R}^+$ such that $v(S) = v^*$ for all $S \supseteq S^*$, and $v(S) = 0$ for all other S. A single-minded bid is the pair (S^*, v^*).

Single-Minded Bidders

Truthful Mechanism

A mechanism for single-minded bidders in which losers pay 0 is incentive compatible if and only if it satisfies the following two conditions:

1. **Monotonicity**: A bidder who wins with bid (S_i^*, v_i^*) keeps winning for any $v_i' > v_i^*$ and for any $S_i' \subset S_i^*$.
2. **Critical Payment**: A bidder who wins pays the minimum value needed for winning: the infimum of all values v_i' such that (S_i, v_i') still wins.

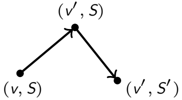

Single-Minded Bidders

The Greedy Algorithm $\mathcal{A}_G$

Initialization:

- $\frac{v_1}{\sqrt{|S_1|}} \geq \ldots \geq \frac{v_n}{\sqrt{|S_n|}}$
- $W \leftarrow \emptyset$

For i = 1..n do: if $S_i \cap \left(\bigcup_{j \in W} \right) = \emptyset$ then $W \leftarrow W \cup \{i\}$

Output:

- Allocation: W
- Payments: if $i \in W$ then $p_i = \frac{v_j}{\sqrt{|S_j|}} \cdot \sqrt{|S_i|}$, where j the smallest index such that $S_i \cap S_j \neq \emptyset$.

Single-Minded Bidders

- Optimal Allocation: OPT
- Winners of $\mathcal{A}_G$: W
- $OPT_i = \{j \in OPT, j \geq i | S_i \cap S_j \neq \emptyset\}$

$$\boxed{SW_{\mathcal{A}_G} \geq \frac{1}{\sqrt{m}} SW_{OPT}} \quad (\sqrt{m}\textbf{-approximation})$$

- $\sum_{j \in OPT_i} v_j \leq \frac{v_i}{\sqrt{|S_i|}} \sum_{j \in OPT_i} \sqrt{|S_j|} \leq \frac{v_i}{\sqrt{|S_i|}} \sqrt{|OPT_i|} \sqrt{\sum_{j \in OPT_i} |S_j|} \leq \frac{v_i}{\sqrt{|S_i|}} \sqrt{|S_i|} \sqrt{\sum_{j \in OPT_i} |S_j|} \leq v_i \sqrt{m}$
- $OPT \subseteq \bigcup_{i \in W} OPT_i$
- $\sum_{i \in W} v_i \geq \frac{1}{\sqrt{m}} \sum_{i \in W} \sum_{j \in OPT_i} v_j \geq \frac{1}{\sqrt{m}} \sum_{i \in OPT} v_i$

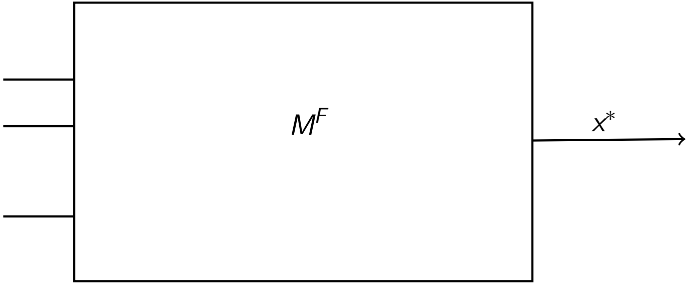

Truthful in Expectation

A randomized mechanism (f, p) is truthful in expectation if for any player i, any $v_{-i} \in V_{-i}$ and any $v_i, v_i' \in V_i$,

$$\mathbb{E}[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \geq \mathbb{E}[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})]$$

Randomized Mechanisms

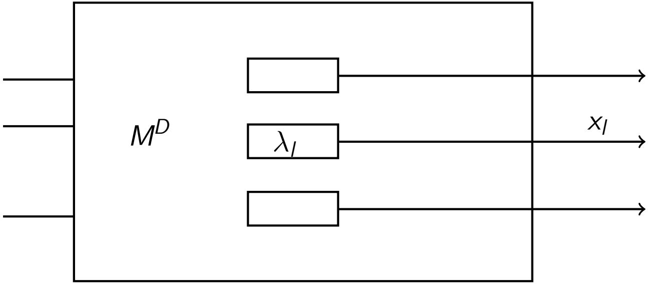

Truthful in Expectation

A randomized mechanism (f, p) is truthful in expectation if for any player i, any $v_{-i} \in V_{-i}$ and any $v_i, v_i' \in V_i$,

$$\mathbb{E}[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \geq \mathbb{E}[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})]$$

Randomized Mechanisms

The Truthful α-approximation support mechanism

1. Use VCG to get a truthful fractional mechanism M^F that outputs allocation $f^F(v) = x^*(v)$. the optimal solution to the LPR, and prices $p^F(v)$.
2. Use $\mathcal{A}$ to obtain the convex decomposition $\frac{x^*}{\alpha} = \sum_{l \in \mathcal{I}} \lambda_l x_l$ with only polynomially many λ_l.
3. Return the support mechanism $M^D = (f^D, p^D)$ with $f^D(v) = \{\lambda_l\}_{l \in \mathcal{I}}$ and $p^D(v) = \frac{p^F(v)}{\alpha}$.

Outline

Other Objective Functions

Setting (Job Scheduling)

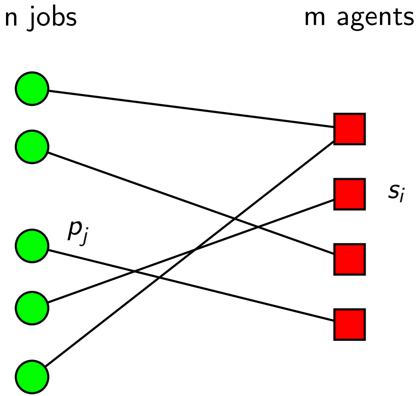

Other Objective Functions

Setting (Job Scheduling)

- Agents' Types: $\{t_i = \frac{1}{s_i}\}$
- Mechanism's Input: $\{b_i\}$
- Mechanism's Output: $o(\mathbf{b})$
- Agent's Profit: $profit_i(t_i, \mathbf{b}) = p_i(\mathbf{b}) - cost_i(t_i, o(\mathbf{b}))$
- Agent's Cost: $cost_i(t_i, o) = t_i w_i(o) = \frac{l_i(o)}{s_i}$

Other Objective Functions

- $p_1 \geq \ldots \geq p_n$
- $s_1 \geq \ldots \geq s_m$
- $T_j = \min_i \max\{\frac{p_j}{s_i}, \frac{\sum_{k=1}^{j} p_k}{\sum_{l=1}^{i} s_l}\}$
- $T_{LB} = \max_j T_j$

$$\boxed{T_{OPT} \geq T_{LB}}$$

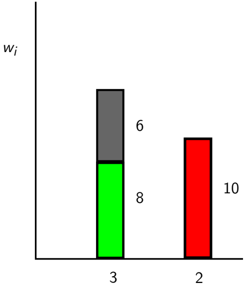

Other Objective Functions

The Fractional Algorithm $\mathcal{A}_F$

1. $j_1 = \min_j \{ j | \frac{\sum_{k=1}^{j} p_k}{s_1} > T_{LB} \}$
2. Assign jobs $1, \ldots, j_1 - 1$ and a fraction of j_1 to machine 1 $(w_1 = T_{LB} \cdot s_1)$

Continue recursively...

Feasibility and Truthfulness

- **1.** All jobs are assigned and **2.** $\frac{p_j}{s_{i(j)}} \leq T_{LB}$.
- $\mathcal{A}_\mathcal{F}$ is monotone.

Other Objective Functions

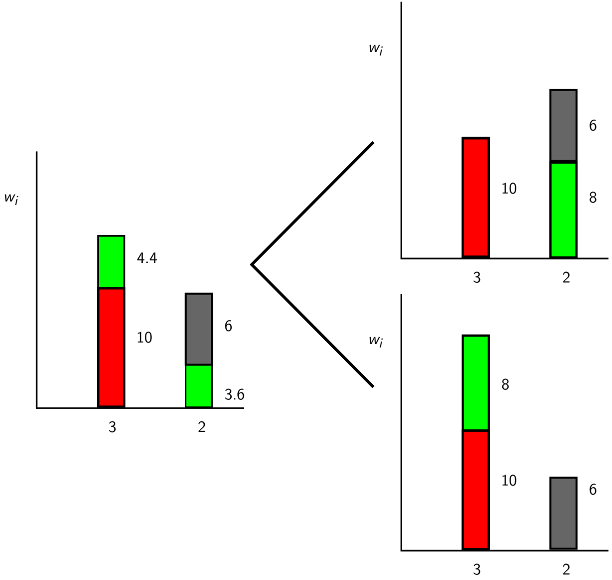

Other Objective Functions

Randomized Rounding

1. Choose $\alpha \in [0,1]$ uniformly
2. if (fraction in i) $\geq \alpha$, then assign to i, otherwise to $i+1$

$\mathcal{A}_F +$ randomized rounding

$\mathcal{A}_F +$ randomized rounding algorithm is truthful in expectation, and obtains a 2-approximation to the optimal makespan in polynomial-time.

Proof.

- Expected Load remains the same.
- Additional Cost $\leq a \cdot p_j + (1-a) \cdot p_k \leq T_{LB}$

□