Truthful Mechanisms and Shortest Paths

Rudolf Müller

Universiteit Maastricht
The Netherlands

joined work with Hongwei Gui and Rakesh V. Vohra
Kellog School of Management, Northwestern University

Guest Lecture
Foundations of Electronic Marketplaces
September 22, 2005

Motivation

- Given a **decision problem** (e.g., a combinatorial optimization problem)
- **Input:** split among different agents, private information.
- **Output:** computed by a center on basis of input reports.
- **Assumption:** agents vary in their valuations over the chosen output.
- **Incentives to manipulate:** report of wrong input might give favored decision.
- **Question:** can we introduce (side-)payments that depend on agent reports which give incentives to tell the truth?

Example: Minimum Spanning Tree

- **Setting:** every agent owns one edge in a graph, no agent owns a cut, edge lengths are private information.
- **Valuation:** if edge is chosen, cost to operate it is equal to the length.
- **Input:** every agent reports his edge length.
- **Output:** center chooses a minimum spanning tree based on these reports, pays cost of chosen edges.
- **Incentives to manipulate:** higher reports increase revenue, while there might be still a chance to be in the minimum spanning tree.
- **Question:** which payment should be made in order to make agents tell the truth?

Example II: Sealed bid, single item auction

- **Setting:** 1 item, n bidders, we want to reward the item to the bidder who values it most.
- **Valuation:** each bidder has a private value for the item.
- **Input:** maximum price each bidder is willing to pay.
- **Output:** the winner of the auction, and a price he has to pay.
- **Incentives to manipulate:** lower reports might still win, but result in a lower price.
- **Question:** which payment gives incentives to tell the truth?
- **Vickrey (1961):** if the payment is equal to the highest loosing bid, then every bidder cannot do better than reporting his valuation.

Mechanism Design Setting

$N = \{1, \ldots, n\}$ set of agents.

T_i set of *types*, $t_i \in T_i$ type of agent i, $T = \times_{i \in N} T_i$

Y set of *outcomes* (output)

$v_i : Y \times T_i \to \mathbb{R}$ *valuation* of agent i for outcome y if of type t_i.
Notation: $v_i(y|t_i)$.

A *social choice function*: $f : T \to Y$.
A *payment function*: $p : T \to \mathbb{R}^n$.

We assume *(quasi-linear) utilities*: agent i, if of type t_i, values outcome y and payment p_i:

$$u_i : Y \times \mathbb{R} \times T_i \to \mathbb{R}$$

$$u_i(y, p_i|t_i) = v_i(y|t_i) - p_i$$

Example: Minimum Spanning Tree

Every agent in $N = \{1, \ldots, n\}$ owns an edge e_i.

$T_i = [0, \infty)$ represents costs c_i of operating edge e_i

Y set of spanning trees

$$v_i(y|c_i) = \begin{cases} -c_i & \text{if } e_i \in y \\ 0 & \text{else} \end{cases}$$

Social choice function: any algorithm that computes a spanning tree.

A *payment function*: $p : T \to \mathbb{R}^n$ rewards agents in the spanning tree.

Mechanism Design Setting (II)

Definition. (f, p) is *dominant strategy incentive compatible* if and only if for all i, for all $s_i, t_i \in T_i$, for all $t_{-i} \in T_{-i}$:

$$v_i(f(t_i, t_{-i})|t_i) - p(t_i, t_{-i}) \geq v_i(f(s_i, t_{-i})|t_i) - p(s_i, t_{-i})$$

Definition. f is *dominant strategy incentive compatible* (or *implementable*) if and only if there exists a payment function p such that (f, p) is dominant strategy incentive compatible.

Example: Choose a weight w_y for every $y \in \Gamma$, and a multiplier q_i for every agent, then the *weighted utilitarian social choice function*

$$f(t) \in \text{argmax}\{w_y + \sum_{i=1}^{n} q_i v_i(y|t_i) \mid y \in Y\}$$

is dominant strategy incentive compatible (Vickrey (1961), Clarke (1970), Groves (1971), Roberts (1979)).

Example: VCG payment for minimum spanning tree

Choose allocation rule *min total cost* (= max -(total cost)) and define

$$p_i(c_i) = c_i(y^{opt}) + (c_{-i}(y_{-i}^{opt}) - c(y^{opt}))$$

Interpretation: Add to declared cost the marginal increase of the cost of a minimum spanning tree if bidder i would not be present.

Let y' be the solution chosen if bidder i reports c_i' rather than his true type c_i, and p' the payment.

Note that:

$$\begin{aligned}
u_i(y', p'|c_i) &= -c_i(y') + (c_i'(y') + (c'_{-i}(y_{-i}^{opt}) - c'(y'))) \\
&= c_{-i}(y_{-i}^{opt}) - c(y') \\
&\leq c_{-i}(y_{-i}^{opt}) - c(y^{opt}) \\
&= -c_i(y^{opt}) + (c_i(y^{opt}) + (c_{-i}(y_{-i}^{opt}) - c(y^{opt}))) \\
&= u_i(y^{opt}, p|c_i)
\end{aligned}$$

Mechanism Design Setting (III)

Problems:

- In many settings (e.g., combinatorial auctions) computing the weighted utilitarian s.c.f. is NP-complete or reporting the types requires exponential communication.
- Sometimes the utilitarian s.c.f. is not in the interest of the center (e.g., task scheduling (Nisan and Ronen, 2001)).
- Sometimes the center might want to modify the utilitarian social choice function in order to increase revenue (e.g., optimal auctions).
- Very little is known about implementable social choice functions (i.e., algorithms) for multi-dimensional type spaces.

This presentation

- Given type spaces $T_i \subseteq \mathbb{R}^k$, outcomes Y, valuations $v_i(y|t_i)$
- Characterize
$$f : \times_{i \in N} T_i \rightarrow Y$$
that are dominant strategy incentive compatible.
- Approach: a construction of payments or a prove that no payments exist.
- Based on relation between shortest paths and negative cycles.

Outline

- Introduction to Mechanism Design
- Projection to the single-agent case
- Allocation graphs
- Necessary condition: no negative 2-cycles
- Environments in which no negative 2-cycles is sufficient:
  - Combinatorial auctions with bounded type domains
  - Multi-item auctions with decreasing marginal utilities

Projection to the single-agent case

Given N, T, v, f, p.
Fix agent i, and report of the other agents t_{-i}.
Define:

$$f_i(t_i) = f(t_i, t_{-i})$$
$$p_i(t_i) = (p(t_i, t_{-i}))_i$$

For simplicity we drop the dependence on t_{-i} in our notation.

Lemma f is dominant strategy incentive compatible if and only if for all $i \in N$ and $t_{-i} \in T_{-i}$, f_i is dominant strategy incentive compatible (or *rationalizable*).

Allocation Graph

From now on: fix i, t^{-i}, drop i in the notation: $f = f_i$.

Define an (infinite) digraph $G = (T, A)$, arc lengths

$$l(s,t) = v(f(t)|t) - v(f(s)|t).$$

Note that f is dominant strategy incentive compatible if and only if there exists p such that for all s, t:

$$\begin{aligned} p(t) - p(s) &\leq v(f(t)|t) - v(f(s)|t) \\ &= l(s,t) \end{aligned}$$

Theorem (Rochet, 1987). f is dominant strategy implementable (or, rationalizable) if and only if G does not contain a negative length cycle.

Proof

Necessity.

Let p be a payment function that makes f d.s.i.c.. Let $C = \{t_0, \ldots, t_{k-1}, t_k = t_0\}$ be a cycle.

$$
\begin{aligned}
l(C) &= \sum_{j=0}^{k-1} l(t_j, t_{j+1}) \\
&\geq \sum_{j=0}^{k-1} p(t_{j+1}) - p(t_j) = 0
\end{aligned}
$$

Sufficiency.

Fix a type t_0 and define $p(t_0) = 0$ and $p(t) = \inf(l(P) \mid P$ is a path from s to $t)$.

Note that $p(t) > -\infty$, since $l(P) - l(t, t_0) \geq 0$.

Finally, observe that $p(t) \leq p(s) + l(s, t)$, that is this payment makes f d.s.i.c..

A special case

- Homogeneous, multi-item auctions, k items on sale.
- Additive valuations: $t \in R^k$, $v(k|t) = \sum_{j=1}^{k} t_j$.
- Length of an edge (if $f(s) \leq f(t)$):

$$l(s,t) = \sum_{j=1}^{f(t)} t_j - \sum_{j=1}^{f(s)} t_j = \sum_{j=f(s)+1}^{f(t)} t_j$$

Theorem (Bikhchandani, Chatterij, Sen, 2003). f is dominant strategy incentive compatible if and only if for all s, t, $f(s) < f(t)$:

$$\sum_{j=f(s)+1}^{f(t)} t_j - \sum_{j=f(s)+1}^{f(t)} s_j \geq 0.$$

$$(\Leftrightarrow \quad l(s,t) + l(t,s) \geq 0).$$

Corollary. G does not have a negative length cycle, if and only if it does not have a negative length 2-cycle.

Question: Does this generalize to other settings?

Weak monotonicity sufficient for d.s.i.c.

When are the following equivalent:

1. G does not have a negative cycle
2. G does not have a negative 2-cycle (called *weak monotonicity*)

Lavi, Mu'alem, and Nisan (2003): Combinatorial auctions, if valuations are non-negative, and monotone (free-disposal).
Gui, M., and Vohra (2004): Holds for the environments:

- Y finite, $T = \mathbb{R}^Y$, or $T = \mathbb{R}^Y_+$, or $T = [0,1]^Y$,
- Combinatorial auctions, $T = \mathbb{R}^{2^S}$, or $T = \mathbb{R}^{2^S}_+$, or $T = [0,1]^{2^S}$,
- Multi-item auctions with *decreasing marginal utilities*, i.e., $T = \{t \in [0,1]^m \mid t_{k+1} \le t_k, k = 1, \ldots, m-1\}$,

Saks and Yu (2004): Holds for any finite Y, $T \subset \mathbb{R}^Y$, $v(y|t) = t_y$, T convex.

Core idea of proofs in Gui, M., and Vohra (2004)

To prove:

$$G \text{ does not have a negative 2-cycle} \\ \Rightarrow \\ G \text{ does not have a negative cycle}$$

Step 1: Study instead of a graph on T a graph on Y:

- $\Gamma = (Y, A)$
- for $\alpha \in Y$ let $R_\alpha = \{t \in T \mid f(t) = \alpha\}$
- for $\alpha, \beta \in Y$ define

$$l(\alpha, \beta) = \inf\{v(f(t)|\, t) - v(f(s)|\, t) \mid s \in R_\alpha, t \in R_\beta\}$$

Core idea of the proofs (II)

Step 2: Verify that Γ has no negative (2-)cycle if and only if G has no negative (2-)cycle.

Step 3: Assume Γ has no negative 2-cycle. Observe that for all $t \in R_\beta$ and for all $\alpha \neq \beta$:

$$\begin{aligned} v(\beta \mid t) - v(\alpha \mid t) &\geq \inf\{v(\beta \mid s) - v(\alpha \mid s) \mid s \in R_\beta\} \\ &= l(\alpha, \beta) \\ &\geq -l(\beta, \alpha) \end{aligned}$$

Step 4: Observe that in our cases $v(\beta|t) - v(\alpha|t)$ is linear in t.

Thus

$R_\beta \subseteq Q_\beta := \{t \in \mathbb{R}^k \mid v(\alpha \mid t) - v(\beta \mid t) \leq l(\beta, \alpha)$ for all $\alpha \in Y\}$,

where Q_β are polyhedra (k is the dimension of T).

Core idea of the proofs (III)

Step 5: Observe that

$$T = \bigcup_{\alpha \in Y} (T \cap Q_\alpha)$$

and that Q_α intersect only on their boundaries.

Step 6: Now suppose there exists a strictly negative cycle C in Γ. Assume for a moment, that this cycle visits every node in Y:

$$C = [\alpha_0, \ldots, \alpha_{k-1}, \alpha_k = \alpha_0].$$

Observe: $l(C) < 0$ implies that the system:

$$z_{i+1} - z_i \geq l(\alpha_i, \alpha_{i+1}) + \epsilon \quad \text{for all } i = 0, \ldots, k-1$$

has a feasible solution (set $z_0 = 0$, compute longest paths).

Core idea of the proofs (IV)

Step 7: Let us look at the case $T = \mathbb{R}^Y$, i.e., $v(y|t) = t_y$.

Consider the following type: $t_i = z_i$ where z_i is a solution of the system above.

Observe:

$$v(\alpha_{i+1}|t) - v(\alpha_i|t) = t_{i+1} - t_i > l(\alpha_i, \alpha_{i+1})$$

Thus $t \notin Q_{\alpha_i}, i = 1, \ldots, k$.

But observe that, in this case, $t \in T$. Therefore, we have found a type for which f doesn't define an allocation. A contradiction.

Complication: C may not contain all nodes in Y.

Solution: add extra arcs such that the extended system does not have a positive cycle.

Core idea of the proofs (V)

This proves the theorem for general valuations and for combinatorial auctions, with arbitrary, and with non-negative valuations.

For the cases of bounded valuations we have to overcome another:

Complication: The solution z may not be an element of T.

Solution: build up a careful line of argumentation that constructs a solution $z \in T$ (doable, but lengthy...)

Literature

Hongwei Gui, R.M., Rakesh V. Vohra (2004), *Dominant Strategy Mechanisms with Multidimensional Types*, http://www.kellogg.northwestern.edu/research/math/dp_list_4.htm

Roberts (1979). If $|Y| = k$ and $T = \mathbb{R}^k$ then:
f is d.s. implementable $\Leftrightarrow$ f is weighted utilitarian maximizer.

Lavi, Mu'alem, and Nisan (2003). Weak-monotonicity is sufficient for combinatorial auctions. (appeared in STOCS Proceedings 2003)

Saks and Yu (2005). Weak monotonicity is sufficient for truthfulness on any convex type domains. (appeared in Proceedings of ACM Conf. Electr. Commerce, 2005).

Related research

Bayesian setting:

- Projection does also work.
- Γ is not finite anymore, but a set of probability distributions.
- studied by Rochet (various papers), Jehiel and Moldovanu (2001)
- Generalizations of previous results can be found in:
  M., Perea, and Wolf. "Weak Monotonicity and Bayes-Nash Incentive Compatibility" (available on request).