CS294 P29: Algorithmic Game Theory **November 7, 2011**

Lecture 9: Mechanism Design

Lecturer: Christos Papadimitriou *Scribe: Yusef Shafi, Michael Landis*

9.1 Mechanism Design

In Lecture 8, we saw how Social Choice Theory (SCT) leads to Mechanism Design (MD): In SCT, we want to reconcile the various preferences of the agents. In MD, we don't know these preferences, but we have an objective that we wish to maximize (typically but not exclusively, total welfare), had we known them. We construct a game whereby the outcome will be chosen. This game always has a dominating outcome, and this outcome is the optimum of our objective for the players' utilities. Vickrey's auction is an interesting example: It compels the agents to reveal their true valuations and achieves the social optimum (the agent who wants the item most gets it). Here we look at the general setting.

Consider n players and a set of A alternative outcomes. Each player has a *private* utility function, $u_i \in U : A \mapsto \mathbb{R}$, where $u_i(a)$ denotes the worth of outcome $a \in A$ to player i. U is some set of utility functions. There is also the designer's objective $C : A \times U^n \mapsto \mathbb{R}$, assigning a real number to all possible situations (utilities for the players and outcomes). One important example is social welfare: $f(a, u_1, \ldots, u_n) = \sum_i u_i(a)$.

We are interested in designing a game, G. Each player's strategy set is U — that is, every player chooses a utility function (for example, her true utility, or a fictional one). We also design a mapping $f : U^n \mapsto A$, mapping all possible plays to outcomes in A.

What are the payoffs of the game?

$$u_i^G(v_1, ..., v_n) = u_i(f(v_1, ..., v_n)),$$

that is, if everybody plays v_i (some utility in U), then the ith player gets a payoff equal to her own true utility for the outcome that results, through the designed map f, from these choices.

Since G depends on the utilities, the designer does not know G. The designer wants G to have a dominant strategy play $(v_1, \ldots, , v_n)$ such that $f(v_1, ..., v_n) \in \arg\max_{a \in A} C(a, u_1, ..., u_n)$. The Revelation Principle (see below) says, we might as well assume that the play is $(u_1, ..., u_n)$, that is, everybody "tells the truth."

A dominant strategy means

$$\forall i \forall u_i \forall u_i' \forall u_{-i} : u_i^G(u_i, u_{-i}) \geq u_i^G(u_i', u_{-i})$$

However, we conclude the following from what we learned in Lecture 8.

Theorem 9.1. *Arrow's Impossibility Theorem ⇒ Gilbert-Satterthwaite ⇒ our ideal game design mechanism is impossible for* $|A| \geq 3$.

9.2 Restricted Domains and the Vickery's Auction

Gilbert-Satterthwaite holds when $U : A \mapsto \mathbb{R}$, the most general form of utility functions. But what happens if we create restrictions on these functions? What opportunities arise? For semilinear domains $A = (A_0, \mathbb{R}^n)$, which are the basic outcomes and the payments by agents, respectively. $u_i(a, p_1, ..., p_n) = v_i(a) - p_i$, where $v_i(a)$ is the private value individual i holds for outcome a, and p_i is the payment individual pays. The game mechanism, $f(v_1, ..., v_n) = (a, p_1, ..., p_n)$, is designed to maximize some objective $C(a, u_1, ..., u_n)$ as defined by the designer.

For example, take Vickery's auction, which attempts to maximize social welfare. Recall $G = S_i$ (by the direct revelation mechanism), where S_i is equal to all possible $v_i : A \mapsto \mathbb{R}$. So, for Vickrey's auction, we define a mechanism, $f : V^n \mapsto (A, p_1, ..., p_n)$, where $V = A \mapsto \mathbb{R}$.

For a single-item auction, $A = 1, 2, ..., n$ and

$$v_i(j) = \begin{cases} 0 & \text{if } i \neq j \\ v_i(a) & \text{if } i = j \end{cases}$$

so the game designer defines $f(v_1, ..., v_n) = (k = \arg\max(v_i))$, where

$$p_i = \begin{cases} 0 & \text{if } i \neq k \\ \max_{j \in \{1,...,n\} \setminus i} v_j(a) & \text{if } i = k \end{cases}$$

where p_i conditionally equals 0 or the second highest valuation among players. Now, a mechanism is incentive compatible (IC) if $\forall i \forall v_i : v_i' \neq v_i$. If $f(v_i) = (a, p_1, ..., p_n)$ and $f(v_i', v_{-i}) = (a', p_1', ..., p_n')$, then $v_i(a) - p_i \geq v_i(a') - p_i'$. So,

Theorem 9.2. *A Vickery auction is IC.*

Why is this result beneficial? Why not design a game which allows players to behave naturally, i.e. lie? It turns out there only IC game design mechanisms exist.

Theorem 9.3. (Revelation Principle) *If there is a mechanism, then it is a truthful mechanism.*

Proof. Suppose you design another mechanism. Because all utilities are private, each player will select a strategy that maximizes her interests given the possible outcomes. Suppose that this is computed by an algorithm M_i for the ith player. That is, the ith player inputs her true utility u_i to the algorithm, and it computes the best strategy, and submits then it to the mechanism G. But then, the game G *together* with the algorithms M_i is an *IC* mechanism! □

Vickery's auction works well for two reasons. First, it is in the semilinear domain. Second, it maximizes $\sum_i v_i(a)$, which is the social welfare. Very few alternatives to Vickery auctions have been discovered for different conditions.

9.3 Vickery-Clarke-Groves Auction

The Vickery-Clarke-Groves (VCG) presents a general-form solution for IC mechanisms. The VCG designer defines a mechanism function, $f(v_1, ..., v_n) = (a, p_1, ..., p_n)$, where $a \in \arg\max_A \sum_i v_i(a)$, and defines the prices as

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

Note, for some a each player i is paid according to $v_i(a)$, the social welfare is $\sum_i v_i(a)$, and something the player cannot control (i.e. some tax suffered regardless of what that player declared). That is

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

so VCG aligns the players' and the designer's interests. Since players want to assist the designer maximize social welfare, we have

Theorem 9.4. *VCG is IC.*

Proof. Recall the definition of IC. Since $v_i(a)-p_i \geq v_i(a')-p'_i$, then $\sum_j v_j(a)-h_i(v_{-i}) \geq \sum_j v_j(a')-h_i(v_{-i})$, since we can add $h_i(v_{-i})$ to both sides of the inequality, and a maximizes social welfare. $\square$

9.4 Clarke Pivot Rule

Now, we consider how to best define h_i. An uninteresting option is to define $h_i = 0$, but this introduces problems when the game payoff is large. Instead, consider the Clarke pivot rule (payment)

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

where prices are defined as

$$p_i = \max_{b \in A} \sum_{j \neq i} v_j(b) - \sum_{j \neq i} v_j(a)$$

where $p_i = 0$ if you are selected to get the item (absorb the cost). Otherwise, you earn value defined.

Definition 9.1. Individually rational (IR) *A mechanism is individually rational (IR) if all players receive nonnegative utility.*

Definition 9.2. No positive transfers (NPT) *A mechanism has no positive transfers if no player is ever paid money.*

Because NTR and IR, then

$$u_i = \sum_i v_i(a) - \sum_{j \neq i} v_j(b) \geq 0$$

is true. So, player i pays an amount equal to the difference in the total social welfare given player i's inclusion and exclusion.

9.5 Examples

9.5.1 Single Item Auction

The Vickrey auction is a special case of a VCG mechanism with the Clarke pivot rule. We have $A = \{i-\text{wins} | i \in I\}$. Each player receives 0 if he doesn't win, and any positive value if he does win. Thus, $V_i = \{v_i | v_i(i - \text{wins}) \geq 0 \text{ and } \forall j \neq i, v_i(j - \text{wins}) = 0\}$. Finding the player with the highest value is the same as maximizing $\sum_i v_i(i)$ because only one player gets a nonzero value. Thus, VCG payments using CPR yield Vickrey's second price auction.

9.5.2 Reverse Auction (Procurement)

Here, the goal is to procure an item from the bidder with the lowest cost. Each contractor has a valuation $V_i = \{c_i | v_i(i - \text{wins}) \leq 0 \text{ and } \forall j \neq i, v_i(j - \text{wins}) = 0\}$. In this case, maximizing the social welfare means choosing the contractor with the least cost (the least negative c_i). Contractors should be truthful. The VCG payment rule is for the mechanism to pay the least cost bidder the magnitude of the second lowest bid and pay nothing to the other bidders. The second best bid is what would happen if the winning bid did not exist, so the absolute difference between the best and second-best bids is the benefit to society.

9.5.3 Public Project

Suppose the government is considering whether or not to build a public project, such as a bridge. The project cost is C. The possible outcomes are that the project is not built or that it is built. Each player has a privately-known and possibly negative valuation v_i. The government will go forward with the project if $\sum_i v_i > C$. Under the CPR, player i with positive valuation will pay only if he or she is pivotal: that is, without i, $\sum_{v \neq i} v_j \leq C$ but $\sum_j v_j > C$. Player i in this case will pay $p_i = C - \sum_{j \neq i} v_j$. However, it can be shown that $\sum_i p_i < C$ unless $\sum_i v_i = C$, and so payments will not cover costs. In reality, no individual in a society would be pivotal. It can be shown rigorously that such a project can only be built with subsidy.

Theorem 9.5. *(Meyerson) The only incentive compatible mechanisms that maximize social welfare are those with VCG payments. The payment function is essentially uniquely determined by the social choice function.*

No ex-post IR mechanism that maximizes social welfare can recover the cost of the project.

9.5.4 Bilateral Trade

Suppose a seller has valuation $0 \leq v_s \leq 1$ and a buyer has valuation $0 \leq v_b \leq 1$. The possible outcomes are no-trade and trade. The outcome must maximize the welfare, and we have $a =$ "$v_b > b_s$". That is, trade happens if $v_b > v_s$. VCG and CPR imply that the buyer pays v_s and the seller receives v_b since it is necessary that no payment be made in the case of no-trade. Thus, CPR implies that the mechanism must subsidize the trade, which follows from the result due to Meyerson in the previous example.

In the bilateral trade problem, the only IC mechanism that maximizes social welfare and makes no payments in the case of no-trade subsidizes trade. If a mechanism satisfies ex-post IR, it cannot dictate positive payments from the players in the case of no-trade and so it subsidizes trade.

9.5.5 Combinatorial Auction

Consider an auction with items $1, \ldots, m$. Suppose that their are bidders $1, \ldots, n$ each of whom want a bundle $S_i \subseteq \{1, \ldots, m\}$ with value v_i. A valuation must have free disposal, that is, be monotone: for $S \subset T$, we have $v(S) \leq v(T)$, with $v() = 0$. The problem is to find $x \in \{0, 1\}^n$ such that xv is maximized and $x_i, x_j = 1 \rightarrow s_i \cap s_j \neq$.

It turns out that this problem is equivalent to the NP-hard weighted set packing problem. In fact, NP-hard approximations often result in violations of IC, or result in the break down of the proof of IC.

9.5.6 Combinatorial Public Project

In this problem , the government will build k of n possible projects. The outcomes for the problem are $A = \{(a_1, \ldots, a_k) : a_i \in \{1, \ldots, n\}\}$. The valuations are $v_i : A \to \mathbb{R}$. It turns out that finding an optimal mechanism for this problem is also NP-Hard, even if v_i is submodular (simple with diminishing returns).

9.5.7 Buying a Path

Consider a graph $G = (V, E)$ with edges $e \in E$ and a source and sink s and t, respectively. With each edge is associated a private cost $c(e)$. Maximum social welfare is achieved by choosing the shortest path. Each edge e should be paid $p = c(e) + \Delta($ shortest path: with vs. without $e)$. (Note that if an edge e were pivotal in the sense that if it were missing, there would be no path from s to t, then the cost would be infinite without the edge.) The problem can be solved with $n+1$ applications of Dijkstra's algorithm. In 2002, Hershberger et al. showed that in the case of undirected graphs, the problem can be solved with only one application of Dijkstra's algorithm.

Frugality Consider a graph with two paths from s to t, one involving a single link of cost 20 and the other ten segments of cost 1 each. Under VCG, the payment would be 110 because each edge would receive a bonus of 10. VCG overpays (the price it pays is much more than the cost of the second-cheapest path), and thus it is not "frugal."

Theorem 9.6. *(Frugality) The payment could be improved, but the ratio of the VCG path to the cheapest Nash path is lower bounded by $\sqrt{n}$.*

9.6 Beyond VCG

Here, we study of the question of whether we can implement any other designer objectives than $\max_A \sum_i v_i(a)$. The answer is a qualified negative.

Theorem 9.7. *(Roberts) If $|A| \geq 3$ and $V = \mathbb{R}^A$, then the only IC mechanisms are those that optimize $\max_{a \in A' \subset A} \sum_i w_i v_i(a)$, where w_i are arbitrary positive weights.*

Theorem 9.8. *(PSS 2009) This also holds for CPP (does NOT hold for combinatorial auctions).*

It seems as though we have arrived at a very negative result in that VCG is all we can do in general. In fact, valuations often have specific structures more restrictive than the general $\mathbb{R}^A$.

Consider the special case $V_i = \mathbb{R}$ (single-parameter domains). Suppose that $W_i \subset A$ are the outcomes where i wins. Let

$$v_i(a) = \begin{cases} -v_i, & a \in W_i \\ 0, & \text{o/w (also: } p_i = 0). \end{cases} \tag{9.1}$$

Theorem 9.9. *The function f is IC if and only if it is monotone (if v_i wins then $v_i' > v_i$ wins)) and the payment is equal to the smallest v_i that wins.*

Thus, for instance, one can optimize $\sum_i v_i^2(a)$.

Theorem 9.10. *Payments are unique up to $\pm h_i(v_{-i})$.*

We can also conclude, as mentioned in the examples, that bilateral trade and public project construction cannot be done without subsidies: one cannot use mechanisms to come up with good designs for all problems.