CS295 Introduction to Algorithmic Game Theory

Instructor: Ioannis Panageas
Scribed by: Hari Kishore Chaparala, Calvin Yeung, Jason Alaya

Lecture 11. Introduction to Mechanism Design.

1 Approaches to Solving the Single Allocation Problem

We start by considering a simple setting given by the single allocation problem.

1.1 Problem Setting

Definition 1.1 *The single item allocation problem is given by:*

1. *A single indivisible item*
2. *n agents competing for the item*
3. *An associated value/valuation v_i for the item for each agent i*

The goal is to maximize the social surplus, which, in this case, is the value of the agent who receives the item or to allocate the item to the user with highest valuation.

1.2 Naive Approach

Given the single allocation problem, one might intuitively come up with the following approach:

1. Ask each agents/bidders to report (bid) their values b_i
2. Select agent i^* with the highest bid, $i^* = \arg\max_i b_i$
3. Allocate item to i^*

This approach is problematic. Since agents do not need to pay the amount they bid even if they get the item, they are incentivized to report an arbitrarily high bid to ensure that they get it. Thus, the outcome of this mechanism is unpredictable, which makes it hard to reason about performance. We will look at more reasonable approaches in the following subsections.

1.3 Lotteries

Definition 1.2 (Lottery) *In the lottery mechanism, we*

1. *Select a uniformly random agent*
2. *Allocate the item to that agent*

This mechanism is predictable and thus easy to reason about. The following theorem gives us a bound on the expected surplus of the lottery mechanism.

Theorem 1.1 *The lottery mechanism has n-approximation ratio (i.e. $\frac{OPT}{alg} \leq n$).*

Proof: Assume $v_1 = 1$ and $v_i = \epsilon$ for $i \geq 2$. Then, the expected surplus is

$$1 \cdot \frac{1}{n} + \epsilon \cdot \frac{n-1}{n} = \frac{1 + \epsilon(n-1)}{n}$$

Since the optimal surplus is 1 (agent 1 gets the item), the approximation ratio is $\frac{n}{1+\epsilon(n-1)}$. This converges to n as $\epsilon \to 0$. ■

This approximation ratio is quite bad. One way to design better mechanisms is to introduce payments. We will discuss such mechanisms next.

1.4 First-Price Auctions

Definition 1.3 (First-Price Auction) *In a first-price auction,*

1. *Each agent report their bids b_i*
2. *Select agent $i^* = \arg\max_i b_i$*
3. *Agent i^* gets the item and pays b_{i^*}*

The problem with first-price auctions is that they are hard to reason about. Thus, it is hard for participants to figure out how to bid and hard for designers to predict what will happen.

1.5 Second-Price Auctions (Vickrey Auctions)

Definition 1.4 (Second-Price Auction) *In a second-price auction,*

1. *Each agent report their bids b_i*
2. *Let agent $i^* = \arg\max_i b_i$ and j^* the agent with the second highest bid*
3. *Agent i^* gets the item and pays b_{j^*}*

Second-price auctions are predictable and easier to reason about. The following theorem is a result about the dominant strategy for participants in a second-price auction.

Definition 1.5 (Dominant strategy) *A strategy is said to be dominant if the utility for playing this strategy is no less than that for playing any other strategy, regardless of what strategies other players play.*

In the case here, for agent i, the utility is $u_i = v_i - b_{j^*}$ if i gets the item and $u_i = 0$ otherwise.

Theorem 1.2 (Vickrey is truthful) *In second price auctions, every bidder i has a dominant strategy, which is to bid truthfully; i.e. set $b_i = v_i$.*

Proof: Fix an agent i and set $B = \max_{j \neq i} b_j$ Consider the cases:

1. if $b_i < B$ then i gets utility 0 as the agent doesn't get the item
2. if $b_i \geq B$ then i wins the item and $u_i = v_i - B$

   Now consider the cases

   (a) $v_i < B$

      i. If the bid is truthful, $u_i = 0$ as $b_i = v_i < B$

      ii. If bid is less than valuation, we have $u_i = 0$ as $b_i < v_i < B$

      iii. If bid is greater than valuation, we have $u_i = 0$ when $B > b_i > v_i$ and $u_i < 0$ when $b_i \geq B > v_i$

   (b) $v_i \geq B$

      i. If the bid is truthful, $u_i = v_i - B \geq 0$ as $b_i = v_i \geq B$ and the agent wins the item.

      ii. If bid is less than valuation, we have $u_i = v_i - B \geq 0$ when $v_i > b_i \geq B$ and $u_i = 0$ when $v_i \geq B > b_i$ as the agent looses the item.

      iii. If bid is greater than valuation, we have $u_i = v_i - B > 0$ as $b_i > v_i \geq B$ and the agent wins the item.

From above, we can see that for scenarios **If the bid is truthful, If bid is less than valuation, and If bid is greater than valuation**, no matter the different cases of $bi \lesseqqgtr B$ and $v_i \lesseqqgtr B$, if the bid is truthful the agent get the maximum utility. So **truth telling** is the **Dominant Strategy** in Vickrey Auction.

■

Corollary 1.3 *Every agent with dominant strategy in Vickrey Auction has non-zero utility.*

2 A General Approach to Auctions

In this section, we will look at more general forms of auction mechanisms.

2.1 Desirable Properties of Auctions

Auctions should satisfy the following properties:

1. Dominant strategy incentive compatible (DSIC); i.e. reporting bids truthfully is a dominant strategy.
2. If bidders are truthful, then the auction achieves maximum surplus $\sum_{i=1}^{n} v_i x_i$, where $x_i = 1$ if bidder i wins, and $x_i = 0$ otherwise.
3. The auction can be implemented in polynomial time.

2.2 An Example

Let us now look at a more complex case.

Example 2.1 *Suppose there is a society with n citizens and a public good G, where*

- *each citizen i has a private valuation v_i for G (here we assume $v_i \geq 0$ for simplicity)*
- *the cost of building G is C and is publicly known.*
- *G should be built if $\sum_{i=1}^{n} v_i \geq C$*

Our goal is to design a mechanism that charges citizens such that G is built only if $\sum_{i=1}^{n} v_i \geq C$.

A naive approach to this example would be to allocate the cost C equally to all citizens. With this payment rule, each citizen would have to pay C/n. Then, if for citizen i, $v_i > C/n$, he/she has incentive for G to be built. Hence, citizen i can just report valuation $C+\epsilon$, such that $\sum_{i=1}^{n} v_i > C$, for G to be built. Hence, this mechanism is not DSIC.

Instead, we should charge citizen i the amount $p_i = \max(0, C - \sum_{j \neq i} v_j)$ if G is built and $p_i = 0$ otherwise. This is DSIC as we will see later.

2.3 VCG Mechanisms

Definition 2.1 (Quasi-Linear environment) *A quasi-linear environment, or Vickrey-Groves-Clark (VCG) environment, is defined as*

- *A set of n agents*
- *A set of finite outcomes $\mathcal{X}$*
- *Valuation functions $v_i : \mathcal{X} \to \mathbb{R}^+$ for each agent i*
- *Utilities $u_i = v_i - p_i$, where p_i is the received payment of agent i (which can be positive or negative)*

Definition 2.2 (VCG Mechanism) *The family of mechanisms is defined as follows:*

- *Agents have private valuations v_i and report their bids b_i*
- *Set $x^* = \arg\max_{x \in \mathcal{X}} \sum_{i=1}^{n} b_i(x)$*
- *Each agent pays $p_i = h_i(b_{-i}) - \sum_{j \neq i} b_j(x^*)$*
- *Each agent has utility $u_i = v_i(x^*) - p_i(x^*)$*

It is important to note that there is a bid $b_i(x)$ for each possible outcome $x \in \mathcal{X}$.

Theorem 2.2 (VCG is DSIC) *Every VCG mechanism is DSIC.*

Proof: Fix agent i, and let $x^* = \arg\max_{x\in\mathcal{X}} \sum_{i=1}^n b_i(x)$ where i reports $b_i = v_i$ (i.e. truthfully). Similarly, let x' be the argmax if i reports $b_i \neq v_i$ (i.e. not truthfully) and all other agents report truthfully $b_j = v_j \forall j \neq i$. Then by the definition of utility,

$$\begin{aligned} u_i &= v_i(x^*) + \sum_{j\neq i} b_j(x^*) - h_i(v_{-i}) \\ &= v_i(x^*) + \sum_{j\neq i} v_j(x^*) - h_i(v_{-i}), \text{by the definition of } x^* \\ &= \sum_j v_j(x^*) - h_i(v_{-i}) \end{aligned}$$

Similarly,

$$\begin{aligned} u_i' &= v_i(x') + \sum_{j\neq i} b_j(x') - h_i(v_{-i}) \\ &= v_i(x') + \sum_{j\neq i} v_j(x') - h_i(v_{-i}), \text{by the definition of } x' \\ &= \sum_j v_j(x') - h_i(v_{-i}) \end{aligned}$$

Here, u_i is i's utility if he/she reports truthfully and u_i' is i's utility if not. Since x^* is the argmax of $\sum_{j=1}^n v_j(x)$, it follows that

$$\sum_{j=1}^n v_j(x^*) \geq \sum_{j=1}^n v_j(x) \ \forall x \in \mathcal{X}$$

It follows trivially that $u_i \geq u_i'$, which proves the theorem. ■

2.4 Clark Pivots

In the definition of VCG mechanisms, the function h_i is not specified. Hence, a question that arises naturally is how we should choose h_i. A requirement would be to choose h_i such that utility is non-negative (if $v_i \geq 0$ for all i), so that agents have incentive to participate in the mechanism. A Clark pivot is one such choice for h_i.

Definition 2.3 (Clark Pivot) *A Clark pivot is a choice for h_i such that*

$$h_i = \max_{x\in\mathcal{X}} \sum_{j\neq i} b_j(x)$$

Lemma 2.3 *VCG with Clark pivots has non-negative utilities if each agent's valuation v_i is non-negative.*

Proof: Assume $v_j \geq 0$ for all j. An agent i participating in a VCG mechanism has utility

$$\begin{aligned} u_i &= v_i(x^*) + \sum_{j \neq i} b_j(x^*) - \max_{x \in \mathcal{X}} \sum_{j \neq i} b_j(x) \\ &= b_i(x^*) + \sum_{j \neq i} b_j(x^*) - \max_{x \in \mathcal{X}} \sum_{j \neq i} b_j(x), \text{VCG is DSIC} \\ &= \max_{x \in \mathcal{X}} \sum_{j} b_j(x) - \max_{x \in \mathcal{X}} \sum_{j \neq i} b_j(x) \geq 0, \text{by definition of } x^* \end{aligned}$$

■

Corollary 2.4 *For single item allocation, a VCG mechanism with Clark pivots give us the second-price auction.*

Proof: Let v_i denote the true valuation of agent i for the item. Since this is a single item allocation, only one person receives the item and all others receive **no item.** If $x \in \mathcal{X}$ is the outcome where agent i receives the item, then $b_i(x)$ is agent i's bid for the item while $b_j(x) = 0$ for all $j \neq i$.

By the definition of VCG we pick the outcome x^* where

$$x^* = \arg\max_{x \in \mathcal{X}} \sum_{i=1}^{n} b_i(x)$$

But since only one agent gets the item, say agent a, we have that

$$\begin{aligned} x^* &= \arg\max_{x \in \mathcal{X}} [b_a(x) + 0 + 0...] \\ &= \arg\max_{x \in \mathcal{X}} b_a(x) \end{aligned}$$

Thus, x^* is the outcome where the bidder with the highest bid is picked. Let $b_a = A$.

By the Clark pivot in the VCG mechanism, the price paid by a is

$$\begin{aligned} p_a &= \max_{x \in \mathcal{X}} \sum_{j \neq a} b_j(x) - \sum_{j \neq a} b_j(x^*) \\ &= \max_{x \in \mathcal{X}} \sum_{j \neq a} b_j(x) \end{aligned}$$

since $b_j(x^*) = 0$ for all $j \neq a$.

Note that $x' = \arg\max_{x \in \mathcal{X}} \sum_{j \neq a} b_j(x)$ is the outcome of the allocation if a were removed from it. Thus x' is the outcome where the bidder, say b, with the second highest bid is picked. Let $b_b = B$. Then,

$$p_a = \max_{x \in \mathcal{X}} \sum_{j \neq a} b_j(x) = B$$

For all other agents $i \neq a$ we have

$$p_i = \max_{x \in \mathcal{X}} \sum_{j \neq i} b_j(x) - \sum_{j \neq i} b_j(x^*)$$
$$= \max_{x \in \mathcal{X}} [A + 0 + 0...] - [A + 0 + 0...] = 0$$

Thus VCG with Clark pivots applied to the single item allocation problem gives us the second-price auction. ■

Remark 2.4 *VCG might not be computationally efficient due to the computation of argmax. Example: Combinatorial auctions.*

Corollary 2.5 *The mechanism described in* ***Example 2.1*** *is DSIC.*

Proof: Assume the mechanism described in the example is VCG (except for the payment rule). We derive a payment rule as specified by the VCG definition and show that it indeed corresponds to the payment rule given in Example 2.1. In the case where x^* is the outcome where G is not built, it is trivially the case that $p_i = 0$. Let x^* be the outcome where the good G is built. By definition,

$$\begin{aligned} p_i &= h_i(b_{-i}) - \sum_{j \neq i} b_j(x^*) \\ &= h_i(v_{-i}) - \sum_{j \neq i} v_j(x^*) \text{ VCG is DSIC} \end{aligned}$$

Since $v_j(x^*) = v_j$ is constant and h is arbitrary, we let h be some constant function, giving us

$$p_i = K - \sum_{j \neq i} v_j$$

As we are considering the case where G is built, it follows that the payments sum to C:

$$\sum_{i=1}^{n} p_i = C = nK - (n-1) \sum_{j=1}^{n} v_j$$

By the problem specification, G should be built if $\sum_{j=1}^{n} v_j \geq C$. Hence,

$$\begin{aligned} (n-1) \sum_{j=1}^{n} v_j \geq (n-1)C &\implies nK - (n-1) \sum_{j=1}^{n} v_j \leq nK - (n-1)C \\ &\implies C \leq nK - (n-1)C \\ &\implies K \leq C \end{aligned}$$

Since K is arbitrary with the constraint given above, we set $K = C$, which gives us $p_i = C - \sum_{j \neq i} v_j$. Since payments should be non-negative, we take the maximum, setting $p_i = \max(0, C - \sum_{j \neq i} v_j)$. This can be done as we can set h accordingly to achieve this result. So, the mechanism with the payment rule specified in Example 2.1 is VCG and thus DSIC. ■

References

[1] Tim Roughgarden *Lectures Notes on Algorithmic Game Theory*. Version: July 28, 2014.