COMP/MATH 533: Algorithmic Game Theory	September 24, 2014
Lecture 7	
Lecturer: Yang Cai	*Scribe: Muhammad Samir Khan and Jonathan Talmi*

NOTE: The content of these notes has not been formally reviewed by the lecturer. It is recommended that they are read critically.

1 Criticism of Myerson's Optimal Auction

Last time we criticised the Myerson's Auction. In Myerson's Auction, we get maximum revenue with DSIC. However, it relies on knowing the probability distribution of the bidders. If there are many bidders and a lot of past data (market is "rich") then assuming we know the underlying distribution is reasonable. But, sometimes there is not enough data to learn the distribution (market is "thin"). Another problem with this mechanism is that we may end up with weird Allocation and Payment Rules. Last time we saw such an example with 2 bidders following $U[0,1]$ and $U[0,100]$. This property is undesirable because in practice the auctioneers might have to explain the allocation and pricing to the bidders (e.g. if Google runs an auction). Last week we saw a simple $\frac{1}{2}$-approximation (using Prophet Inequality) that has a very simple form (Vickrey with reserve) avoiding the weird allocation and pricing rule. This time we will address the other issue of Myerson's auction, its heavy dependence on the distribution.

2 Prior-Independent Auctions

Can we design auctions that know nothing about the probability distribution of the bidders and do (almost) as good as the auctions tailored for specific distributions? This is an active research area, called prior-independent mechanism design. While results are known for single-dimensional case, not much is known for the multi-dimensional scenario.

The classical result in this field is the *Bulow-Klemperer Theorem* below. This was proven by two economists in the 90's, and has later inspired computer scientists to study prior-independent auctions.

Theorem 1 (Bulow-Lemperer '96)**.** *For any regular distribution F and integer n*

$$\mathbb{E}_{v1,\ldots v_{n+1}}[REV(Vickrey)] \geq \mathbb{E}_{v_1,\ldots,v_n}[REV(Myerson)]$$

Remark 1. *Note that Vickrey's auction is prior-independent as we are not using any information about the underlying probability distributions. The theorem says that, for one item, Vickrey's Auction is very close to Myerson's Optimal Revenue Auction (in fact $\frac{n-1}{n}$-close). In practice, sometimes it is better to find more competition than to find the right auction format (but, e.g. for Google's sponsored search auction, almost everyone is in the auction so we can not really get more competition).*

Proof: In order to prove the theorem, we introduce an intermediary auction. Consider an auction M with $n+1$ bidders defined as follows:

1. Run Myerson's Auction on the first n bidders.
2. If the item is unallocated, give it to the last bidder for free.

Note that this auction is introduced as a proof technique, we will not really apply this auction anywhere, so don't try to optimize it. In fact, we just want to use the following two properties of this auction 1) it has the same revenue as Myerson's Auction with n bidders; 2) the item is always allocated (hence giving it out for free if the first n bidders are unsuccessful in securing it).

Notice that the first n bidders are exactly the same as in Myerson's Auction and the last bidder's bid does not matter. Thus, M is a ***DSIC*** mechanism, so its revenue is no greater than the revenue of the optimal DSIC mechanism that always allocates the item. Now, let's ask the following question: among all DSIC mechanisms for $n+1$ bidders that always allocates the item, which one optimizes the revenue. The answer is simple – it must be the one that always allocates the item to the bidder who has the highest virtual value. Since the distributions are i.i.d. and regular, the bidder with the highest virtual value is exactly the bidder with the highest bid. Thus, the auction is exactly the Vickrey auction.

$$\implies \mathbb{E}_{v_1,\ldots v_{n+1}}[\text{REV(Vickrey)}] \geq \mathbb{E}_{v_1,\ldots v_{n+1}}[\text{REV}(M)] = \mathbb{E}_{v_1,\ldots,v_n}[\text{REV(Myerson)}]$$

$\square$

3 General Mechanism Design Problem (Multi-Dimensional)

So far we have focused on single-dimensional environment, where we typically have one item or one item with multiple copies. However, in practice bidders usually have different values for different items. For example, for an auction of paintings, every bidder will have its own unique value for each painting. This is when we enter the multi-dimensional environment where many things that we have learnt no longer hold. The multi-dimensional environment has:

- $\boldsymbol{n}$ strategic bidders/participants/agents
- a set of possible outcomes $\boldsymbol{\Omega}$
- a private value $\boldsymbol{v_i(\omega)}$ for each bidder $\boldsymbol{i}$ and each $\boldsymbol{\omega \in \Omega}$ (abstract and could be large)

As an example of the multi-dimensional environment, consider the single-item auction. In the single-dimensional setting there are $n+1$ outcomes in Ω (the $(n+1)$th outcome is to not give the item to anybody). Bidder i has positive value for the outcome in which he wins the item and value 0 for the other n outcomes. We can make this scenario multi-dimensional. There are still $n+1$ outcomes in Ω. However, bidder i now has different values (besides the outcome where i wins) for the other n outcomes. For example, imagine that the item is a startup company who has a lot of valuable patents. The bidder will not want the competitor to get the patents and, in case he can not get the item, he would prefer someone in a different business to get the company.

As an example for Ω, consider two bidders, 1 and 2, and two items, A and B. Then $\Omega = \{1 : A \text{ and } 2 : \emptyset, 1 : A \text{ and } 2 : B, 1 : B \text{ and } 2 : \emptyset, 1 : B \text{ and } 2 : A, 1 : AB \text{ and } 2 : \emptyset, 1 : \emptyset \text{ and } 2 : \emptyset, 1 : \emptyset \text{ and } 2 : A, 1 : \emptyset \text{ and } 2 : B, 1 : \emptyset \text{ and } 2 : AB\}$

Optimizing the social welfare is the same as solving the following optimization problem:

$$\boldsymbol{\omega^* := \mathbf{argmax}_{\omega} \sum_{i} v_i(\omega)}$$

Note that Ω can now be exponential in size. In this setting, how do we design a DSIC mechanism that optimizes social welfare? We can take the same two-step approach as the single-dimensional environment. However, the bid is no longer a single number. Instead each bid $\boldsymbol{b_i}$ is a vector indexed by Ω. For the allocation rule, assume $\boldsymbol{b_i}$'s are the true values and solve the optimization problem above to maximize the social welfare (the optimization problem may not be tractable but more on that later). In single-dimensional setting Myerson's Lemma would give the payment rule. However, Myerson's Lemma does not apply to the multi-dimensional setting (how do we even define monotone allocation rule when bids are vectors?). It is not immediately clear if there is even a payment rule that makes the mechanism DSIC.

3.1 Vickrey-Clarke-Groves (VCG) Mechanism

Theorem 2. *[The Vickrey-Clarke-Groves (VCG) Mechanism] In every general mechanism design environment, there is a DSIC mechanism that maximizes the social welfare. In particular the allocation rule is*

$$x(b) = argmax_\omega \sum_i b_i(\omega) \tag{1}$$

and the payment rule is

$$p_i(b) = \max_\omega \sum_{i \neq j} b_j(\omega) - \sum_{i \neq j} b_j(\omega^*) \tag{2}$$

were $\omega^* = argmax_\omega \sum_i b_i(\omega)$ is the outcome chosen in (1)

We try and explain what the payment rule means. $\max_\omega \sum_{j \neq i} b_j(\omega)$ is the optimal social welfare when bidder i is not participating. ω^* is the optimal social welfare outcome and $\sum_{j \neq i} b_j(\omega^*)$ is the welfare from all agents except i. So the difference $\max_\omega \sum_{i \neq j} b_j(\omega) - \sum_{i \neq j} b_j(\omega^*)$ can be viewed as the "*welfare loss*" inflicted on the other $n-1$ bidders by i's presence (also known as "externality" in Economics). For example, for the single-item single-dimensional setting, if i is the winner then $\max_\omega \sum_{j \neq i} b_j(\omega)$ is the second largest bid and $\sum_{j \neq i} b_j(\omega^*) = 0$. In other words, it is exactly the second-price.

Proof: We argue that for every bidder, i, and for every bid vector (except i's bid), b_{-i}, setting $b_i = v_i$ is the optimal choice.

Let $\omega^* := \text{argmax}_{\omega \in \Omega} \sum_i b_i(\omega)$. The utility for bidder i for the bid vector b is:

$$u_i(b) = v_i(\omega^*) - [\text{max}_\omega \sum_{j \neq i} b_j(\omega) - \sum_{j \neq i} b_j(\omega^*)]$$

$$= v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) - \text{max}_\omega \sum_{j \neq i} b_j(\omega)$$

Bidder i has no control on $\text{max}_\omega \sum_{j \neq i} b_j(\omega)$ so it can only optimize $v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*)$. Now imagine if bidder i is given the power to choose any ω for ω^*, then it will choose $\text{argmax}_\omega v_i(\omega) + \sum_{j \neq i} b_j(\omega)$. This is clearly the best outcome i could hope for. The surprising fact is that although i does not have the power to choose ω^*, by setting $b_i = v_i$ the allocation rule (1) will pick $\omega^* \in \text{argmax}_\omega v_i(\omega) + \sum_{j \neq i} b_j(\omega)$. $\square$

We have shown a DSIC mechanism that optimizes social welfare for any mechanism design problem. However, if Ω is really large then the problem of finding ω^* can be intractable. For example:

- We have m items, n bidders, each bidder wants only one item. This can be done in polynomial time (Ask yourself why.).
- We have m items, n bidders, each bidder is single-minded (only likes a particular set of items). This is **NP**-hard.
- We have m items, n bidders, each bidder can take any set of items. This is obviously also **NP**-hard, and probably harder than the previous one.

We might hope to find approximation algorithms for the allocation rule in polytime but the mechanism may no longer be DSIC. The problem to find a computationally efficient approximation algorithm that is DSIC is a very active research area.