CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

Allan Borodin (instructor)
Tyrone Strangway and Young Wu (TAs)

October 24, 2016

Lecture 12

- Announcements
  - Office hours: Tuesdays 3:30-4:30 SF 2303B
    or schedule meeting; or drop by.
  - I have posted the three questions for Assignment 2 which is due this Friday, October 28. This is now the entire assignment.
  - Assignment 1 has been graded. Some statistics:
    average 73.9%, median 78.9%, 6 below 50%
    One person did not submit via Markus.
- Todays agenda
  - Concluding comments on Chapter 14
  - Chapter 15
    - Chapter 15: The single parameter win-lose setting with respect to social welfare/surplus
    - VCG in the single parameter individually rational setting.
    - Some examples
    - sponsored search; VCG vs GSP

A concluding comment of Chapter 14

In example 14.11.1, we see that the situation for trying to optimally sell two items to one Bayesian buyer becomes much more complex than when optimally selling one item (i.e. by Myerson's optimal auction).

A concluding comment of Chapter 14

In example 14.11.1, we see that the situation for trying to optimally sell two items to one Bayesian buyer becomes much more complex than when optimally selling one item (i.e. by Myerson's optimal auction).

The setting is as simple as one can imagine for two items. Namely, the buyer's valuation function is additive; i.e. if his value is $v(1)$ for item 1 and $v(2)$ for item 2, then if he is allocated both items, his valuation is $v(1) + v(2)$. Moreover, the prior distribution for each item is the same.

A concluding comment of Chapter 14

In example 14.11.1, we see that the situation for trying to optimally sell two items to one Bayesian buyer becomes much more complex than when optimally selling one item (i.e. by Myerson's optimal auction).

The setting is as simple as one can imagine for two items. Namely, the buyer's valuation function is additive; i.e. if his value is $v(1)$ for item 1 and $v(2)$ for item 2, then if he is allocated both items, his valuation is $v(1) + v(2)$. Moreover, the prior distribution for each item is the same.

Suppose that $Prob[v(j) = 1] = Prob[v(j) = 2] = \frac{1}{2}$ for $j = 1, 2$.

A concluding comment of Chapter 14

In example 14.11.1, we see that the situation for trying to optimally sell two items to one Bayesian buyer becomes much more complex than when optimally selling one item (i.e. by Myerson's optimal auction).

The setting is as simple as one can imagine for two items. Namely, the buyer's valuation function is additive; i.e. if his value is $v(1)$ for item 1 and $v(2)$ for item 2, then if he is allocated both items, his valuation is $v(1) + v(2)$. Moreover, the prior distribution for each item is the same.

Suppose that $Prob[v(j) = 1] = Prob[v(j) = 2] = \frac{1}{2}$ for $j = 1, 2$.
Then if each item is sold separately, the expected revenue is 1 from each item and hence total expected revenue 2 for both items.

A concluding comment of Chapter 14

In example 14.11.1, we see that the situation for trying to optimally sell two items to one Bayesian buyer becomes much more complex than when optimally selling one item (i.e. by Myerson's optimal auction).

The setting is as simple as one can imagine for two items. Namely, the buyer's valuation function is additive; i.e. if his value is $v(1)$ for item 1 and $v(2)$ for item 2, then if he is allocated both items, his valuation is $v(1) + v(2)$. Moreover, the prior distribution for each item is the same.

Suppose that $Prob[v(j) = 1] = Prob[v(j) = 2] = \frac{1}{2}$ for $j = 1, 2$.
Then if each item is sold separately, the expected revenue is 1 from each item and hence total expected revenue 2 for both items.

Alternatively, if the seller just offers a bundle price of 3 for both items, then the expected revenue is $3 \cdot prob[V(1) + v(2) \geq 3] = 3 \cdot \frac{3}{4} = 2.25$.

A concluding comment of Chapter 14

In example 14.11.1, we see that the situation for trying to optimally sell two items to one Bayesian buyer becomes much more complex than when optimally selling one item (i.e. by Myerson's optimal auction).

The setting is as simple as one can imagine for two items. Namely, the buyer's valuation function is additive; i.e. if his value is $v(1)$ for item 1 and $v(2)$ for item 2, then if he is allocated both items, his valuation is $v(1) + v(2)$. Moreover, the prior distribution for each item is the same.

Suppose that $Prob[v(j) = 1] = Prob[v(j) = 2] = \frac{1}{2}$ for $j = 1, 2$.
Then if each item is sold separately, the expected revenue is 1 from each item and hence total expected revenue 2 for both items.

Alternatively, if the seller just offers a bundle price of 3 for both items, then the expected revenue is $3 \cdot prob[V(1) + v(2) \geq 3] = 3 \cdot \frac{3}{4} = 2.25$.

In other valuations, it is better not to bundle and in some valuations it may be best to offer the buyer a choice of individual prices or a bundle.

Selling multiple items to an additive buyer: a little good news

There are now a few results within the spirit of "simple (and approximate) vs optimal auctions". This is very much a topic within the field of algorithmic game theory (AGT) given our interest in computationally efficiency and conceptual simplicity

For those interested I suggest looking at the following:

1. Hart and Nisan (2014) show that if two items are distributed independently, then selling them separately obtains at least a fraction $\frac{1}{2}$ of the optimal revenue. Moreover, if the items are i.i.d., then pricing individually obtains at least a fraction $\frac{e}{e-1}$ of the optimal revenue.
2. Hart and Nisan also show that for an arbitrary number of items, neither pricing separately or as a grand bundle can guarantee a constant fraction of the optimal revenue.
3. However, in contrast, Babaioff et al (2014) show that the maximum of separate priceis and grand bundle pricing guarantees a constant fraction independent of the number of items.

Begin chapter 15: Social welfare for auctions in the single parameter setting

In Lecture 9, for definteness, we described a general auction type setting where we postulated a set of items, and buyers and sellers for those items. As I remarked then, the text (and other texts such as the AGT book edited by Nisan, Roughgarden, Tardos and Vazirani) do not mention "items" but rather frame the outcomes in terms of feasible allocations. We'll adopt that viewpoint and terminology now.

Begin chapter 15: Social welfare for auctions in the single parameter setting

In Lecture 9, for defintеness, we described a general auction type setting where we postulated a set of items, and buyers and sellers for those items. As I remarked then, the text (and other texts such as the AGT book edited by Nisan, Roughgarden, Tardos and Vazirani) do not mention "items" but rather frame the outcomes in terms of feasible allocations. We'll adopt that viewpoint and terminology now.

What we are assuming throughout our discussions is *quasi linear utility*, namely that an agents utility $u_i(a)$ for an outcome a having value $v_i(a)$ is $u_i(a) = v_i(a) + m$ where m is some quantity of (say) money. In our auctions thus far, m is negative as it represents a payment from the buyers. In a *procurement auction* (example 15.1.2 and section 15.4.2), m is positive representing a payment to the agents who are providers of a service which has some private intrinsic cost (i.e. the "value" is negative).

The setting for Chapter 15

Chapetr 14 was concerned with the *revenue* of a seller/mechanism and we were mainly focused on the sale of a single item in the Bayesian setting where the value v_i of the item was drawn from a prior distribution.

The setting for Chapter 15

Chapetr 14 was concerned with the *revenue* of a seller/mechanism and we were mainly focused on the sale of a single item in the Bayesian setting where the value v_i of the item was drawn from a prior distribution.

In contrast, for the most part, the goal of the mechanism in Chapter 15 is to maximize social welfare/surplus. Moreover, we are not in the Bayesian setting (i.e. there do no exist prior distributions for the agents). The chapter concludes with a section on worst case revenue maximization (i.e. without prior distributions).

The setting for Chapter 15

Chapetr 14 was concerned with the *revenue* of a seller/mechanism and we were mainly focused on the sale of a single item in the Bayesian setting where the value v_i of the item was drawn from a prior distribution.

In contrast, for the most part, the goal of the mechanism in Chapter 15 is to maximize social welfare/surplus. Moreover, we are not in the Bayesian setting (i.e. there do no exist prior distributions for the agents). The chapter concludes with a section on worst case revenue maximization (i.e. without prior distributions).

The single paramter setting for Chapter 15 (Definition 15.2.1)

- There is a set U of agents, and a set $\mathcal{L} \subset 2^U$ of *feasible allocations*.
- We say that agent i is a "winner" in a feasible allocation $L \in \mathcal{L}$ if $i \in L$. Each agent i has a single private value v_i which is the value for agent i if i is a winner.

Chapter 16 considers the more general setting where an agent can have different values for different allocations.

Social welfare: deterministic and randomized mechanisms

The social welfare of an allocation L is defined as the sum of the valuations of the winning agents.

Social welfare: deterministic and randomized mechanisms

The social welfare of an allocation L is defined as the sum of the valuations of the winning agents.

We are considering sealed bid (direct) auctions where each agent i offers a bid b_i and the mechanism uses an allocation rule to determine the winners and the payment rule for each (winning) agent. We will assume that the payment of losing agents is 0. We will also assume a finite set U of n agents.

Social welfare: deterministic and randomized mechanisms

The social welfare of an allocation L is defined as the sum of the valuations of the winning agents.

We are considering sealed bid (direct) auctions where each agent i offers a bid b_i and the mechanism uses an allocation rule to determine the winners and the payment rule for each (winning) agent. We will assume that the payment of losing agents is 0. We will also assume a finite set U of n agents.

For a deterministic mechanism, an *allocation rule* is a mapping $\alpha : \mathbf{b} \rightarrow \mathcal{L}$. Equivalently, $\alpha : \mathbf{b} \rightarrow \{0,1\}^n : \alpha_i = \alpha(\mathbf{b})_i = 1$ iff agent i is a winner.

Social welfare: deterministic and randomized mechanisms

The social welfare of an allocation L is defined as the sum of the valuations of the winning agents.

We are considering sealed bid (direct) auctions where each agent i offers a bid b_i and the mechanism uses an allocation rule to determine the winners and the payment rule for each (winning) agent. We will assume that the payment of losing agents is 0. We will also assume a finite set U of n agents.

For a deterministic mechanism, an *allocation rule* is a mapping $\alpha : \mathbf{b} \rightarrow \mathcal{L}$. Equivalently, $\alpha : \mathbf{b} \rightarrow \{0,1\}^n : \alpha_i = \alpha(\mathbf{b})_i = 1$ iff agent i is a winner.

The payment rule $p : \mathbf{b} \rightarrow \mathbb{R}^n$ determines the payment for each winning agent where we assume the payment $p_i = p_i(\mathbf{b})_i = 0$ if agent i is not a winner.
The utiltiy of agent of agent i with value v_i when bidding b_i is

$$u_i(\mathbf{b}|v_i) = v_i\alpha_i(\mathbf{b}) - p_i(\mathbf{b})$$

Randomized mechanisms

Although the mechanism examples in this chapter are deterministic, mechanisms can be randomized and that is how the text initially defines the relevant concepts.

Randomized mechanisms

Although the mechanism examples in this chapter are deterministic, mechanisms can be randomized and that is how the text initially defines the relevant concepts.

The discussion in the text is not clear but lets say that the randomization is just in terms of the feasible allocation that the mechanism produces. This then means that the allocation rule is a vector of probabiities; that is $\alpha(\mathbf{b}) = (\alpha_1(\mathbf{b}), \ldots, \alpha_n(\mathbf{b}))$ where $\alpha_i(\mathbf{b})$ is the probability that agent i is a winner.

Randomized mechanisms

Although the mechanism examples in this chapter are deterministic, mechanisms can be randomized and that is how the text initially defines the relevant concepts.

The discussion in the text is not clear but lets say that the randomization is just in terms of the feasible allocation that the mechanism produces. This then means that the allocation rule is a vector of probabiities; that is $\alpha(\mathbf{b}) = (\alpha_1(\mathbf{b}), \ldots, \alpha_n(\mathbf{b}))$ where $\alpha_i(\mathbf{b})$ is the probability that agent i is a winner.

The utility $u_i(\mathbf{b}|v_i) = v_i\alpha_i(\mathbf{b}) - p_i(\mathbf{b})$ is then the expectation over the randomness of the mechanism.

Randomized mechanisms

Although the mechanism examples in this chapter are deterministic, mechanisms can be randomized and that is how the text initially defines the relevant concepts.

The discussion in the text is not clear but lets say that the randomization is just in terms of the feasible allocation that the mechanism produces. This then means that the allocation rule is a vector of probabiities; that is $\alpha(\mathbf{b}) = (\alpha_1(\mathbf{b}), \ldots, \alpha_n(\mathbf{b}))$ where $\alpha_i(\mathbf{b})$ is the probability that agent i is a winner.

The utility $u_i(\mathbf{b}|v_i) = v_i\alpha_i(\mathbf{b}) - p_i(\mathbf{b})$ is then the expectation over the randomness of the mechanism.

For randomized mechanisms, we will also have to distinguish between worst case truthfulness vs truthfulness in the sense of the expectation.

Some examples

- A combinatorial auction CA consists of a set of items M where each agent i has a value $v_i(S)$ for differents desired sets $S \subseteq M$. In a sinlge parameter CA, the value $v_i = v_i(S)$ is the same for each desired set and we assume that the mechanism knows what particular sets agent i desires. $(S_1, \ldots, S_n)$ with $S_i \subseteq M$ is a feasible allocation if $S_i \cap S_j = \varnothing$ for all $i \neq j$. Here the winners are those agents i such that $S_i \neq \varnothing$.

Some examples

- A combinatorial auction CA consists of a set of items M where each agent i has a value $v_i(S)$ for differents desired sets $S \subseteq M$. In a sinlge parameter CA, the value $v_i = v_i(S)$ is the same for each desired set and we assume that the mechanism knows what particular sets agent i desires. $(S_1, \ldots, S_n)$ with $S_i \subseteq M$ is a feasible allocation if $S_i \cap S_j = \emptyset$ for all $i \neq j$. Here the winners are those agents i such that $S_i \neq \emptyset$.

  The spectrum auction in Example 15.1.1 would be such an example where the items might be wavelengths (for transmitting) and agents want known bands (i.e. intervals) of the spectrum. These intervals cannot intersect in a feasible allocation. .

Some examples

- A combinatorial auction CA consists of a set of items M where each agent i has a value $v_i(S)$ for differents desired sets $S \subseteq M$. In a sinlge parameter CA, the value $v_i = v_i(S)$ is the same for each desired set and we assume that the mechanism knows what particular sets agent i desires. $(S_1, \ldots, S_n)$ with $S_i \subseteq M$ is a feasible allocation if $S_i \cap S_j = \emptyset$ for all $i \neq j$. Here the winners are those agents i such that $S_i \neq \emptyset$.

  The spectrum auction in Example 15.1.1 would be such an example where the items might be wavelengths (for transmitting) and agents want known bands (i.e. intervals) of the spectrum. These intervals cannot intersect in a feasible allocation. .
- The feasible allocations in a single item auction (as in chapter 14) are those with just one winner.

More examples

- In the shared communcation channel example in Example 15.1.3, the agents have known bandwidth requirements w_i and a feasible allocation is one in which $\sum_{i:i \text{ is a winner}} w_i \leq C$ for some known capacity C.

More examples

- In the shared communcation channel example in Example 15.1.3, the agents have known bandwidth requirements w_i and a feasible allocation is one in which $\sum_{i:i \text{ is a winner}} w_i \leq C$ for some known capacity C.

These examples are all different in terms of the underlying combinatorial allocation problem.

- As we mentioned before, the allocation problem for a general CA is the set packing problem, which is NP-hard to approximate to within a factor $\min(n, m^{\frac{1}{2}-\epsilon})$ where $m = |M|$.
- For the special case of the spectrum auction, this is the interval selection problem which can be optimally solved (for the single-minded case) by an efficient (greedy) algorithm.
- The single item case is a rather trivial allocation problem.
- The allocation problem for the shared communication channel is the knapsack problem which is NP-hard but can be approximately solved to within any factor $(1 - \epsilon)$ by a computationally efficient (dynamic programming) algorithm.

The procurement auction in section 15.4.2

In a procurement auction, the mechanism pays the agents for some service or product.

The procurement auction in section 15.4.2

In a procurement auction, the mechanism pays the agents for some service or product.

In the spanning tree auction problem, a service provider (such as Netflix) needs to build a spanning tree by purchasing edges in a network from different agents. Assume each edge e_i is owned by a different agent i who has some private cost if that edge is used by the provider.

The procurement auction in section 15.4.2

In a procurement auction, the mechanism pays the agents for some service or product.

In the spanning tree auction problem, a service provider (such as Netflix) needs to build a spanning tree by purchasing edges in a network from different agents. Assume each edge e_i is owned by a different agent i who has some private cost if that edge is used by the provider.

The feasible solutions are spanning trees and the mechanism is trying to minimize the cost of the spanning tree it computes. Each agent i is trying to maximize the price p_i it receives minus its (private) cost. This becomes a social welfare maximization problem when setting $v_i = -c_i$.

The truthful VCG mechanism

THe Vickery Clarke Groves (VCG) mechanism is a deterministic mechanism that is truthful (ex-post incentive compatible) and individually rational. In Chapter 16, it is stated and proven to be truthful for the general case (i.e. not restricted to single parameter auctions) whenever we have a quasi-linear utiltiy function $u_i = v_i - p_i$.

The truthful VCG mechanism

THe Vickery Clarke Groves (VCG) mechanism is a deterministic mechanism that is truthful (ex-post incentive compatible) and individually rational. In Chapter 16, it is stated and proven to be truthful for the general case (i.e. not restricted to single parameter auctions) whenever we have a quasi-linear utiltiy function $u_i = v_i - p_i$.

The general proof of truthfulness is given in Theorem 16.2.6. The proof is remarkably simple given the generality of this result.

The truthful VCG mechanism

THe Vickery Clarke Groves (VCG) mechanism is a deterministic mechanism that is truthful (ex-post incentive compatible) and individually rational. In Chapter 16, it is stated and proven to be truthful for the general case (i.e. not restricted to single parameter auctions) whenever we have a quasi-linear utiltiy function $u_i = v_i - p_i$.

The general proof of truthfulness is given in Theorem 16.2.6. The proof is remarkably simple given the generality of this result.

As the name might suggest, the Vickery auction is VCG applied to the special case of a single item auction. It is also the special case of the Vickery auction for the mutli item case when there are k copies of the same item and the winners are the agents with the top k values and the price is the $k+1$-st value. (Do not just quote result for question 2 of assignment.)

The truthful VCG mechanism

THe Vickery Clarke Groves (VCG) mechanism is a deterministic mechanism that is truthful (ex-post incentive compatible) and individually rational. In Chapter 16, it is stated and proven to be truthful for the general case (i.e. not restricted to single parameter auctions) whenever we have a quasi-linear utiltiy function $u_i = v_i - p_i$.

The general proof of truthfulness is given in Theorem 16.2.6. The proof is remarkably simple given the generality of this result.

As the name might suggest, the Vickery auction is VCG applied to the special case of a single item auction. It is also the special case of the Vickery auction for the mutli item case when there are k copies of the same item and the winners are the agents with the top k values and the price is the $k+1$-st value. (Do not just quote result for question 2 of assignment.)

Since we are considering social welfare (and not revenue), we do not have to consider reserve prices but as we have mentioned, reserve prices do not cause a problem for truthfulness.

The VCG mechanism in a figure

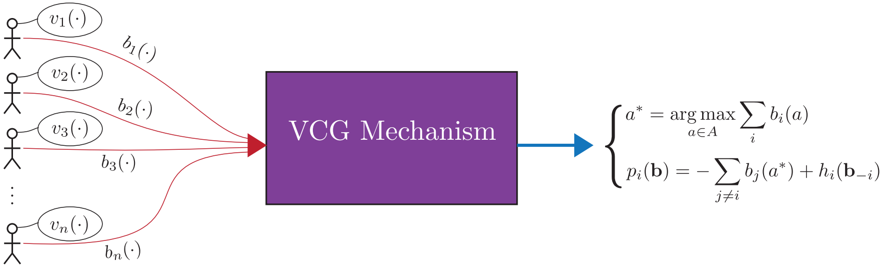

Figure : The depiction of VCG for arbitrary function $h_i(\mathbf{b}_{-i})$.

Here A is the set of feasible allocations. VCG is truthful for any h_i that does not depend on b_i. To make VCG IR and all payments non-negative, we use the "Clarke pivot":

$$h_i(\mathbf{b}_{-i}) = \max_{a \in A} \sum_{j \neq i} b_j(a)$$

VCG in words

VCG is a deterministic sealed bid auction. Each agent i submits their bids $b_i(a_i)$ for each desired outcome a_i.

Note: In the single-minded case, the agent has just one desired outcome but (say) in the general CA, an agent could desire many different sets each having its own valuation.

VCG in words

VCG is a deterministic sealed bid auction. Each agent i submits their bids $b_i(a_i)$ for each desired outcome a_i.

Note: In the single-minded case, the agent has just one desired outcome but (say) in the general CA, an agent could desire many different sets each having its own valuation.

The mechanism computes an optimal feasible allocation a^* with regard to social welfare. Each agent is allocated whatever is determined by a^*. For the CA problem, an agent receives a desired subset or does not get allocated any set. (Ties can be broken arbitrarily.)

VCG in words

VCG is a deterministic sealed bid auction. Each agent i submits their bids $b_i(a_i)$ for each desired outcome a_i.

Note: In the single-minded case, the agent has just one desired outcome but (say) in the general CA, an agent could desire many different sets each having its own valuation.

The mechanism computes an optimal feasible allocation a^* with regard to social welfare. Each agent is allocated whatever is determined by a^*. For the CA problem, an agent receives a desired subset or does not get allocated any set. (Ties can be broken arbitrarily.)

Given this optimal allocation a^*, the payment $p_i(\mathbf{b})$ for agent i is :

$$p_i(\mathbf{b}) = \max_{a \in A} \sum_{j \neq i} b_j(a) - \sum_{j \neq i} b_j(a^*)$$

That is, agent i is being charged his impact on the social welfare of the other agents.

VCG in words

VCG is a deterministic sealed bid auction. Each agent i submits their bids $b_i(a_i)$ for each desired outcome a_i.

Note: In the single-minded case, the agent has just one desired outcome but (say) in the general CA, an agent could desire many different sets each having its own valuation.

The mechanism computes an optimal feasible allocation a^* with regard to social welfare. Each agent is allocated whatever is determined by a^*. For the CA problem, an agent receives a desired subset or does not get allocated any set. (Ties can be broken arbitrarily.)

Given this optimal allocation a^*, the payment $p_i(\mathbf{b})$ for agent i is :

$$p_i(\mathbf{b}) = \max_{a \in A} \sum_{j \neq i} b_j(a) - \sum_{j \neq i} b_j(a^*)$$

That is, agent i is being charged his impact on the social welfare of the other agents. Be sure you understand this statement!