Algorithmic Game Theory

Mechanism Design: Multi Parameter Environments

Georgios Birmpas
birbas@diag.uniroma1.it

Based on slides by Alexandros Voudouris

General environments

- A set of n **agents**
- A finite set Ω of **outcomes**

General environments

- A set of n **agents**
- A finite set Ω of **outcomes**
- Each agent i has a private non-negative value $v_i(\omega)$ for every outcome $\omega \in \Omega$

General environments

- A set of n **agents**
- A finite set Ω of **outcomes**
- Each agent i has a private non-negative value $v_i(\omega)$ for every outcome $\omega \in \Omega$
- The social welfare of an outcome $\omega \in \Omega$ is $\sum_i v_i(\omega)$

General environments

- A set of n **agents**
- A finite set Ω of **outcomes**
- Each agent i has a private non-negative value $v_i(\omega)$ for every outcome $\omega \in \Omega$
- The social welfare of an outcome $\omega \in \Omega$ is $\sum_i v_i(\omega)$

- **Our goals:**
  - Incentivize the agents **to truthfully report their values**
  - Choose an outcome that **maximizes the social welfare**

Single-item auctions

- There are only $n + 1$ outcomes, corresponding to the number of possible winners (if any)

Single-item auctions

- There are only $n + 1$ outcomes, corresponding to the number of possible winners (if any)
- In the standard model, the value of each agent is 0 in all n outcomes in which she loses

Single-item auctions

- There are only $n + 1$ outcomes, corresponding to the number of possible winners (if any)
- In the standard model, the value of each agent is 0 in all n outcomes in which she loses
- This leaves only **one unknown parameter** per agent, her value for the outcome in which she wins

Single-item auctions

- There are only $n + 1$ outcomes, corresponding to the number of possible winners (if any)
- In the standard model, the value of each agent is 0 in all n outcomes in which she loses
- This leaves only **one unknown parameter** per agent, her value for the outcome in which she wins

- In general, the agents might have different values for the possible winners of the item

Combinatorial auctions

- Multiple indivisible items for sale

Combinatorial auctions

- Multiple indivisible items for sale
- The agents might have complex preferences over the possible item combinations

Combinatorial auctions

- Multiple indivisible items for sale
- The agents might have complex preferences over the possible item combinations
- For n agents and a set M of m items, the set of outcomes consists of all n-vectors $(X_1, \ldots, X_n)$ such that $\bigcup_i X_i \subseteq M$ and $X_i \cap X_j = \emptyset, \forall i \neq j$

Combinatorial auctions

- Multiple indivisible items for sale
- The agents might have complex preferences over the possible item combinations
- For n agents and a set M of m items, the set of outcomes consists of all n-vectors $(X_1, \dots, X_n)$ such that $\bigcup_i X_i \subseteq M$ and $X_i \cap X_j = \emptyset, \forall i \neq j$
  - There are $(n+1)^m$ different outcomes

Combinatorial auctions

- Multiple indivisible items for sale
- The agents might have complex preferences over the possible item combinations
- For n agents and a set M of m items, the set of outcomes consists of all n-vectors $(X_1, \dots, X_n)$ such that $\bigcup_i X_i \subseteq M$ and $X_i \cap X_j = \emptyset, \forall i \neq j$
  - There are $(n+1)^m$ different outcomes
- Each agent i has a private value $v_i(S)$ for every possible bundle $S \subseteq M$ of items

Combinatorial auctions

- Multiple indivisible items for sale
- The agents might have complex preferences over the possible item combinations
- For n agents and a set M of m items, the set of outcomes consists of all n-vectors $(X_1, \ldots, X_n)$ such that $\bigcup_i X_i \subseteq M$ and $X_i \cap X_j = \emptyset, \forall i \neq j$
  - There are $(n+1)^m$ different outcomes
- Each agent i has a private value $v_i(S)$ for every possible bundle $S \subseteq M$ of items
  - Each agent i has 2^m parameters

Valuation Functions

A function f is

- Submodular: $f(S \cup \{j\}) - f(S) \geq f(T \cup \{j\}) - f(T)$
  for any $S \subseteq T$, and $j \notin T$
- Supermodular: $f(T \cup \{j\}) - f(T) \geq f(S \cup \{j\}) - f(S)$
  for any $S \subseteq T$, and $j \notin T$
- Symmetric: $f(S)=f(T)$ when $|S|=|T|$
- Symmetric Submodular: Submodular and Symmetric
- Subadditive: $f(S \cup T) \leq f(S) + f(T)$, for any S, T

Symmetric Submodular ⊆ Submodular ⊆ Subadditive

VCG mechanisms

- A general solution for any environment

VCG mechanisms

- A general solution for any environment
- The VCG (Vickrey-Clarke-Groves) mechanisms implement (truthfully) the social welfare maximizing outcome

VCG mechanisms

- A general solution for any environment
- The VCG (Vickrey-Clarke-Groves) mechanisms implement (truthfully) the social welfare maximizing outcome

- **Allocation rule:** Maximize the social welfare according to the input

$$\boldsymbol{x}(\boldsymbol{b}) = \arg\max_{\omega \in \Omega} \sum_i b_i(\omega)$$

VCG mechanisms

- A general solution for any environment
- The VCG (Vickrey-Clarke-Groves) mechanisms implement (truthfully) the social welfare maximizing outcome

- **Allocation rule:** Maximize the social welfare according to the input

$$\boldsymbol{x}(\boldsymbol{b}) = \arg\max_{\omega \in \Omega} \sum_i b_i(\omega)$$

- **Payment rule:** For a set of functions $h_1, \dots, h_n$ such that h_i is independent of the bid of agent i,

$$p_i(\boldsymbol{b}) = h_i(\boldsymbol{b}_{-i}) - \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)$$

VCG mechanisms

Theorem

Every VCG mechanism is truthful and maximizes the social welfare

VCG mechanisms

Theorem
Every VCG mechanism is truthful and maximizes the social welfare

- The utility of agent i is

$$u_i(\boldsymbol{b}) = v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - p_i(\boldsymbol{b})$$

VCG mechanisms

Theorem
Every VCG mechanism is truthful and maximizes the social welfare

- The utility of agent i is

$$u_i(\boldsymbol{b}) = v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - p_i(\boldsymbol{b})$$

$$= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - \left(h_i(\boldsymbol{b}_{-i}) - \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)\right)$$

VCG mechanisms

Theorem
Every VCG mechanism is truthful and maximizes the social welfare

- The utility of agent i is

$$
\begin{aligned}
u_i(\boldsymbol{b}) &= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - p_i(\boldsymbol{b}) \\
&= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - \left(h_i(\boldsymbol{b}_{-i}) - \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)\right) \\
&= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big) - h_i(\boldsymbol{b}_{-i})
\end{aligned}
$$

VCG mechanisms

Theorem

Every VCG mechanism is truthful and maximizes the social welfare

- The utility of agent i is

$$u_i(\boldsymbol{b}) = v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - p_i(\boldsymbol{b})$$

$$= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - \left(h_i(\boldsymbol{b}_{-i}) - \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)\right)$$

$$= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big) - h_i(\boldsymbol{b}_{-i})$$

independent of b_i

VCG mechanisms

Theorem
Every VCG mechanism is truthful and maximizes the social welfare

- The utility of agent i is

$$
\begin{aligned}
u_i(\boldsymbol{b}) &= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - p_i(\boldsymbol{b}) \\
&= v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) - \left(h_i(\boldsymbol{b}_{-i}) - \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)\right) \\
&= \underbrace{v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)}_{\text{The social welfare according to the true value of agent } i \text{ and the bids of the other agents}} - h_i(\boldsymbol{b}_{-i})
\end{aligned}
$$

independent of b_i

VCG mechanisms

- Agent i cares about the welfare of all agents (based on the reported valuations) and aims to maximize the quantity

$$v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)$$

VCG mechanisms

- Agent i cares about the welfare of all agents (based on the reported valuations) and aims to maximize the quantity

$$v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)$$

- Since $\boldsymbol{x}(\boldsymbol{b})$ is such that

$$\boldsymbol{x}(\boldsymbol{b}) \in \arg\max_{\omega \in \Omega} \left\{ b_i(\omega) + \sum_{j \neq i} b_j(\omega) \right\}$$

the best response of agent i is to set $b_i = v_i$

VCG mechanisms

- Agent i cares about the welfare of all agents (based on the reported valuations) and aims to maximize the quantity

$$v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)$$

- Since $\boldsymbol{x}(\boldsymbol{b})$ is such that

$$\boldsymbol{x}(\boldsymbol{b}) \in \arg\max_{\omega \in \Omega} \left\{ b_i(\omega) + \sum_{j \neq i} b_j(\omega) \right\}$$

  the best response of agent i is to set $b_i = v_i$

- Therefore every agent i truthfully reports her true values

VCG mechanisms

- Agent i cares about the welfare of all agents (based on the reported valuations) and aims to maximize the quantity

$$v_i\big(\boldsymbol{x}(\boldsymbol{b})\big) + \sum_{j \neq i} b_j\big(\boldsymbol{x}(\boldsymbol{b})\big)$$

- Since $\boldsymbol{x}(\boldsymbol{b})$ is such that

$$\boldsymbol{x}(\boldsymbol{b}) \in \arg\max_{\omega \in \Omega} \left\{ b_i(\omega) + \sum_{j \neq i} b_j(\omega) \right\}$$

  the best response of agent i is to set $b_i = v_i$

- Therefore every agent i truthfully reports her true values
- The mechanism is designed so that the incentives of the agents are aligned with the goal of maximizing the social welfare $\square$

Clarke payments

- There are a lot of different VCG mechanisms, depending on how we choose the h-functions

Clarke payments

- There are a lot of different VCG mechanisms, depending on how we choose the h-functions
- We would like to have reasonable payment rules, that satisfy a couple of properties:
  - **Individual rationality:** Every agent has non-negative utility, and therefore incentive to participate
  - **No positive transfers:** The mechanism does not pay the agents, the agents pay the mechanism

Clarke payments

- Clarke payments: define

$$h_i(\boldsymbol{v}_{-i}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega)$$

Clarke payments

- Clarke payments: define

$$h_i(\boldsymbol{v}_{-i}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega)$$

and, hence

$$p_i(\boldsymbol{v}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) - \sum_{j \neq i} v_j(\boldsymbol{x}(\boldsymbol{v}))$$

Clarke payments

- Clarke payments: define

$$h_i(\boldsymbol{v}_{-i}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega)$$

and, hence

$$p_i(\boldsymbol{v}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) - \sum_{j \neq i} v_j(\boldsymbol{x}(\boldsymbol{v}))$$

- The payment of agent i is the difference between the maximum social welfare of the other agents when she does not participate, and the social welfare when she participates
- Agent i pays the loss in welfare due to her participation

Clarke payments

<u>Theorem</u>
A VCG mechanism with Clarke payments satisfies the properties of individual rationality and no positive transfers

Clarke payments

Theorem
A VCG mechanism with Clarke payments satisfies the properties of individual rationality and no positive transfers

- No positive transfers:

$$p_i(\boldsymbol{v}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) - \sum_{j \neq i} v_j\big(\boldsymbol{x}(\boldsymbol{v})\big) \geq 0$$

Clarke payments

Theorem
A VCG mechanism with Clarke payments satisfies the properties of individual rationality and no positive transfers

- No positive transfers:

$$p_i(\boldsymbol{v}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) - \sum_{j \neq i} v_j\big(\boldsymbol{x}(\boldsymbol{v})\big) \geq 0$$

- Individual rationality:

$$u_i(\boldsymbol{v}) = \sum_{j} v_j\big(\boldsymbol{x}(\boldsymbol{v})\big) - \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega)$$

Clarke payments

Theorem

A VCG mechanism with Clarke payments satisfies the properties of individual rationality and no positive transfers

- No positive transfers:

$$p_i(\boldsymbol{v}) = \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) - \sum_{j \neq i} v_j\big(\boldsymbol{x}(\boldsymbol{v})\big) \geq 0$$

- Individual rationality:

$$u_i(\boldsymbol{v}) = \sum_{j} v_j\big(\boldsymbol{x}(\boldsymbol{v})\big) - \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega)$$

$$= \max_{\omega \in \Omega} \sum_{j} v_j(\omega) - \max_{\omega \in \Omega} \sum_{j \neq i} v_j(\omega) \geq 0 \qquad \square$$

Drawbacks of VCG mechanisms

- **Preference elicitation:** VCG mechanisms demand from each agent to communicate her values for every possible outcome
  - Not practical in many situations: communicating 2^m parameters in the case of combinatorial auctions is impossible, even for small m

Drawbacks of VCG mechanisms

- **Preference elicitation:** VCG mechanisms demand from each agent to communicate her values for every possible outcome
  - Not practical in many situations: communicating 2^m parameters in the case of combinatorial auctions is impossible, even for small m

- **Social welfare maximization might be a hard problem**
- Knapsack auctions:
  - each agent i demands w_i items and has a private value v_i
  - the seller has a total amount of W items
  - Even though every agent has only one private parameter, maximizing the social welfare is equivalent to the Knapsack problem, which is NP-hard

- Exercise: Consider the following setting, we have n players and m items and we want to allocate the items to the players. By allocating we mean that the players will get the items without paying something (this is a problem without payments). Each player has a value for each of the items and these values might be different. Consider the following mechanism (Round-Robin Mechanism): The players are ordered in an arbitrary way and the mechanism runs in rounds following this ordering. In each round a player, when his order comes, chooses his most desirable item among the remaining ones i.e. the first player gets his most desirable item, the second player gets his most desirable item among the ones that remain and so on. So if we have n players {1,2,…, n} the mechanism runs as follows and with each agent getting his most desirable item among the ones that remain, 1-> 2-> 3->…-> n-> 1-> 2->…, until we run out of items. Is this mechanism truthful? Explain your answer.

- Exercise 4: Consider the previous problem once again but now under the following mechanism: The players are ordered in an arbitrary way and the mechanism runs following this ordering. In each round a player, when his order comes, chooses his most desirable item among the remaining ones and the last player gets all the remaining items. Thus, there is only one round this time and the last player is the only one that might get more than one items. Is this mechanism truthful? Explain your answer.