L11 Introduction to Mechanism Design

CS 280 Algorithmic Game Theory

Ioannis Panageas

Inspired by
by J. Hartline and T. Roughgarden notes

Warm-up: Single-item allocation

Definition (Single-item)**.** *The single-item allocation problem is given by*

- *a single indivisible item,*
- *n agents competing for the item,*
- *each agent i has an associated value/valuation v_i for getting the item.*

Warm-up: Single-item allocation

Definition (Single-item)**.** *The single-item allocation problem is given by*

- *a single indivisible item,*
- *n agents competing for the item,*
- *each agent i has an associated value/valuation v_i for getting the item.*

Goal: maximize the social surplus, i.e., the value of the agent who receives the item.

Warm-up: Single-item allocation

Definition (Single-item). *The single-item allocation problem is given by*

- *a single indivisible item,*
- *n agents competing for the item,*
- *each agent i has an associated value/valuation v_i for getting the item.*

Goal: maximize the social surplus, i.e., the value of the agent who receives the item.

Example:

- Ask the agents/bidders to report their values (**bids**), each agent reports b_i.
- Select the agent i^* with highest report ($i^* = \text{argmax}_i\, b_i$).
- Allocate item to i^*.

Warm-up: Single-item allocation

Definition (Single-item). *The single-item allocation problem is given by*

- *a single indivisible item,*
- *n agents competing for the item,*
- *each agent i has an associated value/valuation v_i for getting the item.*

Goal: maximize the social surplus, i.e., the value of the agent who receives the item.

Example:

- Ask the agents/bidders to report their values (**bids**), each agent reports b_i.
- Select the agent i^* with highest report ($i^* = \operatorname{argmax}_i b_i$).
- Allocate item to i^*.

If your $v_i = 1\$$, how would you play?

Warm-up: Single-item allocation

Definition (Single-item). *The single-item allocation problem is given by*

- *a single indivisible item,*
- *n agents competing for the item,*
- *each agent i has an associated value/valuation v_i for getting the item.*

Goal: maximize the social surplus, i.e., the value of the agent who receives the item.

Example:

- Ask the agents/bidders to report their values (**bids**), each agent reports b_i.
- Select the agent i^* with highest report ($i^* = \text{argmax}_i\ b_i$).
- Allocate item to i^*.

If your $v_i = 1\$$, how would you play?
You should always bid the highest number you can think of!

Warm-up: Single-item allocation

Definition (Single-item). *The single-item allocation problem is given by*

- *a single indivisible item,*
- *n agents competing for the item,*
- *each agent i has an associated value/valuation v_i for getting the item.*

Goal: maximize the social surplus, i.e., the value of the agent who receives the item.

Example:

- Ask the agents/bidders to report their values (**bids**), each agent reports b_i.
- Select the agent i^* with highest report ($i^* = \text{argmax}_i\, b_i$).
- Allocate item to i^*.

If your $v_i = 1\$$, how would you play?
You should always bid the highest number you can think of!
Outcome of mechanism is unpredictable, hard to reason about performance

Warm-up: Single-item allocation

Approach 1 (Lottery). *select a uniformly random agent, and allocate the item to that agent.*

Warm-up: Single-item allocation

Approach 1 (Lottery). *select a uniformly random agent, and allocate the item to that agent.*

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

Warm-up: Single-item allocation

Approach 1 (Lottery). *select a uniformly random agent, and allocate the item to that agent.*

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

Proof. Consider $v_1 = 1$ and $v_i = \epsilon$ for $i \geq 2$. Expected surplus is

$$\frac{1}{n}\left(1 + (n-1)\epsilon\right).$$

Thus approximation ratio is $\frac{n}{1+(n-1)\epsilon} \to n$ as $\epsilon \to 0$.

Warm-up: Single-item allocation

Approach 1 (Lottery). *select a uniformly random agent, and allocate the item to that agent.*

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

Proof. Consider $v_1 = 1$ and $v_i = \epsilon$ for $i \geq 2$. Expected surplus is

$$\frac{1}{n}\left(1 + (n-1)\epsilon\right).$$

Thus approximation ratio is $\frac{n}{1+(n-1)\epsilon} \to n$ as $\epsilon \to 0$.

Charge payments, proportionally to the agents' bids => **Discourage low-valued** agents from making high bids.

Single-item auctions

Approach 2 (First-price auction). *The first-price auction is defined:*

- *Agents report their bids* b_i.
- *Select agent* $i^* = \arg\max_i b_i$ *(highest bid).*
- i^* *gets the item and pays the amount of* b_{i^*}.

Single-item auctions

Approach 2 (First-price auction)**.** *The first-price auction is defined:*

- *Agents report their bids* b_i.
- *Select agent* $i^* = \arg\max_i b_i$ (*highest bid*).
- i^* *gets the item and pays the amount of* b_{i^*}.

First-price auctions are **hard to reason** about.
As a participant, it's hard to figure out **how to bid.**
As an auction designer, it's **hard to predict** what will happen.

Single-item auctions

Approach 3 (Second-price auction). *The second-price or Vickrey auction is defined:*

- *Agents report their bids* b_i.
- *Let agent* $i^* = \arg\max_i b_i$ *and let* j^* *be the agent with (second) highest bid.*
- i^* *gets the item and pays the amount of* b_{j^*}.

Single-item auctions

Approach 3 (Second-price auction). *The second-price or Vickrey auction is defined:*

- *Agents report their bids* b_i.
- *Let agent* $i^* = \arg\max_i b_i$ *and let* j^* *be the agent with (second) highest bid.*
- i^* *gets the item and pays the amount of* b_{j^*}.

Theorem (Vickrey is truthful). *In second price auctions, every bidder i has a dominant strategy: Set her bid* b_i *equal to her private* v_i *(report truthfully). Dominant means the utility of bidder i is maximized no matter what other bidders do.*

Single-item auctions

Approach 3 (Second-price auction). *The second-price or Vickrey auction is defined:*

- *Agents report their bids* b_i.
- *Let agent* $i^* = \arg\max_i b_i$ *and let* j^* *be the agent with (second) highest bid.*
- i^* *gets the item and pays the amount of* b_{j^*}.

Theorem (Vickrey is truthful). *In second price auctions, every bidder i has a dominant strategy: Set her bid* b_i *equal to her private* v_i *(report truthfully). Dominant means the utility of bidder i is maximized no matter what other bidders do.*

Remark: Utility of bidder i is $u_i := v_i - p_i$ **if he gets the item** and $u_i := 0$ **otherwise.**

Second price auctions

Proof. Fix an agent i and set $B = \max_{j \neq i} b_j$.

Consider the cases:

- If $b_i < B$ then i gets utility 0.
- If $b_i \geq B$ then i wins the item and $u_i = v_i - B$.

Second price auctions

Proof. Fix an agent i and set $B = \max_{j \neq i} b_j$.

Consider the cases:

- If $b_i < B$ then i gets utility 0.
- If $b_i \geq B$ then i wins the item and $u_i = v_i - B$.

Assume that $v_i < B$.

- If i reports truthfully, she gets utility 0 (did not win the item).
- Assume not and $b_i < v_i$ then the utility of i will still be 0.
- Assume not and $B > b_i > v_i$ then the utility will still be 0.
- Assume not and $b_i \geq B > v_i$ then the utility will be negative.

Second price auctions

Proof cont. Fix an agent i and set $B = \max_{j \neq i} b_j$.

Assume that $v_i \geq B$.

- If i reports truthfully, she gets utility $v_i - B \geq 0$ (won the item).
- Assume not and $b_i > u_i$ then the utility of i will still be $v_i - B \geq 0$.
- Assume not and $v_i > b_i \geq B$ then the utility will still be $v_i - B \geq 0$.
- Assume not and $v_i \geq B > b_i$ then the utility will be 0.

Second price auctions

Proof cont. Fix an agent i and set $B = \max_{j \neq i} b_j$.

Assume that $v_i \geq B$.

- If i reports truthfully, she gets utility $v_i - B \geq 0$ (won the item).
- Assume not and $b_i > u_i$ then the utility of i will still be $v_i - B \geq 0$.
- Assume not and $v_i > b_i \geq B$ then the utility will still be $v_i - B \geq 0$.
- Assume not and $v_i \geq B > b_i$ then the utility will be 0.

No matter what other bidders do, truthtelling is best strategy (**Dominant strategy**).

A general approach

An auction should satisfy following properties:

- Dominant strategy incentive compatible (DSIC), i.e., truthtelling is dominant strategy.

A general approach

An auction should satisfy following properties:

- Dominant strategy incentive compatible (DSIC), i.e., truthtelling is dominant strategy.
- (Auctioneer incentives) If bidders are truthful, auction miximizes surplus

$$\sum_{i=1}^{n} v_i x_i$$

where x_i is 1 if i wins the item and 0 otherwise.

A general approach

An auction should satisfy following properties:

- Dominant strategy incentive compatible (DSIC), i.e., truthtelling is dominant strategy.
- (Auctioneer incentives) If bidders are truthful, auction miximizes surplus

$$\sum_{i=1}^{n} v_i x_i$$

  where x_i is 1 if i wins the item and 0 otherwise.
- (Computational Efficiency) The auction can be implemented in polynomial time.

An Example

Problem: Consider a society of n citizens and public good G.

- Each agent has (private) valuation v_i for the good.
- Cost of building G is (publicly known) C.

- G should be built if $\sum_{i=1}^{n} v_i > C$.

An Example

Problem: Consider a society of n citizens and public good G.

- Each agent has (private) valuation v_i for the good.
- Cost of building G is (publicly known) C.

- G should be built if $\sum_{i=1}^{n} v_i > C$.

Goal: Design a mechanism that charges citizens in a way that G is built only

if $\sum_{i=1}^{n} v_i > C$.

An Example

Problem: Consider a society of n citizens and public good G.

- Each agent has (private) valuation v_i for the good.
- Cost of building G is (publicly known) C.

- G should be built if $\sum_{i=1}^{n} v_i > C.$

Goal: Design a mechanism that charges citizens in a way that G is built only

if $\sum_{i=1}^{n} v_i > C.$

Question: Allocating the cost C equally works? Why not?

An Example

Problem: Consider a society of n citizens and public good G.

- Each agent has (private) valuation v_i for the good.
- Cost of building G is (publicly known) C.

- G should be built if $\sum_{i=1}^{n} v_i > C$.

Goal: Design a mechanism that charges citizens in a way that G is built only

if $\sum_{i=1}^{n} v_i > C$.

Question: Allocating the cost C equally works? Why not?

Answer: For citizen i, if $v_i > \frac{C}{n}$, i should report $C + \epsilon$ so G will be built!

An Example

Problem: Consider a society of n citizens and public good G.

- Each agent has (private) valuation v_i for the good.
- Cost of building G is (publicly known) C.

- G should be built if $\sum_{i=1}^{n} v_i > C$.

Goal: Design a mechanism that charges citizens in a way that G is built only

if $\sum_{i=1}^{n} v_i > C$.

Question: Allocating the cost C equally works? Why not?

Answer: **No DSIC!** $+ \epsilon$ so G will be built!

An Example

Solution: Charge citizen i the amount $p_i := \max(0, C - \sum_{j \neq i} v_i)$.
Similarly can be shown that is DSIC.

An Example

Solution: Charge citizen i the amount $p_i := \max(0, C - \sum_{j \neq i} v_i)$.
Similarly can be shown that is DSIC.

Definition (Quasi-linear environments). *Also known as Vickrey-Groves-Clark (VCG) environments:*

- *n agents,*
- *Set of outcomes (finite) $\mathcal{X}$,*
- *Each agent i has a valuation $v_i : \mathcal{X} \to \mathbb{R}^+$,*
- *Each agent has utility $u_i = v_i - p_i$ where p_i is the received payment (positive or negative).*

An Example

Solution: Charge citizen i the amount $p_i := \max(0, C - \sum_{j \neq i} v_i)$.
Similarly can be shown that is DSIC.

Definition (Quasi-linear environments). *Also known as Vickrey-Groves-Clark (VCG) environments:*

- *n agents,*
- *Set of outcomes (finite) $\mathcal{X}$,*
- *Each agent i has a valuation $v_i : \mathcal{X} \to \mathbb{R}^+$,*
- *Each agent has utility $u_i = v_i - p_i$ where p_i is the received payment (positive or negative).*

Remark: This framework is called mechanism design with money.

VCG mechanisms

Definition (VCG mechanism). *The family of mechanisms is defined as follows:*

- *Agents have 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* $p_i := \underbrace{h_i(b_{-i})}_{\text{without } i} - \underbrace{\sum_{j \neq i} b_i(x^*)}_{\text{with } i}$
- *Each agent has utility* $u_i = v_i(x^*) - p_i(x^*)$.

VCG mechanisms

Definition (VCG mechanism). *The family of mechanisms is defined as follows:*

- *Agents have 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* $$p_i := \underbrace{h_i(b_{-i})}_{\text{without } i} - \underbrace{\sum_{j \neq i} b_i(x^*)}_{\text{with } i}$$
- *Each agent has utility $u_i = v_i(x^*) - p_i(x^*)$.*

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

Proof. Fix an agent i and let $x^* = \arg\max \sum_{i=1}^{n} v_i(x)$. Assume that i reports $b_i \neq v_i$ and x' be the maximum if i reports b_i.

VCG mechanisms

Definition (VCG mechanism). *The family of mechanisms is defined as follows:*

- *Agents have 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* $p_i := \underbrace{h_i(b_{-i})}_{\text{without } i} - \sum_{j \neq i} b_i(x^*)$

By definition of x^* $u_i \geq u'_i$

- *Each agent has utility* $u_i = v_i(x^*) - p_i(x^*)$.

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

Proof. Fix an agent i and let $x^* = \arg\max \sum_{i=1}^{n} v_i(x)$. Assume that i reports $b_i \neq v_i$ and x' be the maximum if i reports b_i. Observe that
$u_i = v_i(x^*) + \sum_{j \neq i} v_i(x^*) - h_i(v_{-i})$ if i reports v_i and
$u'_i = v_i(x') + \sum_{j \neq i} v_i(x') - h_i(v_{-i})$ if i reports b_i.

VCG mechanisms

Question: What is an appropriate choice for h_i?
Answer: So that utility is non-negative (payment at most the valuation).

VCG mechanisms

Question: What is an appropriate choice for h_i?
Answer: So that utility is non-negative (payment at most the valuation).

Definition (VCG mechanism)**.** *The family of mechanisms is defined as follows:*

- *Agents have 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* $p_i := \max_x \sum_{j \neq i} b_i(x) - \sum_{j \neq i} b_i(x^*)$
- *Each agent has utility* $u_i = v_i(x^*) - p_i(x^*)$.

Clark pivot

VCG mechanisms

Question: What is an appropriate choice for h_i?
Answer: So that utility is non-negative (payment at most the valuation).

Definition (VCG mechanism). *The family of mechanisms is defined as follows:*

- *Agents have 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 $p_i := \max_x \sum_{j \neq i} b_i(x) - \sum_{j \neq i} b_i(x^*)$*
- *Each agent has utility $u_i = v_i(x^*) - p_i(x^*)$.*

Clark pivot

Remark: For single-item, this is the second-price auction!
VCG might not be efficiently computable (e.g., combinatorial auctions)