CSC304
Algorithmic Game Theory & Mechanism Design

Nisarg Shah

Announcements

- Assignment 1 solutions uploaded

- Additional office hour
  - Tomorrow (Wed, Oct 19), 2-3pm ET, same Zoom link

- Midterm 1
  - Thursday, Oct 20, 4:10pm – 5:00pm (tutorial slot)
  - In-person
  - EX 100 (Exam Centre)
  - Aid: One 8.5” x 11” sheet of handwritten notes on one side
  - Syllabus: Game theory (first lecture to end of game theory portion in today’s lecture)

Resuming discussion on VCG mechanism

Mathematical Setup

- A = finite set of outcomes
- Each agent i has a private valuation $v_i : A \to \mathbb{R}$
  - Agent i might report $\tilde{v}_i$ instead of the true v_i
- Mechanism consist of a pair of rules (f, p)
  - Input: reported valuations $\tilde{v} = (\tilde{v}_1, \dots, \tilde{v}_n)$
  - $f(\tilde{v}) \in A$ is the outcome implemented
  - $p(\tilde{v}) = (p_1, \dots, p_n)$ are the payments
    - $p_i(\tilde{v})$ is the amount agent i needs to pay
    - Each agent's payment depends on everyone's reports
- Utility to agent i : $u_i(\tilde{v}) = v_i\big(f(\tilde{v})\big) - p_i(\tilde{v})$

Value minus payment

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- Truthfulness/strategyproofness
  - For all agents i, all v_i, and all $\tilde{v}$,
    $$u_i(v_i, \tilde{v}_{-i}) \geq u_i(\tilde{v}_i, \tilde{v}_{-i})$$
  - "Every agent is at least as happy reporting the truth as telling any lie, irrespective of what other agents report"
  - Almost same as telling the truth being a weakly dominant action
    - What's the difference?

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- Individual rationality
  - For all agents i and for all $\tilde{v}_{-i}$,
$$u_i(v_i, \tilde{v}_{-i}) \geq 0$$
  - "No agent should regret participating if she tells the truth."
  - Assumes that the utility from not participating is 0

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- No payments to agents
  - For all agents i and for all $\tilde{v}$,
    $$p_i(\tilde{v}) \geq 0$$
  - "Agents pay the center. Not the other way around."
  - Common for auctions, but we may want the reverse in other settings

Desiderata

- We want the mechanism (f, p) to satisfy some nice properties

- Welfare maximization
  - $f(\tilde{v})$ must be in $\text{argmax}_a \sum_i \tilde{v}_i(a)$
    - Important when making the users happy matters more than the immediate short-term revenue
    - Or think of the auctioneer as "agent $n+1$" with utility equal to the total payment received $\sum_i p_i(\tilde{v})$, and look at total utility

$$\left(\sum_i v_i\big(f(\tilde{v})\big) - p_i\big(f(\tilde{v})\big)\right) + \left(\sum_i p_i\big(f(\tilde{v})\big)\right) = \sum_i v_i\big(f(\tilde{v})\big)$$

Single-item Vickrey Auction

- Simplifying notation:
- $f(\tilde{v})$: give the item to agent $i^* \in \text{argmax}_i\ \tilde{v}_i$
- $p(\tilde{v})$: $p_{i^*} = \max_{j \neq i^*} \tilde{v}_j$, other agents pay nothing

VCG Auction

- Single-item
  - Simplified notation: v_i = value of agent i for the item
  - $f(\tilde{v})$: give the item to agent $i^* \in \text{argmax}_i \, \tilde{v}_i$
  - $p(\tilde{v})$: $p_{i^*} = \max_{j \neq i^*} \tilde{v}_j$, other agents pay nothing
- General setup
  - $f(\tilde{v}) = a^* \in \text{argmax}_{a \in A} \sum_i \tilde{v}_i(a)$ ← Maximize welfare
  - $p_i(\tilde{v}) = \left[\max_a \sum_{j \neq i} \tilde{v}_j(a)\right] - \left[\sum_{j \neq i} \tilde{v}_j(a^*)\right]$

i's payment = welfare that others lost due to presence of i

VCG: Simple Example

- Suppose each agent has a value XBox and a value for PS4.
- Their value for $\{XBox, PS4\}$ is the max of their two values.

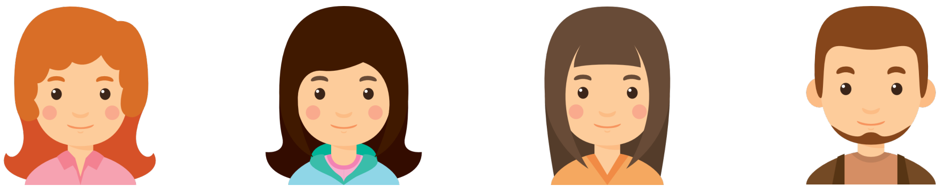

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Q: Who gets the xbox and who gets the PS4?

Q: How much do they pay?

VCG: Simple Example

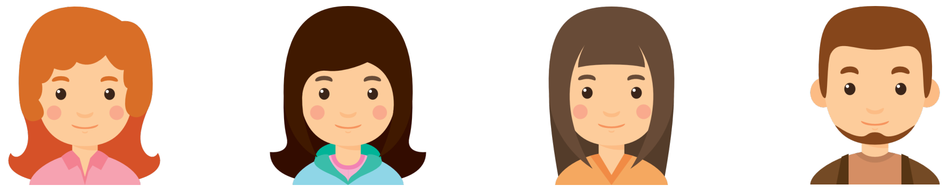

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Allocation:

- A4 gets XBox, A3 gets PS4
- Achieves maximum welfare of $7 + 6 = 13$

VCG: Simple Example

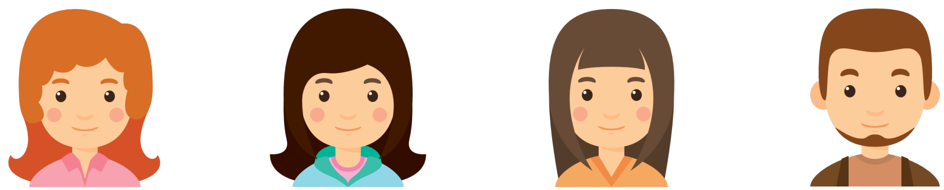

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Zero payments charged to A1 and A2
  - “Deleting” either does not change the outcome/payments for others
- Can also be seen by individual rationality

VCG: Simple Example

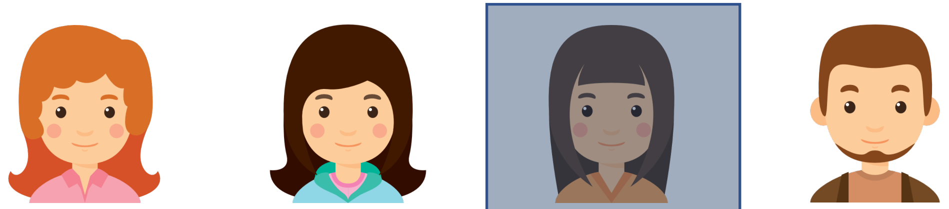

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Payment charged to A3 $= 11 - 7 = 4$
  - Max welfare to others if A3 absent: $7 + 4 = 11$
    - Give XBox to A4 and PS4 to A1
  - Welfare to others if A3 present: 7

VCG: Simple Example

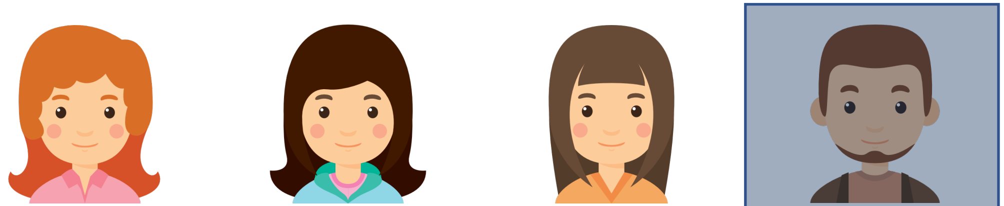

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Payments:

- Payment charged to A4 $= 12 - 6 = 6$
  - Max welfare to others if A4 absent: $8 + 4 = 12$
    - Give XBox to A3 and PS4 to A1
  - Welfare to others if A4 present: 6

VCG: Simple Example

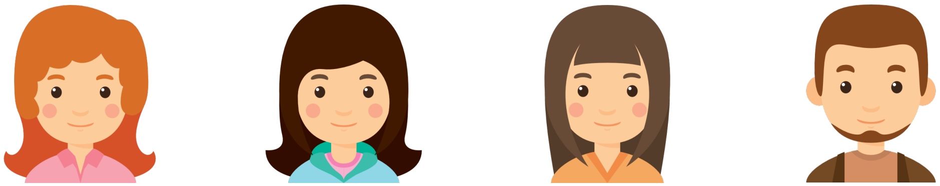

	A1	A2	A3	A4
XBox	3	4	8	7
PS4	4	2	6	1

Final Outcome:

- Allocation: A3 gets PS4, A4 gets XBox
- Payments: A3 pays 4, A4 pays 6
- Net utilities: A3 gets $6 - 4 = 2$, A4 gets $7 - 6 = 1$

Properties of VCG Auction

- Strategyproofness:
  - Suppose agents other than i report $\tilde{v}_{-i}$.
  - Agent i reports $\tilde{v}_i \Rightarrow$ outcome chosen is $f(\tilde{v}) = a$
  - Utility to agent i = $v_i(a) - \left(\blacksquare - \sum_{j \neq i} \tilde{v}_j(a)\right)$

Term that agent i cannot affect

  - Agent i wants a to maximize $v_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
  - f chooses a to maximize $\tilde{v}_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
  - Hence, agent i is best off reporting $\tilde{v}_i = v_i$
    - f chooses a that maximizes the utility to agent i

Properties of VCG Auction

- Individual rationality:
  - $a^* \in \text{argmax}_{a \in A}\, v_i(a) + \sum_{j \neq i} \tilde{v}_j(a)$
  - $\tilde{a} \in \text{argmax}_{a \in A}\, \sum_{j \neq i} \tilde{v}_j(a)$

$$u_i(v_i, \tilde{v}_{-i})$$
$$= v_i(a^*) - \left(\sum_{j \neq i} \tilde{v}_j(\tilde{a}) - \sum_{j \neq i} \tilde{v}_j(a^*)\right)$$
$$= \left[v_i(a^*) + \sum_{j \neq i} \tilde{v}_j(a^*)\right] - \left[\sum_{j \neq i} \tilde{v}_j(\tilde{a})\right]$$

$$= \text{Max welfare to all agents}$$
$$- \text{max welfare to others when } i \text{ is absent}$$
$$\geq 0$$

Properties of VCG Auction

- No payments to agents:
  - Suppose the agents report $\tilde{v}$
  - $a^* \in \text{argmax}_{a \in A} \sum_j \tilde{v}_j(a)$
  - $\tilde{a} \in \text{argmax}_{a \in A} \sum_{j \neq i} \tilde{v}_j(a)$

$$p_i(\tilde{v})$$

$$= \sum_{j \neq i} \tilde{v}_j(\tilde{a}) - \sum_{j \neq i} \tilde{v}_j(a^*)$$

$= \text{Max welfare to others when } i \text{ is absent}$
$\quad - \text{welfare to others when } i \text{ is present}$
≥ 0

Properties of VCG Auction

- Welfare maximization:
  - By definition, since f chooses the outcome maximizing the sum of reported values

- Informal result:
  - Under minimal assumptions, VCG is the unique auction satisfying these properties.

Example: Seller as Agent

- Seller (S) wants to sell his car (c) to buyer (B)

- Seller has a value for his own car: $v_S(c)$
  - Individual rationality for the seller mandates that seller must get revenue at least $v_S(c)$

- Idea: Add seller as another agent and make his values part of the welfare calculations!

Seller as Agent

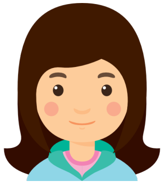

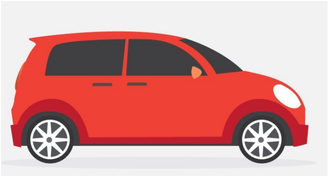

$v_S(c) = 3$

$v_B(c) = 5$

- What if…
  - We give the car to buyer when $v_B(c) > v_S(c)$ and
  - Buyer pays seller $v_B(c)$: Not strategyproof for buyer!
  - Buyer pays seller $v_S(c)$: Not strategyproof for seller!
  - Hmm…what would VCG do?

What would VCG do?

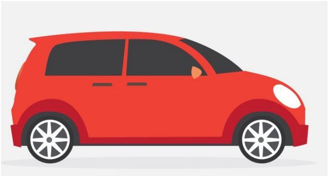

$v_S(c) = 3$

$v_B(c) = 5$

- Allocation?
  - Buyer gets the car (welfare = 5)
- Payment?
  - Buyer pays: $3 - 0 = 3$
  - Seller pays: $0 - 5 = -5$

Mechanism takes \$3 from buyer and gives \$5 to the seller!

- Need external subsidy

Problems with VCG

- Difficult to understand
  - Need to reason about what welfare maximizing allocation in agent i's absence
- Does not care about revenue
  - Although we can lower bound its revenue
- With sellers as agents, need subsidy
  - With no subsidy, cannot get the other three properties
- Might be NP-hard to compute

Single-Minded Bidders

- Selling a set S of m items
- Each agent i has two private values (v_i, S_i)
  - $S_i \subseteq S$ is the subset of items desired by agent i
  - When given a bundle of items A_i, agent i has value v_i if $S_i \subseteq A_i$ and 0 otherwise
  - "Single-minded"
- Welfare-maximizing allocation
  - Agent i either gets S_i or nothing
  - Find a subset of players with the highest total value such that their desired sets are disjoint

Single-Minded Bidders

- Weighted Independent Set (WIS) problem
  - Given a graph with weights on nodes, find an independent set of nodes with the maximum weight
  - Known to be NP-hard
  - Easy to reduce our problem to WIS
    - Not even $O(m^{0.5-\epsilon})$ approximation of welfare unless $NP \subseteq ZPP$

- We will see an algorithm that is:
  - $\sqrt{m}$-approximation
    - Approximation = $\frac{\text{maximum possible welfare}}{\text{welfare achieved by algo}}$ on the worst instance
  - Still strategyproof!

Greedy Algorithm

- Input: (v_i, S_i) for each agent i
- Output: Agents with mutually independent S_i

- Greedy algorithm:
  - Sort the agents in a specific order (we'll see).
  - Relabel them as $1,2, \dots, n$ in this order.
  - $W \leftarrow \emptyset$
  - For $i = 1, \dots, n$:
    - If $S_i \cap S_j = \emptyset$ for every $j \in W$, then $W \leftarrow W \cup \{i\}$
  - Give agents in W their desired items.

Greedy Algorithm

- Sort by what?
- We want to satisfy agents with higher values.
  - $v_1 \geq v_2 \geq \cdots \geq v_n \Rightarrow m$-approximation ☹
- But we don't want to exhaust too many items.
  - $\frac{v_1}{|S_1|} \geq \frac{v_2}{|S_2|} \geq \cdots \frac{v_n}{|S_n|} \Rightarrow m$-approximation ☹
- $\sqrt{m}$-approximation : $\frac{v_1}{\sqrt{|S_1|}} \geq \frac{v_2}{\sqrt{|S_2|}} \geq \cdots \frac{v_n}{\sqrt{|S_n|}}$?

[Lehmann et al. 2011]

Proof of Approximation

- Definitions
  - OPT = Agents satisfied by the optimal algorithm
  - W = Agents satisfied by the greedy algorithm
  - For $i \in W$, $OPT_i = \{j \in OPT, j \geq i : S_i \cap S_j \neq \emptyset\}$
- Claim 1: $OPT \subseteq \cup_{i \in W} OPT_i$
- Claim 2: It is enough to show that $\forall i \in W$

$$\sqrt{m} \cdot v_i \geq \Sigma_{j \in OPT_i} v_j$$

- Observation: For $j \in OPT_i$, $v_j \leq v_i \cdot \frac{\sqrt{|S_j|}}{\sqrt{|S_i|}}$

Proof of Approximation

- Summing over all $j \in OPT_i$:

$$\Sigma_{j \in OPT_i} v_j \leq \frac{v_i}{\sqrt{|S_i|}} \cdot \Sigma_{j \in OPT_i} \sqrt{|S_j|}$$

- Using Cauchy-Schwarz ($\Sigma_i x_i y_i \leq \sqrt{\Sigma_i x_i^2} \cdot \sqrt{\Sigma_i y_i^2}$)

$$\Sigma_{j \in OPT_i} \sqrt{1 \cdot |S_j|} \leq \sqrt{|OPT_i|} \cdot \sqrt{\Sigma_{j \in OPT_i} |S_j|}$$

$$\leq \sqrt{|S_i|} \cdot \sqrt{m}$$

Strategyproofness

- Agent i pays $p_i = v_{j^*} \cdot \sqrt{\frac{|S_i|}{|S_{j^*}|}}$
  - j^* is the smallest index j such that j is currently not selected by greedy but would be selected if we remove (v_i, S_i) from the system
  - Exercise: Show that we must have $j^* > i$
  - Exercise: Show that $S_i \cap S_{j^*} \neq \emptyset$
  - Another interpretation: p_i = lowest value i can report and still win

Strategyproofness

NOT IN SYLLABUS

- Critical payment
  - Charge each agent the lowest value they can report and still win
- Monotonic allocation
  - If agent i wins when reporting (v_i, S_i), she must win when reporting $v_i' \geq v_i$ and $S_i' \subseteq S_i$.
  - Greedy allocation rule satisfies this.
- Theorem: Critical payment + monotonic allocation rule imply strategyproofness.

Moral

- VCG can sometimes be too difficult to implement
  - May look into approximately maximizing welfare
  - As long as the allocation rule is monotone, we can charge critical payments to achieve strategyproofness
  - Note: approximation is needed for computational reasons
- Later in mechanism design without money…
  - We will not be able to use payments to achieve strategyproofness
  - Hence, we will need to approximate welfare just to get strategyproofness, even without any computational restrictions