CSC2556

Lecture 6

Impartial Selection & Facility Location

Announcements

- Reminder
  - Assignment 1 has been posted and is due by 11:59pm ET on Feb 27 (i.e., at the end of the reading week)
  - The assignment is long, so start working on it as soon as possible

- Project
  - This would be a good time to start looking for teammates (Piazza can be useful) and start brainstorming some preliminary project ideas
  - If you want my quick thought on your preliminary idea, you can email me; to discuss it in more detail, email me to set up a 1-1 meeting
  - Proposals will be due in the first week of March

Impartial Selection

Impartial Selection

- "How can we select k people out of n people?"
  - Applications: electing a student representation committee, selecting k out of n grant applications to fund using peer review, …
- Model
  - Input: a *directed* graph $G = (V, E)$
  - Nodes $V = \{v_1, \dots, v_n\}$ are the n people
  - Edge $e = (v_i, v_j) \in E$: v_i supports/approves of v_j
    - We do not allow or ignore self-edges (v_i, v_i)
  - Output: a subset $V' \subseteq V$ with $|V'| = k$
  - $k \in \{1, \dots, n-1\}$ is given

Impartial Selection

- Impartiality: A k-selection rule f is *impartial* if whether or not $v_i \in f(G)$ does not depend on the outgoing edges of v_i
  - v_i cannot manipulate his outgoing edges to get selected
  - Q: But the definition says v_i can neither go from $v_i \notin f(G)$ to $v_i \in f(G)$, nor from $v_i \in f(G)$ to $v_i \notin f(G)$. Why?

- Societal goal: maximize the sum of in-degrees of selected agents $\sum_{v \in f(G)} |in(v)|$
  - $in(v)$ = set of nodes that have an edge to v
  - $out(v)$ = set of nodes that v has an edge to
  - Note: OPT will pick the k nodes with the highest indegrees

Optimal ≠ Impartial

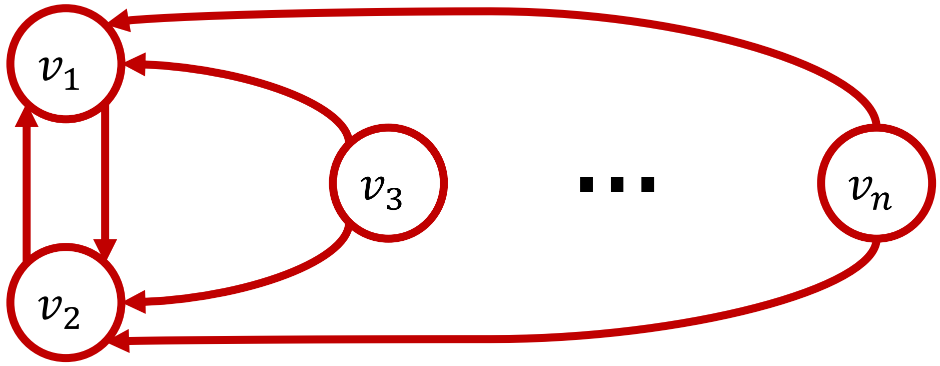

- An optimal 1-selecton rule must select v_1 or v_2
- The other node can remove his edge to the winner, and make sure the optimal rule selects him instead
- This violates impartiality

Goal: Approximately Optimal

- α-approximation: We want a k-selection system that always returns a set with total indegree at least α times the total indegree of the optimal set

- Q: For $k = 1$, what about the following rule?

  Rule: "Select the lowest index vertex in $out(v_1)$.
  If $out(v_1) = \emptyset$, select v_2."

  - A. Impartial + constant approximation
  - B. Impartial + bad approximation
  - C. Not impartial + constant approximation
  - D. Not impartial + bad approximation

No Finite Approximation ☹

- Theorem [Alon et al. 2011]
  For every $k \in \{1, \ldots, n-1\}$, there is no impartial k-selection rule with a finite approximation ratio.
- Proof:
  - For small k, this is trivial. E.g., consider $k = 1$.
    - Consider G that has two nodes v_1 and v_2 that point to each other, and there are no other edges
    - For finite approximation, the rule must choose either v_1 or v_2
    - Say it chooses v_1. If v_2 now removes his edge to v_1, the rule must choose v_2 for any finite approximation, which violates impartiality

No Finite Approximation ☹

- Theorem [Alon et al. 2011]
  For every $k \in \{1, \dots, n-1\}$, there is no impartial k-selection rule with a finite approximation ratio.
- Proof:
  - Proof is more intricate for larger k. Let's do $k = n - 1$.
    - $k = n - 1$: given a graph, "eliminate" a node.
  - Suppose for contradiction that there is such a rule f.
  - W.l.o.g., say v_n is eliminated in the empty graph.
  - Consider a family of graphs in which a subset of $\{v_1, \dots, v_{n-1}\}$ have edges to v_n.

No Finite Approximation ☹

- Proof ($k = n - 1$ continued):
  - Consider *star graphs*
    - A non-empty subset of $\{v_1, \ldots, v_{n-1}\}$ has an edge to v_n and there are no other edges
    - Represented by bit strings $\{0,1\}^{n-1} \backslash \{\vec{0}\}$
  - v_n cannot be eliminated in any star graph (Why?)
  - $f : \{0,1\}^{n-1} \backslash \{\vec{0}\} \rightarrow \{1, \ldots, n-1\}$
    - "Who will be eliminated?"

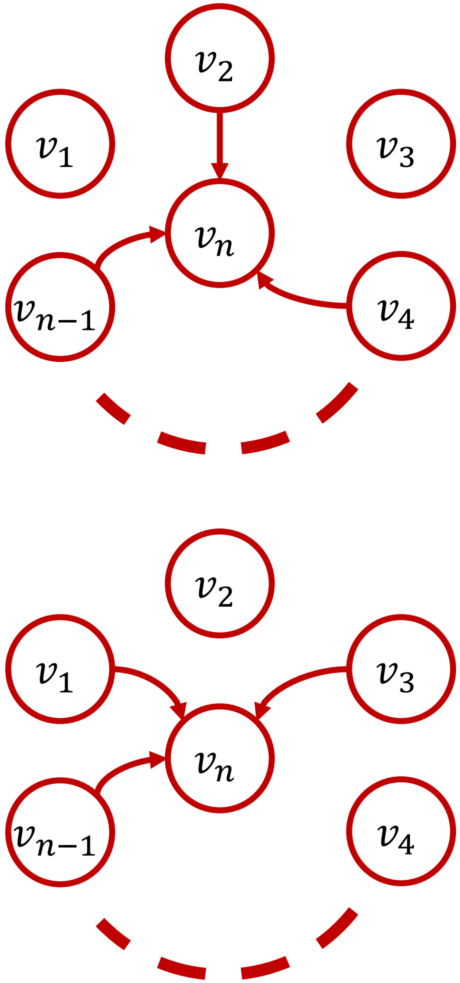

No Finite Approximation ☹

- Proof ($k = n - 1$ continued):
  - Impartiality: $f(\vec{x}) = i \Leftrightarrow f(\text{flip}_i(\vec{x})) = i$
    - flip_i flips the i^{th} coordinate
    - "i cannot add/remove his edge to v_n to change whether he is eliminated"
  - For each i, strings on which f outputs i are paired
    - So, for each i, the number of strings on which f outputs i is even
    - But this is impossible (Why?)
  - So, impartiality must be violated

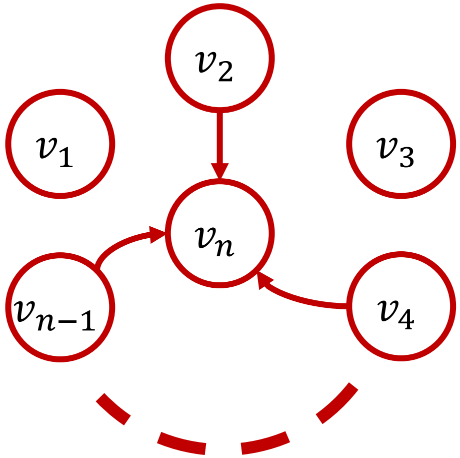

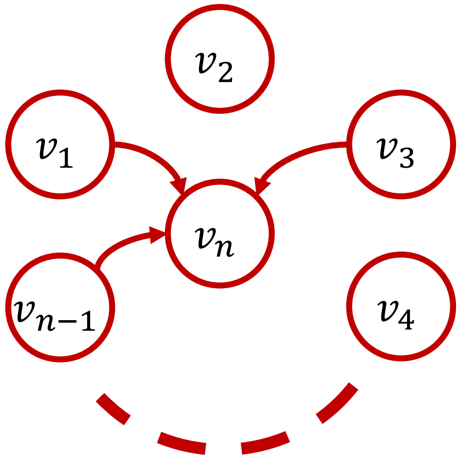

Back to Impartial Selection

- So what *can* we do to select impartially? Randomize!
- Impartiality for randomized mechanisms
  - An agent cannot change the probability of her getting selected by changing her outgoing edges
- Example
  - Choose k nodes uniformly at random
  - Impartial by design
  - Question: What is its approximation ratio?
  - Good when $k \approx n$ but bad when $k \ll n$

Random Partition

- Idea
  - Partition V into V_1 and V_2 and select k nodes from V_1 based only on edges coming to from V_2
  - For impartiality, agents shouldn't be able to affect whether they end up in V_1
  - But a deterministic partition would be bad in the worst case

- Mechanism
  - Assign each node to V_1 or V_2 i.i.d. with probability ½
  - Choose k nodes from V_1 that have most incoming edges from nodes in V_2

Random Partition

- Analysis:
  - OPT = optimal set of k nodes
  - We pick X = k nodes in V_1 with most incoming edges from V_2
  - I = # $V \rightarrow OPT$ edges
  - I' = # $V_2 \rightarrow OPT \cap V_1$ edges
  - Note: $E[I'] = I/4$ (Why?)
  - # incoming edges to $X \geq I'$
    - E[#incoming edges to X] $\geq E[I'] = \frac{I}{4}$

Random Partition

- Generalization
  - Divide into ℓ parts, pick k/ℓ nodes from each part based on incoming edges from all other parts

- Theorem [Alon et al. 2011]:
  - $\ell = 2$ gives a 4-approximation
  - For $k \geq 2$, $\ell \sim k^{1/3}$ gives $1 + O\left(\frac{1}{k^{1/3}}\right)$ approximation

Better Approximations

- Alon et al. [2011]'s conjecture
  - There should be a randomized 1-selection mechanism that achieves 2-approximation
  - Settled by Fischer & Klimm [2014]
  - Permutation mechanism:
    - Select a random permutation $(\pi_1, \pi_2, \dots, \pi_n)$ of the vertices
    - Start by selecting $y = \pi_1$ as the "current answer"
    - At any iteration t, let $y \in \{\pi_1, \dots, \pi_t\}$ be the current answer
    - From $\{\pi_1, \dots, \pi_t\} \setminus \{y\}$, if there are more edges to π_{t+1} than to y, change the current answer to $y = \pi_{t+1}$

Better Approximations

- 2-approximation is tight
  - In an n-node graph, fix u and v, and suppose no other nodes have any incoming/outgoing edges
  - Three cases: only $u \rightarrow v$ edge, only $v \rightarrow u$, or both.
    - The best impartial mechanism selects u and v with probability ½ in every case, and achieves 2-approximation

- Worst case is a bit eccentric
  - $n - 2$ nodes are not voting.
  - What if every node must have an outgoing edge?
  - Fischer & Klimm [2014]
    - In that case, permutation mechanism gives between ${}^{12}/_{7}$ and ${}^{3}/_{2}$ approximation, and no mechanism can do better than ${}^{4}/_{3}$

Facility Location

Facility Location

- Set of agents N
- Each agent i has a true location $x_i \in \mathbb{R}$
- Mechanism f
  - Takes as input reports $\tilde{x} = (\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$
  - Returns a location $y \in \mathbb{R}$ for the new facility
- Cost to agent i : $c_i(y) = |y - x_i|$
- Social cost $C(y) = \sum_i c_i(y) = \sum_i |y - x_i|$

Facility Location

- Social cost $C(y) = \sum_i c_i(y) = \sum_i |y - x_i|$
- Q: Ignoring incentives, what choice of y would minimize the social cost?
- A: The median location $\text{med}(x_1, \dots, x_n)$
  - n is odd → the unique "(n+1)/2"th smallest value
  - n is even → "n/2"th or "(n/2)+1"st smallest value
  - Why?

Facility Location

- Social cost $C(y) = \sum_i c_i(y) = \sum_i |y - x_i|$
- Median is optimal (i.e., 1-approximation)
- What about incentives?
  - Median is also strategyproof (SP)!
  - Irrespective of the reports of other agents, agent i is best off reporting x_i

Informal Proof of SP

No manipulation can help

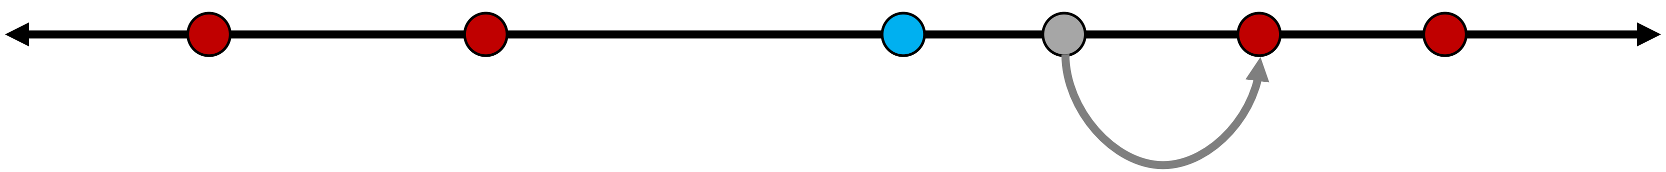

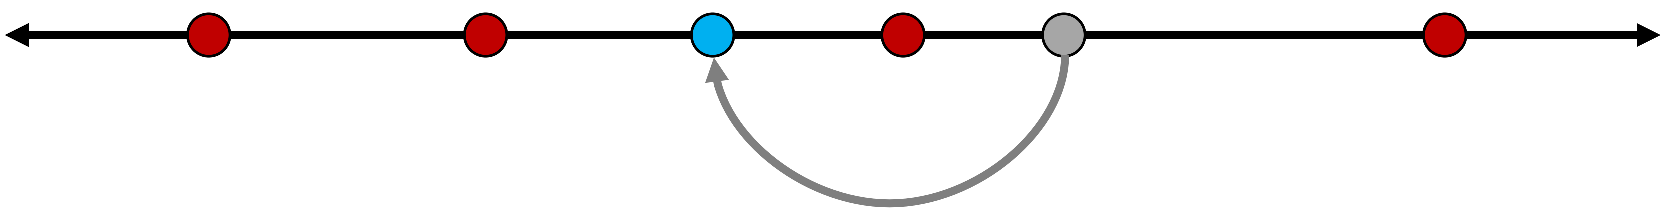

Max Cost

- A different objective function $C(y) = \max_i |y - x_i|$
- Q: Again ignoring incentives, what value of y minimizes the maximum cost?
  - A: The midpoint of the leftmost ($\min_i x_i$) and the rightmost ($\max_i x_i$) locations
- Q: Is this optimal rule strategyproof?
  - A: No!

Max Cost

- $C(y) = \max_i |y - x_i|$
- We want to use a strategyproof mechanism
  - Note: Strategyproofness has nothing to do with the objective function, so median is still SP
- Question: What is the approximation ratio of median for maximum cost?
  1. $\in [1,2)$
  2. $\in [2,3)$
  3. $\in [3,4)$
  4. $\in [4, \infty)$

Max Cost

- Answer: 2-approximation

- Other SP mechanisms that are 2-approximation
  - Leftmost: Choose the leftmost reported location
  - Rightmost: Choose the rightmost reported location
  - Dictatorship: Choose the location reported by agent 1
  - …

Max Cost

- Theorem [Procaccia & Tennenholtz, '09]
  - No deterministic SP mechanism has approximation ratio <2 for maximum cost
- Proof:

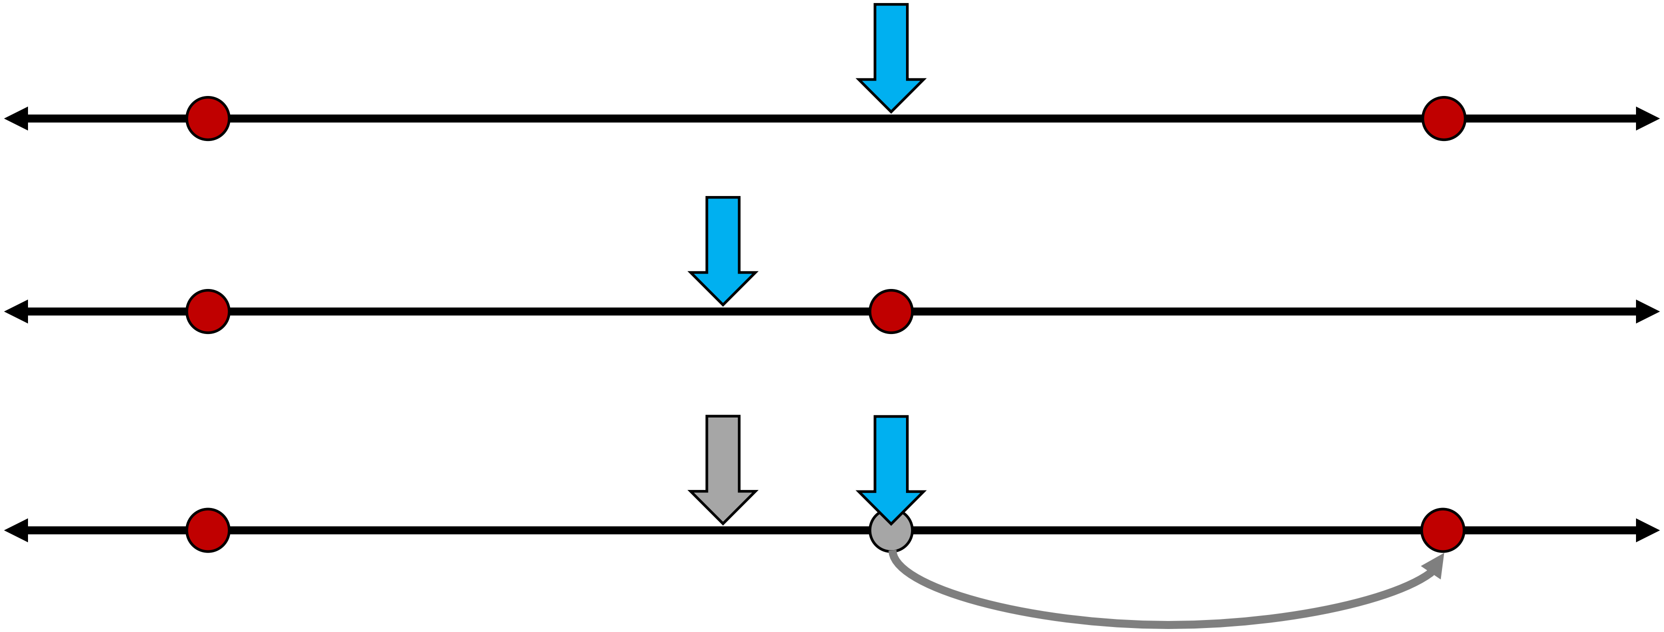

Max Cost + Randomized

- The Left-Right-Middle (LRM) Mechanism
  - Choose $\min_i x_i$ with probability ¼
  - Choose $\max_i x_i$ with probability ¼
  - Choose $(\min_i x_i + \max_i x_i)/2$ with probability ½

- Question: What is the approximation ratio of LRM for maximum cost?

- At most $\frac{(1/4)*2C+(1/4)*2C+(1/2)*C}{C} = \frac{3}{2}$

Max Cost + Randomized

- Theorem [Procaccia & Tennenholtz, '09]:
  The LRM mechanism is strategyproof
- Informal Proof:

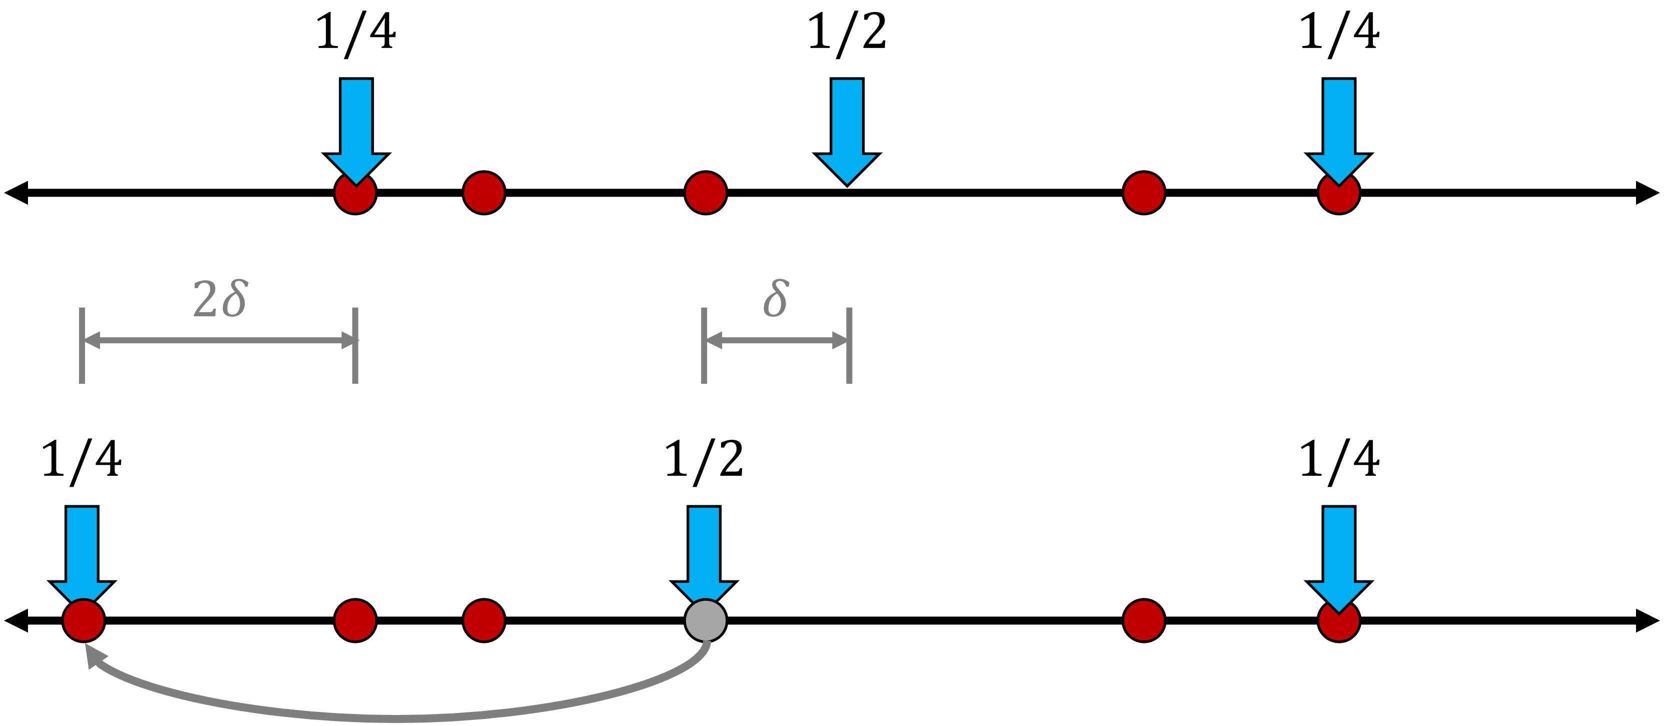

Max Cost + Randomized

- Exercise for you!
  - Try showing that no randomized SP mechanism can achieve approximation ratio $< 3/2$.