CSC304 Lectures 4 & 5

Game Theory (PoA, PoS, Cost sharing & congestion games, Potential function, Braess' paradox)

Recap

- Nash equilibria (NE)
  - No agent wants to change their strategy
  - Guaranteed to exist if mixed strategies are allowed
  - Could be multiple

- Pure NE through best-response diagrams
- Mixed NE through the indifference principle

Worst and Best Nash Equilibria

- What can we say after we identify all Nash equilibria?
  - Compute how "good" they are in the best/worst case

- How do we measure "social good"?
  - Game with only rewards?
    Higher total reward of players = more social good
  - Game with only penalties?
    Lower total penalty to players = more social good
  - Game with rewards and penalties?
    No clear consensus…

Price of Anarchy and Stability

- Price of Anarchy (PoA)

  "Worst NE vs optimum"

  $$\frac{\text{Max total reward}}{\text{Min total reward in any NE}}$$

  or

  $$\frac{\text{Max total cost in any NE}}{\text{Min total cost}}$$

- Price of Stability (PoS)

  "Best NE vs optimum"

  $$\frac{\text{Max total reward}}{\text{Max total reward in any NE}}$$

  or

  $$\frac{\text{Min total cost in any NE}}{\text{Min total cost}}$$

$$\boxed{\text{PoA} \geq \text{PoS} \geq 1}$$

Revisiting Stag-Hunt

Hunter 1 \ Hunter 2	Stag	Hare
Stag	(4 , 4)	(0 , 2)
Hare	(2 , 0)	(1 , 1)

- Max total reward $= 4 + 4 = 8$
- Three equilibria
  - (Stag, Stag) : Total reward = 8
  - (Hare, Hare) : Total reward = 2
  - ($^1/_3$ Stag – $^2/_3$ Hare, $^1/_3$ Stag – $^2/_3$ Hare)
    - Total reward = $\frac{1}{3} * \frac{1}{3} * 8 + \left(1 - \frac{1}{3} * \frac{1}{3}\right) * 2 \in (2,8)$
- Price of stability? Price of anarchy?

Revisiting Prisoner's Dilemma

Sam \ John	Stay Silent	Betray
Stay Silent	(-1 , -1)	(-3 , 0)
Betray	(0 , -3)	(-2 , -2)

- Min total cost $= 1 + 1 = 2$
- Only equilibrium:
  - (Betray, Betray) : Total cost $= 2 + 2 = 4$

- Price of stability? Price of anarchy?

Cost Sharing Game

- n players on directed weighted graph G

- Player i
  - Wants to go from s_i to t_i
  - Strategy set S_i = {directed $s_i \rightarrow t_i$ paths}
  - Denote his chosen path by $P_i \in S_i$

- Each edge e has cost c_e (weight)
  - Cost is split among all players taking edge e
  - That is, among all players i with $e \in P_i$

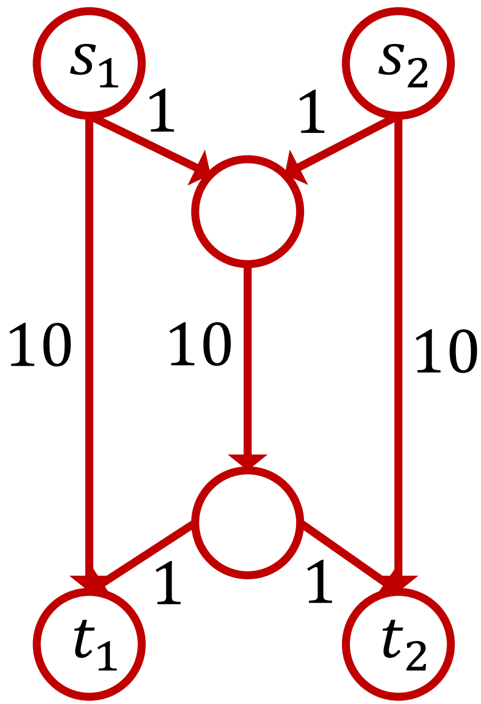

Cost Sharing Game

- Given strategy profile $\vec{P}$, cost $c_i(\vec{P})$ to player i is sum of his costs for edges $e \in P_i$
- Social cost $C(\vec{P}) = \sum_i c_i(\vec{P})$
- Note: $C(\vec{P}) = \sum_{e \in E(\vec{P})} c_e$, where…
  - $E(\vec{P})$={edges taken in $\vec{P}$ by at least one player}
  - Why?

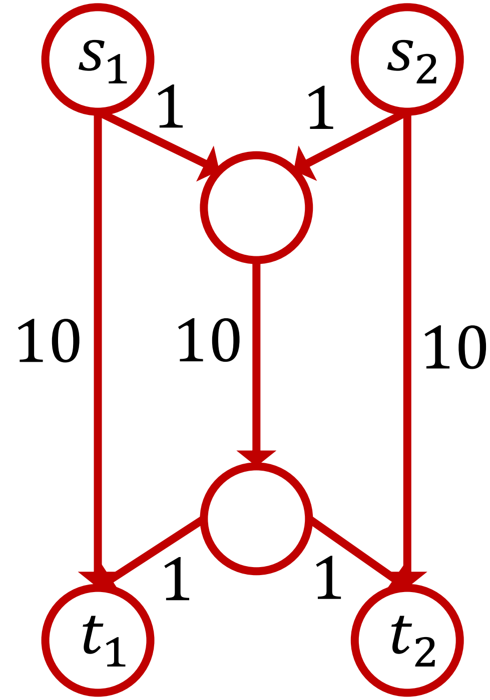

Cost Sharing Game

- In the example on the right:
  - What if both players take direct paths?
  - What if both take middle paths?
  - What if one player takes direct path and the other takes middle path?

- Pure Nash equilibria?

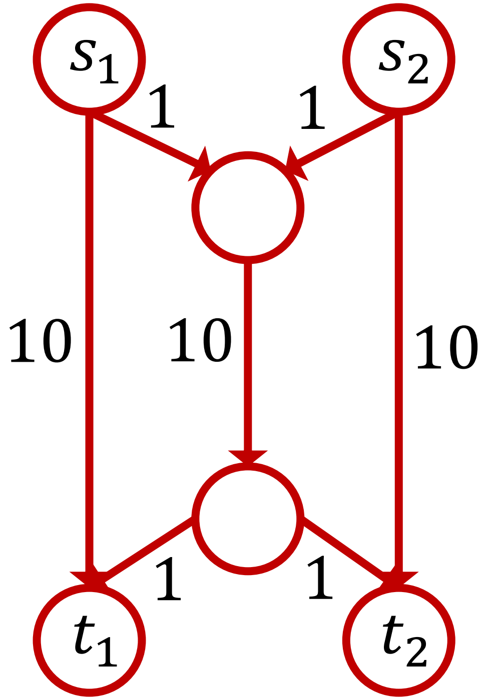

Cost Sharing: Simple Example

- Example on the right: n players

- Two pure NE
  - All taking the n-edge: social cost = n
  - All taking the 1-edge: social cost = 1
    - Also the social optimum

- Price of stability: 1
- Price of anarchy: n
  - We can show that price of anarchy $\leq n$ in *every* cost-sharing game!

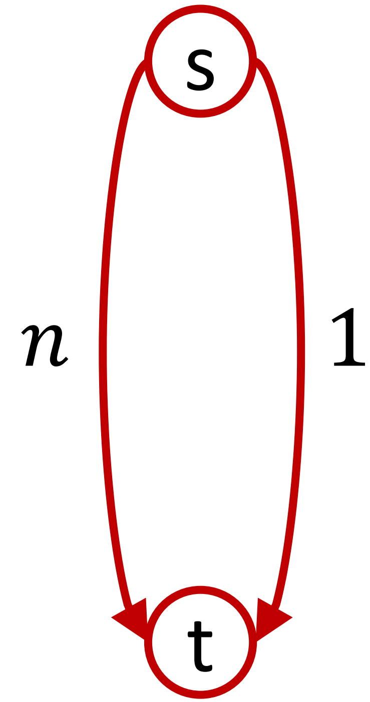

Cost Sharing: PoA

- Theorem: The price of anarchy of a cost sharing game is at most n.
- Proof:
  - Suppose the social optimum is $(P_1^*, P_2^*, \ldots, P_n^*)$, in which the cost to player i is c_i^*.
  - Take any NE with cost c_i to player i.
  - Let c_i' be his cost if he switches to P_i^*.
  - NE $\Rightarrow c_i' \geq c_i$ (Why?)
  - But : $c_i' \leq n \cdot c_i^*$ (Why?)
  - $c_i \leq n \cdot c_i^*$ for each i $\Rightarrow$ no worse than $n \times$ optimum ■

Cost Sharing

- Price of anarchy
  - Every cost-sharing game: PoA $\leq n$
  - Example game with PoA $= n$
  - Bound of n is tight.

- Price of stability?
  - In the previous game, it was 1.
  - In general, it can be higher. How high?
  - We'll answer this after a short detour.

Cost Sharing

- Nash's theorem shows existence of a mixed NE.
  - Pure NE may not always exist in general.

- But in both cost-sharing games we saw, there was a PNE.
  - What about a more complex game like the one on the right?

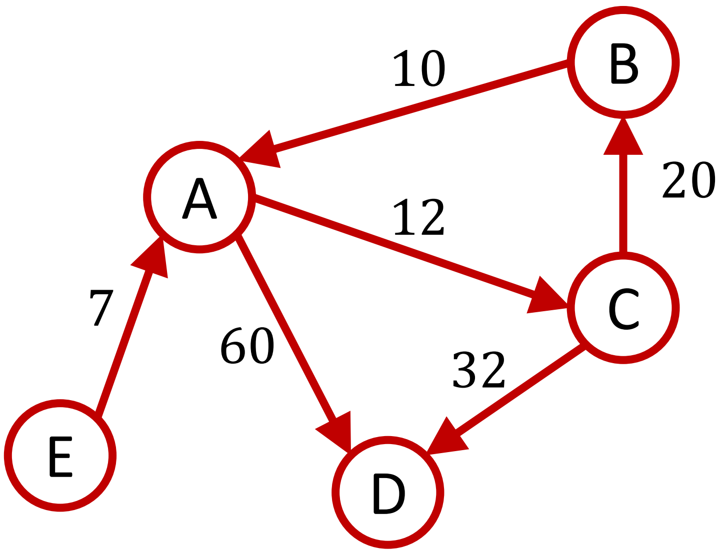

10 players: $E \to C$
27 players: $B \to D$
19 players: $C \to D$

Good News

- Theorem: Every cost-sharing game has a pure Nash equilibrium.

- Proof:
  - Via "potential function" argument

Step 1: Define Potential Fn

- Potential function: $\Phi : \prod_i S_i \rightarrow \mathbb{R}_+$
  - This is a function such that for every pure strategy profile $\vec{P} = (P_1, \ldots, P_n)$, player i, and strategy P_i' of i,

$$c_i(P_i', \vec{P}_{-i}) - c_i(\vec{P}) = \Phi(P_i', \vec{P}_{-i}) - \Phi(\vec{P})$$

  - When a single player i changes her strategy, the change in potential function equals the change in cost to i!
  - Note: In contrast, the change in the social cost C equals the total change in cost to all players.
    - Hence, the social cost will often not be a valid potential function.

Step 2: Potential $F^n \rightarrow$ pure Nash Eq

- A potential function exists ⇒ a pure NE exists.
  - Consider a $\vec{P}$ that minimizes the potential function.
  - Deviation by any single player i can only (weakly) increase the potential function.
  - But change in potential function = change in cost to i.
  - Hence, there is no beneficial deviation for any player.

- Hence, every pure strategy profile minimizing the potential function is a pure Nash equilibrium.

Step 3: Potential F^n for Cost-Sharing

- Recall: $E(\vec{P}) = \{\text{edges taken in } \vec{P} \text{ by at least one player}\}$

- Let $n_e(\vec{P})$ be the number of players taking e in $\vec{P}$

$$\Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k}$$

- Note: The cost of edge e to each player taking e is $c_e / n_e(\vec{P})$. But the potential function includes all fractions: $c_e/1, c_e/2, \dots, c_e/n_e(\vec{P})$.

Step 3: Potential F^n for Cost-Sharing

$$\Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k}$$

- Why is this a potential function?
  - If a player changes path, he pays $\frac{c_e}{n_e(\vec{P})+1}$ for each new edge e, gets back $\frac{c_f}{n_f(\vec{P})}$ for each old edge f.
  - This is precisely the change in the potential function too.
  - So $\Delta c_i = \Delta\Phi$.

■

Potential Minimizing Eq.

- Minimizing the potential function gives some pure Nash equilibrium
  - Is this equilibrium special? Yes!

- Recall that the price of anarchy can be up to n.
  - That is, the worst Nash equilibrium can be up to n times worse than the social optimum.

- A potential-minimizing pure Nash equilibrium is better!

Potential Minimizing Eq.

- $$\sum_{e \in E(\vec{P})} c_e \le \Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k} \le \sum_{e \in E(\vec{P})} c_e * \sum_{k=1}^{n} \frac{1}{k}$$

Social cost

- $$\forall \vec{P},\ C(\vec{P}) \le \Phi(\vec{P}) \le C(\vec{P}) * H(n)$$

Harmonic function $H(n)$ $= \sum_{k=1}^{n} 1/n = O(\log n)$

- $$C(\vec{P}^*) \le \Phi(\vec{P}^*) \le \Phi(OPT) \le C(OPT) * H(n)$$

Potential minimizing eq.

Social optimum

Potential Minimizing Eq.

- Potential-minimizing PNE is $O(\log n)$-approximation to the social optimum.

- Thus, in every cost-sharing game, the price of stability is $O(\log n)$.
  - Compare to the price of anarchy, which can be n

Congestion Games

- Generalize cost sharing games
- n players, m resources (e.g., edges)
- Each player i chooses a set of resources P_i (e.g., $s_i \to t_i$ paths)
- When n_j player use resource j, each of them get a cost $f_j(n_j)$
- Cost to player is the sum of costs of resources used

Congestion Games

- Theorem [Rosenthal 1973]: Every congestion game is a potential game.
- Potential function:

$$\Phi(\vec{P}) = \sum_{j \in E(\vec{P})} \sum_{k=1}^{n_j(\vec{P})} f_j(k)$$

- Theorem [Monderer and Shapley 1996]: Every potential game is equivalent to a congestion game.

Potential Functions

- Potential functions are useful for deriving various results
  - E.g., used for analyzing amortized complexity of algorithms

- Bad news: Finding a potential function that works may be hard.

The Braess' Paradox

- In cost sharing, f_j is decreasing
  - The more people use a resource, the less the cost to each.
- f_j can also be increasing
  - Road network, each player going from home to work
  - Uses a sequence of roads
  - The more people on a road, the greater the congestion, the greater the delay (cost)
- Can lead to unintuitive phenomena

The Braess' Paradox

- Parkes-Seuken Example
  - 2000 players want to go from 1 to 4
  - $1 \rightarrow 2$ and $3 \rightarrow 4$ are "congestible" roads
  - $1 \rightarrow 3$ and $2 \rightarrow 4$ are "constant delay" roads

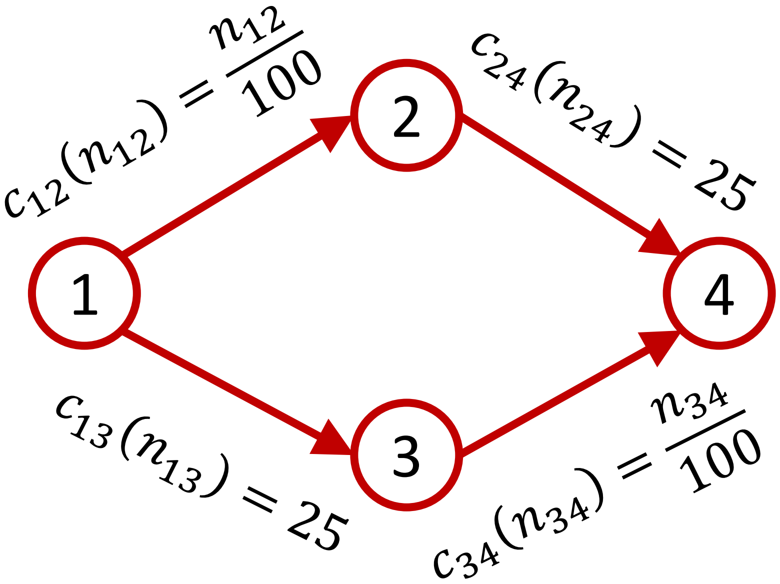

The Braess' Paradox

- Pure Nash equilibrium?
  - 1000 take $1 \to 2 \to 4$, 1000 take $1 \to 3 \to 4$
  - Each player has cost $10 + 25 = 35$
  - Anyone switching to the other creates a greater congestion on it, and faces a higher cost

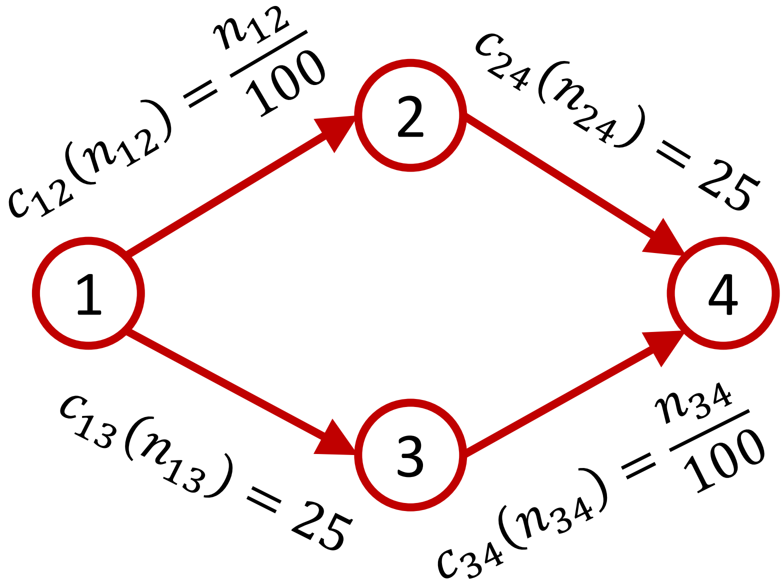

The Braess' Paradox

- What if we add a zero-cost connection $2 \rightarrow 3$?
  - Intuitively, adding more roads should only be helpful
  - In reality, it leads to a greater delay for everyone in the unique equilibrium!

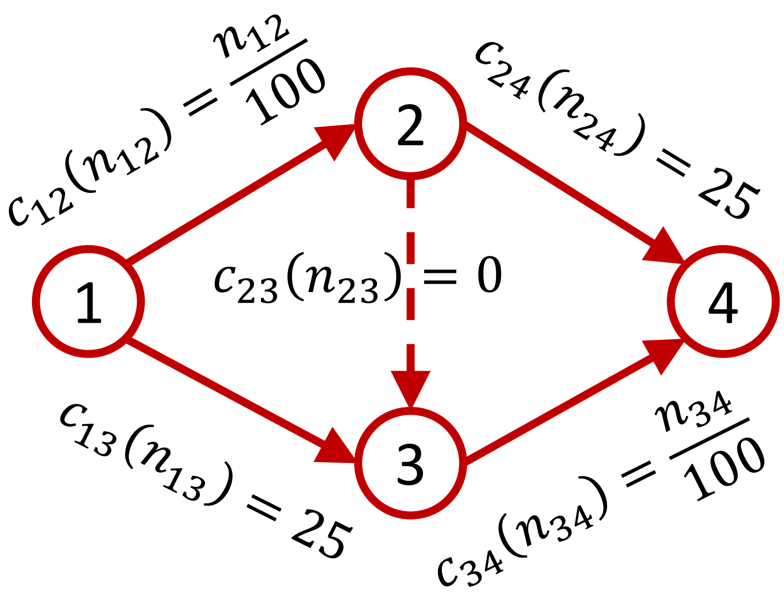

The Braess' Paradox

- Nobody chooses $1 \rightarrow 3$ as $1 \rightarrow 2 \rightarrow 3$ is better irrespective of how many other players take it
- Similarly, nobody chooses $2 \rightarrow 4$
- Everyone takes $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$, faces delay = 40!

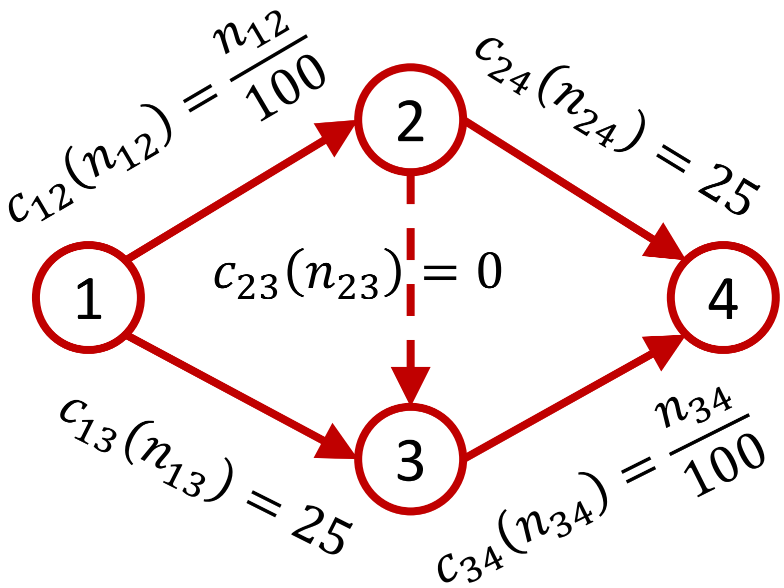

The Braess' Paradox

- In fact, what we showed is:
  - In the new game, $1 \to 2 \to 3 \to 4$ is a strictly dominant strategy for each player!