The Price of Anarchy

—

Lecture 4

1 Routing Games

We'll start off this lecture by looking at a surprising real-world phenomenon: Sometimes, closing a road can speed up traffic!

There are many documented occurrences of this phenomenon in the real world. A well-known instance of this is in Seoul, South Korea, where the city closed a major highway along the Cheonggyecheon Stream to revitalize the area. Counterintuitively, among other benefits, traffic flow across the city improved as drivers redistributed across other routes.

Indeed, a 2008 paper conjectured that we don't even need to look this far to find an example: They claimed that traffic would speed up in Cambridge, if Main Street was closed between Central and Kendall Square. Currently, the fastest route between Harvard Square (the red 'H' node below) and Downtown Boston (the green 'B' node) is to use the yellow route: Take Mass Ave to Central Square, then Main Street to Longfellow Bridge. However, this causes timely congestion on the single-lane Mass Ave and Longfellow Bridge. If, instead, Main Street (dashed) was closed, traffic would be distributed across Mass Ave, Broadway, and Cambridge Street, as well as across Harvard Bridge, Longfellow Bridge, and the Charles River Dam Bridge. The conjecture is that this would lead to commute times decreasing for all drivers.

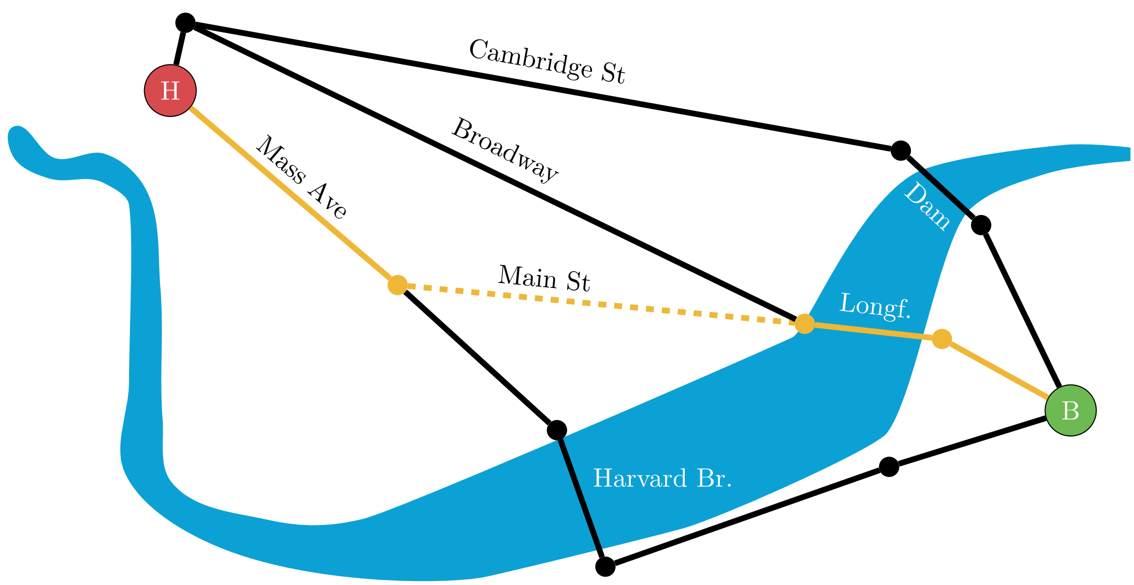

We can model this paradox mathematically:

Example 1 (Braess' Paradox)**.** Consider a situation in which 200 travelers want to get from work (A) to home (D). Initially, each traveler can take one of two paths:

- Go from A to B, taking $k/10$ minutes where k is the number of travelers who take this path, and then go from B to D, which takes 20 minutes, independently of the number of travelers.

- Go from A to C, which takes 20 minutes, independently of the number of travelers, and then go from C to D, taking $k/10$ minutes where k is the number of travelers who take this path.

A graphical depiction of these paths is provided below.

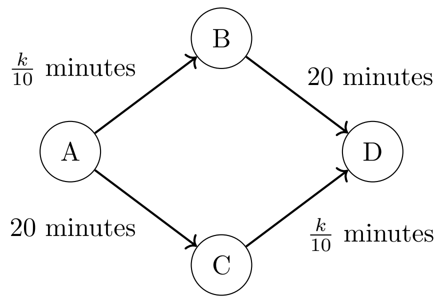

Assume 100 travelers take each route. Then, the travel time for each traveler is 30 minutes, because both A to B and C to D take $\frac{100}{10} = 10$ minutes. This is a Nash equilibrium: If any traveler changed to the other path, their traveling time increases to $20 + \frac{101}{10} = 30.1$ minutes. It is not hard to check that this is indeed the only (pure) Nash equilibrium.

Now, assume a new, very fast 2-way street between B and C is built, allowing to travel between these nodes in (approximately) 0 minutes as shown below:

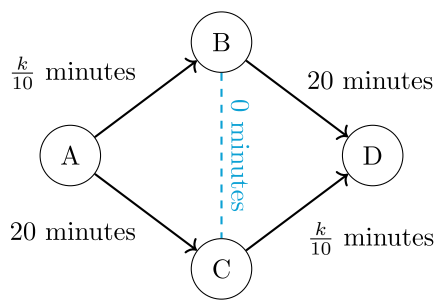

In this case, 100 travelers choosing $A \to B \to D$ and $A \to C \to D$, respectively, is no longer a Nash equilibrium: Any traveler would be better of choosing $A \to B \to C \to D$, bringing their travel time down to 20 minutes.

However, this becomes a problem if all travelers take $A \to B \to C \to D$, since now every traveler takes $\frac{200}{10} + \frac{200}{10} = 40$ minutes to get to D. However, this is a Nash equilibrium![1] No traveler has a useful unilateral deviation because flipping to any of $A \to B \to D$, $A \to C \to D$, or $A \to C \to B \to D$ all result in the same 40 minutes of travel time. Indeed, this is the unique Nash equilibrium.

This phenomenon is called *Braess' paradox*: The Nash equilibrium outcome after adding a new road results in longer travel times for *all* travelers than the Nash equilibrium outcome before the road was built.

We now formalize the game played by the travelers picking a route to get a generalized version of Braess' paradox.

Definition 1 (Routing Games)**.** An (atomic) *routing game* consists of

- a set of *players* $N = \{1, \dots, n\}$,
- a directed graph $G = (V, E)$,
- for each edge $e \in E$, a nonnegative and nondecreasing *cost function* $c_e : \mathbb{N} \to \mathbb{R}^+$, and
- for each player $i \in N$, a *source* and *sink* vertex $a_i, b_i \in V$.

[1] In fact, this is another example of the tragedy of the commons, introduced in lecture 1.

The strategy set S_i of each player i is the set of all paths from their source a_i to their sink b_i in G. In a strategy profile $\boldsymbol{s} = (s_1, ..., s_n)$, we use $n_e(\boldsymbol{s})$ to denote the number players who are using edge e in their strategy s_i under $\boldsymbol{s}$. The cost of player i is the sum over the costs (corresponding to congestion) for all edges they use,

$$\text{cost}_i(\boldsymbol{s}) = \sum_{e \in s_i} c_e(n_e(\boldsymbol{s})).$$

Intuitively, this corresponds to the player's total 'travel time.' Finally, the social cost is the total cost across all players,

$$\text{cost}(\boldsymbol{s}) = \sum_{i \in N} \text{cost}_i(\boldsymbol{s}).$$

Example 2 (Routing game)**.** Consider the following routing game between 3 players. For each $i \in \{1, 2, 3\}$, player i's source is labeled as a_i and their sink is labeled as b_i. Each edge's cost function c_e is written next to the edge as a function of the number of players k who traverse that edge.

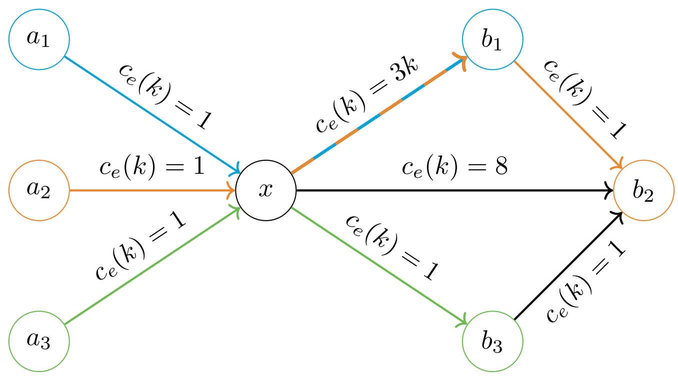

Let's assume that player 1 takes the blue path, player 2 takes the orange path, and player 3 takes the green path. Then player 1's cost is $1 + 3 \cdot 2 = 7$, player 2's cost is $1 + 3 \cdot 2 + 1 = 8$, and player 3's cost is $1 + 1 = 2$, so we get a social cost of 17.

In the solution with lowest social cost, player 1 takes $a_1 \to x \to b_1$, player 2 takes $a_2 \to x \to b_3 \to b_2$, and player 3 takes $a_3 \to x \to b_3$. Then, player 1's cost is $1 + 3 \cdot 1 = 4$, player 2's cost is $1 + 1 + 1 = 3$, and player 3's cost is $1 + 1 = 2$, resulting in a social cost of 9.

In Example 1 of Braess' paradox, there always existed a pure Nash equilibrium in the routing game. Is this the case in all routing games? To answer this question, we make use of the fact that routing games are a subset of a larger class of games called exact potential game.

Definition 2 (Exact Potential Games)**.** A game is an *exact potential game* if there exists a function $\Phi : \prod_{i=1}^n S_i \to \mathbb{R}$ such that for all players $i \in N$, for all strategy profiles $\boldsymbol{s} \in \prod_{i=1}^n S_i$, and for all alternate strategies of player i, $s_i' \in S_i$, it holds that

$$\text{cost}_i(s_i', \boldsymbol{s}_{-i}) - \text{cost}_i(\boldsymbol{s}) = \Phi(s_i', \boldsymbol{s}_{-1}) - \Phi(\boldsymbol{s})$$

In other words, for any strategy profile $\boldsymbol{s}$, when player i unilaterally deviates to s_i', the change in value of Φ equals the change in player i's cost. We call Φ an *exact potential function* for the game.

All exact potential games have a pure Nash equilibrium, namely any pure strategy profile $\boldsymbol{s}$ that minimizes $\Phi(\boldsymbol{s})$. At this minimal $\boldsymbol{s}$, no change to $\boldsymbol{s}$ can further decrease $\Phi(\boldsymbol{s})$, so no deviation of any player from $\boldsymbol{s}$ can reduce their cost — $\boldsymbol{s}$ is a pure Nash equilibrium.

It turns out that routing games actually do fall under this category of games, and that therefore the existence of a pure Nash equilibrium is guaranteed.

Theorem 1. *Routing games are exact potential games.*

We prove this in the Appendix; reading this proof is fully optional.

2 Price of Anarchy

In Example 1, we saw that the social cost in a Nash equilibrium can be larger than the optimal (lowest-possible) social cost. We will now try to quantify this, by investigating how bad the social cost of an equilibrium can be as compared to the optimal solution.

Definition 3 (Price of Anarchy)**.** The *price of anarchy (PoA)* of a game with some objective function (e.g., social cost) is the ratio between the worst objective function value in any equilibrium of the game and the optimal value of the objective function.

For the rest of this lecture, our objective function will be the social cost, and our equilibrium concept will be pure Nash equilibria. In that case, the price of anarchy is

$$\frac{\max_{\boldsymbol{s}\in \mathrm{PNE}} \mathrm{cost}(\boldsymbol{s})}{\min_{\boldsymbol{s}\in S^n} \mathrm{cost}(\boldsymbol{s})},$$

where PNE is the set of all pure Nash equilibria and S^n is the set of all possible pure strategy profiles.

Example 3 (PoA of routing games 1)**.** Consider the following routing game, where n players wish to go from A to C. The edges are labeled by their cost function when k players go through the path.

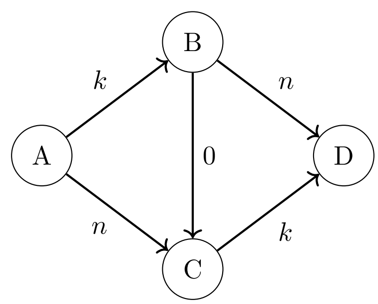

In this game, the unique pure Nash equilibrium is all n players taking the path $A \to B \to C \to D$, resulting in a travel time of $2n$ for everyone. If any single player deviates to $A \to B \to D$ or $A \to C \to D$, their cost is still the same $2n$. The social cost is $2n \cdot n = 2n^2$.

In contrast, the solution with optimal social cost is to send $n/2$ players along the path $A \to B \to D$ and $n/2$ players along the path $A \to C \to D$, resulting in a social cost of $n/2 \cdot n/2 + n/2 \cdot n + n/2 \cdot n + n/2 \cdot n/2 = 3n^2/2$.

Therefore, the price of anarchy in this case is

$$\frac{2n^2}{3n^2/2} = \frac{4}{3}.$$

Example 4 (PoA of routing games 2)**.** Let's consider another routing game with 4 players and source-sink pairs (a_i, b_i) being

$$\{(A,B),(A,C),(B,C),(C,B)\}$$

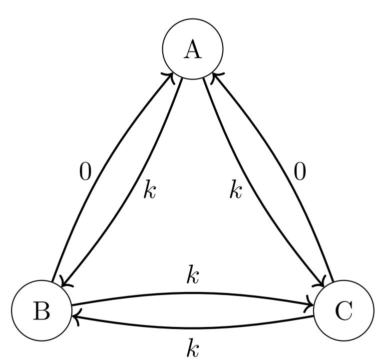

In this game, the solution with optimal social cost is for all players to take the path with cost k directly to their sink: Player 1 takes $A \to B$, player 2 takes $A \to C$, player 3 takes $B \to C$, and player 4 takes $C \to B$. This results in a social cost of

$$1 \cdot 1 + 1 \cdot 1 + 1 \cdot 1 + 1 \cdot 1 = 4$$

This optimal solution is also a Nash equilibrium. For any player, consider any other path they could take to their sink. In every path, the player must traverse at least one edge with cost k, so there is no beneficial deviation for them.

However, this is not the only Nash equilibrium in the game. Consider player 1 taking $A \to C \to B$, player 2 taking $A \to B \to C$, player 3 taking $B \to A \to C$, and player 4 taking $C \to A \to B$. Here, player 1's cost is $2+1=3$, player 2's cost is $2+1=3$, player 3's cost is $0+2=2$, and player 4's cost is $0+2=2$, resulting in an overall social cost of 10. To see that this is a Nash equilibrium, note that visiting any node more than once is never optimal, so we only need to consider paths of length 1 and 2. All players are currently taking the unique path of length 2, and it can quickly be checked that unilaterally deviating to the 'direct' path of length 1 is not beneficial for any player. For example, if player 1 deviated to $A \to B$, now players 1, 2, and 4 would be taking $A \to B$, so player 1 still has a cost of 3.

Therefore, we get that a lower bound (there may be even worse equilibria, for all we know so far) for the price of anarchy in this case is

$$^{10}/_{4} = 2.5.$$

It turns out that this lower bound is tight:

Theorem 2. *For any routing game with linear cost functions, the price of anarchy is at most* 2.5.

In other words, for any routing game with linear cost functions and Nash equilibrium $\boldsymbol{s}$, we know that $\text{cost}(\boldsymbol{s}) \leq 2.5 \cdot \text{cost}(\text{OPT})$. Furthermore, by Example 4, we know that there exists a routing game with linear cost functions that has a Nash equilibrium $\boldsymbol{s}$ where $\text{cost}(\boldsymbol{s}) \geq 2.5 \cdot \text{cost}(\text{OPT})$. In summary, we know that the worst-case PoA for any routing game with linear cost functions is exactly 2.5.

To conclude this section, we will see that such a constant upper bound is not possible if the cost functions can be arbitrary, nondecreasing and nonnegative functions.

Example 5 (Routing game with non-linear cost functions)**.** Consider a routing game with n players consisting of two edges from a common source A to a common sink B, one with cost 1 and one with cost $(k/n)^p$.

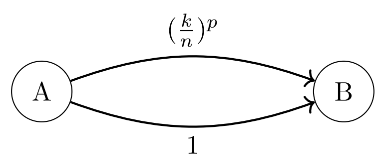

It is Nash equilibrium for all n people to take the edge with cost $(k/n)^p$: Currently, each individual player has cost $(n/n)^p = 1$, so deviating to the other edge for the fixed cost of 1 is not beneficial. This Nash equilibrium results in a total social cost of

$$n \cdot \left(\frac{n}{n}\right)^p = n \cdot 1 = n.$$

On the other hand, consider the strategy profile where one player chooses to take the edge with a fixed cost of 1 and the remaining $n-1$ players take the path with a cost of $(k/n)^p$. Then, the total social cost is

$$n \cdot \left(\frac{n-1}{n}\right)^p + 1$$

As $p \to \infty$, the first term goes to 0. Thus, for any number of player n, there exists a routing game with price of anarchy as close to $\frac{n}{1} = n$ as desired (by making p sufficiently large).

3 Cost Sharing Games

We will now introduce a similar class of games where congestion is beneficial instead of harmful.

Definition 4 (Cost Sharing Games)**.** Just like a routing game, a *cost sharing game* consists of a set of players N, a directed graph G, and a source and sink vertex, a_i and b_i, for each player. Each edge has a fixed cost c_e that is split among the players using it. Thus, the cost of player i taking path s_i in strategy profile $\boldsymbol{s}$ is

$$\text{cost}_i(\boldsymbol{s}) = \sum_{e \in s_i} \frac{c_e}{n_e(\boldsymbol{s})}.$$

As an example, consider a group of travelers deciding where to build roads to all get from their origin to their destination, with the cost of building a road being evenly split among all travelers that use it.

Example 6 (Basic cost sharing game)**.** In the following cost sharing game played between two players, the first player wishes to travel from A to B and the second player wishes to travel from C to D.

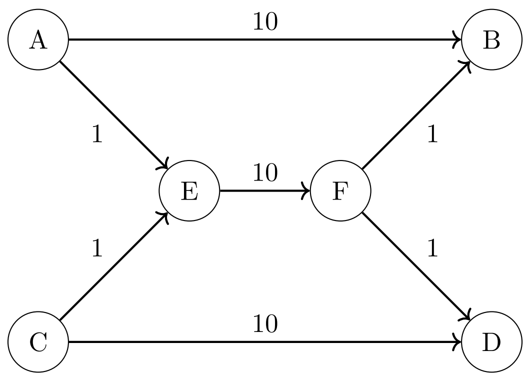

In the strategy profile where player 1 takes $A \to B$ and player 2 takes $C \to D$, the social cost is $10 + 10 = 20$. It is also Nash equilibrium: If either player deviates to taking the middle path, their cost will increase from 10 to 12.

However, the solution with lowest social cost is when player 1 takes $A \to E \to F \to B$ and player 2 takes $C \to E \to F \to D$, for a total social cost of $2(1 + 5 + 1) = 14$.

Example 7 (Lower bound on PoA for cost sharing games)**.** Consider a cost sharing game with n players and two edges between a common source A and a common sink B with costs 1 and n.

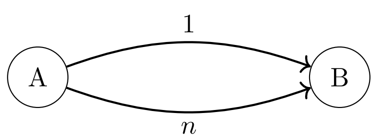

In the strategy profile where all n players take the edge with cost n, the social cost is n: The cost of this edge is split across the n players, each individually having a cost of 1. This strategy profile is also a Nash equilibrium, as no player would benefit from unilaterally deviating to the path with a cost of 1.

However, the socially optimal solution is for all players to take the path with a cost of 1 achieving a total social cost of 1. Thus, we know that the price of anarchy of this cost sharing game is at least $\frac{n}{1} = n$.

It turns out that this lower bound is tight:

Theorem 3. *The price of anarchy of any cost sharing game is at most n.*

Proof. Let $\boldsymbol{s}$ be a Nash equilibrium and let $\boldsymbol{s}^\star$ be a solution minimizing social cost. For all i, it holds that $\text{cost}_i(\boldsymbol{s}) \leq \text{cost}_i(s_i^\star, \boldsymbol{s}_{-i})$ because i can't gain from unilaterally deviating in a Nash equilibrium. But $\text{cost}_i(s_i^\star, \mathbf{s}_{-i}) \leq n \cdot \text{cost}_i(\mathbf{s}^\star)$, because in the worst case i pays for its path alone in the former, $(s_i^\star, \mathbf{s}_{-i})$, and in the best case splits each edge cost n ways in the latter, $\boldsymbol{s}^*$. Thus, we get that

$$\frac{\text{cost}(\boldsymbol{s})}{\text{cost}(\boldsymbol{s}^\star)} = \frac{\sum_{i \in N} \text{cost}_i(\boldsymbol{s})}{\text{cost}(\boldsymbol{s}^\star)} \leq \frac{\sum_{i \in N} \text{cost}_i(s_i^*, \boldsymbol{s}_{-i})}{\text{cost}(\boldsymbol{s}^\star)} \leq \frac{n \sum_{i \in N} \text{cost}_i(\boldsymbol{s}^*)}{\text{cost}(\boldsymbol{s}^\star)} \leq \frac{n\text{cost}(\boldsymbol{s}^*)}{\text{cost}(\boldsymbol{s}^\star)} = n. \qquad \square$$

4 Price of Stability

We have seen that in some cases, like Example 7, there is an equilibrium that achieves the optimal social cost, even though there exists another, suboptimal, equilibrium. In other cases, like Example 1, we have seen that any equilibrium outcome will be far from optimal. We now investigate distinguishing these two cases.

Definition 5 (Price of Stability)**.** The *price of stability (PoS)* of a game with some objective function (e.g., social cost) is the ratio between the best objective function value in any equilibrium of the game and the optimal value of the objective function.

Since we tak social cost as our objective function and our equilibrium concept are pure Nash equilibria, we get that the price of stability is

$$\frac{\min_{\boldsymbol{s}\in \text{PNE}} \text{cost}(\boldsymbol{s})}{\min_{\boldsymbol{s}\in S^n} \text{cost}(\boldsymbol{s})}.$$

Note that the only component that changed from the price of anarchy is that the max in the numerator is now a min, since we consider the best equilibrium instead of the worst equilibrium.

Example 8 (Price of stability)**.** Let's revisit the cost sharing game from Example 7.

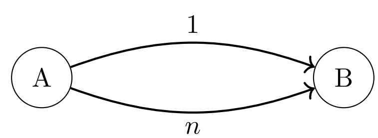

Previously, we identified that the strategy profile where all players take the path with cost 1 is the socially optimal solution. We further can see that this is a Nash equilibrium: If any player deviates unilaterally to the other path, their individual cost will increase from $1/n$ to n. Thus, the price of stability in this game is 1.

Example 9 (Lower bound on PoS for cost sharing games)**.** In the following cost sharing game with n players, each player i wishes to travel from node A_i to B.

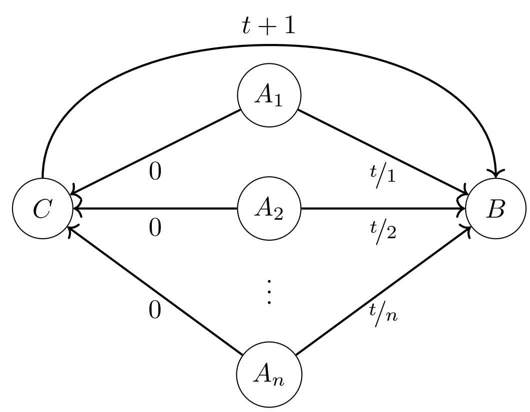

Every player i has the option of choosing the direct path for a cost of t/i, independently of the other players, or they can join the path via C that has a total cost of $t+1$. Here, every player taking their direct path is a Nash equilibrium since no individual player would benefit from paying $t+1$ by going through node C.

Surprisingly, this is the only pure Nash equilibrium. We show this by proving that no player can go through C in a pure Nash equilibrium. Assume towards a contradiction that in some pure Nash equilibrium, player i is the highest-index player going through C, i.e., all players j for $j > i$ use the direct path to B but player i goes through C. Thus, at most i players go through C, so the cost of player i is at least $t+1/i$. Player i has a beneficial deviation of instead going directly to B, for a cost of t/i—a contradiction to the original strategy profile being a Nash Equilibrium. We find that there does not exists a pure Nash Equilibrium with a player going through C.

The social cost of the unique pure Nash equilibrium is

$$\frac{t}{1}+\frac{t}{2}+\cdots+\frac{t}{n}=tH(n)$$

where $H(n)$ is the nth harmonic number. In contrast, when all n players go through x, the social cost is only $t+1$. Thus, as $t \to \infty$, we get cost sharing games with a price of stability as close as desired to $H(n)$.

It turns out that this lower bound is tight:

Theorem 4. *The price of stability of any cost sharing game is at most $H(n)$.*

We prove this in the Appendix; reading this proof is fully optional.

In summary, we found that for cost sharing games the worst-case PoA is exactly n and the worst-case PoS is exactly $H(n) = \Theta(\log n)$.

Appendix: Optional Proofs

This appendix contains proofs omitted from the lecture. Reading and understanding these is fully optional—knowing them will not be required to answer problem set or exam questions. You are more than welcome to ask any questions you may have about these proofs in Office Hours (as, of course, holds true for any part of the course material).

Theorem 1. *Routing games are exact potential games.*

Proof. Define the potential function

$$\Phi(\boldsymbol{s})=\sum_{e\in E}\sum_{k=1}^{n_e(\mathbf{s})} c_e(k)$$

Suppose that in some strategy profile $\boldsymbol{s}$, player i deviates from path s_i to path s_i'. Then, the difference $\text{cost}_i(s_i', s_{-i}) - \text{cost}_i(\mathbf{s})$ is

$$\sum_{e\in s_i'\setminus s_i} c_e(n_e(\mathbf{s})+1) - \sum_{e\in s_i\setminus s_i'} c_e(n_e(\mathbf{s})). \tag{1}$$

The first term is the new cost due to the edges added to the path of player i, denoted $s_i' \setminus s_i$; the second term is the reduced cost due to the edges removed from the path of player i, denoted $s_i \setminus s_i'$.

We want to show that this is precisely $\Phi(s_i', s_{-i}) - \Phi(\mathbf{s})$. This follows because when player i switches from s_i to s_i', only the terms corresponding to edges that are in s_i' and not in s_i, and the edges that are in s_i but not in s_i', are affected. For the edges $e \in s_i' \setminus s_i$, the number of players using it increases by 1, i.e., $n_e(s_i', \boldsymbol{s}_{-1}) = n_e(\boldsymbol{s}) + 1$, so that $c_e(n_e(\boldsymbol{s}+1))$ gets added to the sum for such e. Further, for the edges $e \in s_i \setminus s_i'$,the number of players using it decreases by 1, so that $c_e(n_e(\boldsymbol{s}))$ gets subtracted from the sum for such e. This gives exactly Equation (1). □

Theorem 4. *The price of stability of any cost sharing game is at most $H(n)$.*

Proof. Like routing games, cost sharing games are exact potential games with potential function

$$\Phi(\boldsymbol{s})=\sum_{e\in E}\sum_{k=1}^{n_e(\mathbf{s})}\frac{c_e}{k},$$

the proof of this is analogous to above.

Observe that $\text{cost}(\boldsymbol{s}) \le \Phi(\boldsymbol{s}) \le H(n)\cdot\text{cost}(\boldsymbol{s})$. Let $\boldsymbol{s}^\star$ be a pure strategy profile that minimizes Φ; as described earlier we know that $\boldsymbol{s}^\star$ is a Nash equilibrium. We get that

$$\text{cost}(\boldsymbol{s}^\star) \le \Phi(\boldsymbol{s}^\star) \le \Phi(\text{OPT}) \le H(n)\cdot\text{cost}(\text{OPT})$$

as desired. □