Congestion Games

Example: Network Routing

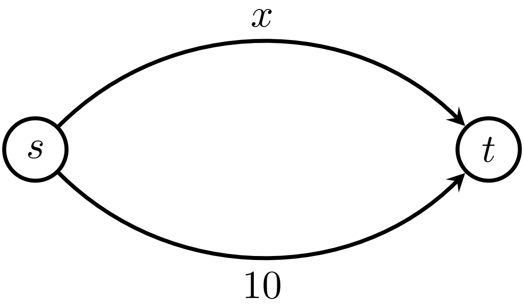

- $n = 10$ players want to travel from s to t
- Each edge e is labeled with its (flow-dependent) delay function $d_e(x)$

Question: *What are the pure Nash equilibria here?*

Example: Network Routing

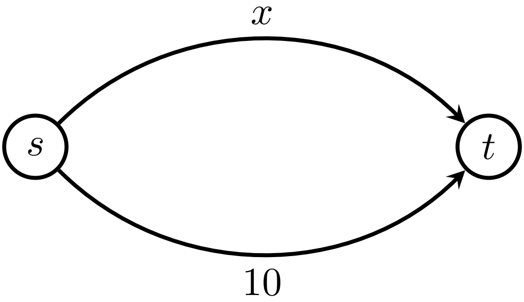

- $n = 10$ players want to travel from s to t
- Each edge e is labeled with its (flow-dependent) delay function $d_e(x)$

Question: *What are the pure Nash equilibria here?*

Example: Network Routing

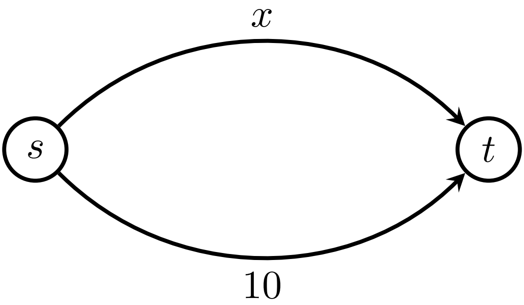

- $n = 10$ players want to travel from s to t
- Each edge e is labeled with its (flow-dependent) delay function $d_e(x)$

Question: *What are the pure Nash equilibria here?*

Example: Network Routing

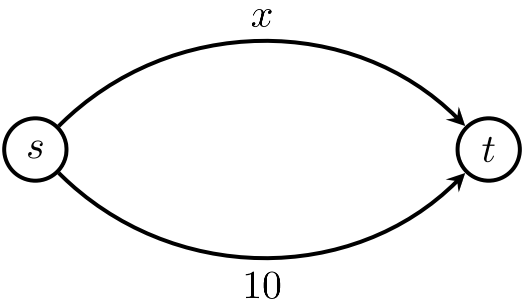

- $n = 10$ players want to travel from s to t
- Each edge e is labeled with its (flow-dependent) delay function $d_e(x)$

Question: *What are the pure Nash equilibria here?*

Example: Network Routing

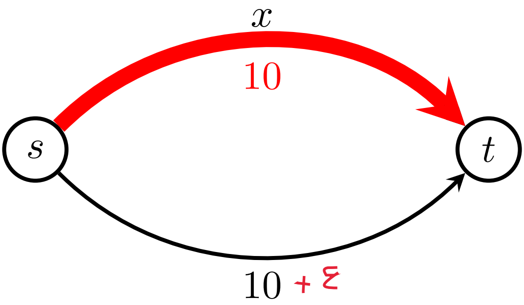

- $n = 10$ players want to travel from s to t
- Each edge e is labeled with its (flow-dependent) delay function $d_e(x)$

Question: *What are the pure Nash equilibria here?* $(10, 0)$

Example: Network Routing

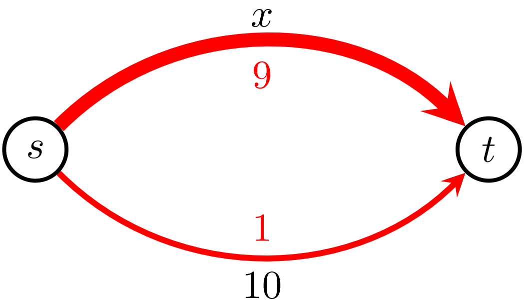

- $n = 10$ players want to travel from s to t
- Each edge e is labeled with its (flow-dependent) delay function $d_e(x)$

Question: *What are the pure Nash equilibria here?* $(10, 0)$, $(9, 1)$

Example: The El Farol Bar Problem

100 people consider visiting the *El Farol* Bar on a Thursday night. They all have identical preferences:

- If 60 or more people show up, it's nicer to be at home.
- If fewer than 60 people show up, it's nicer to be at the bar.

Question: *What are the pure Nash equilibria here?*

Example: The El Farol Bar Problem

100 people consider visiting the *El Farol* Bar on a Thursday night. They all have identical preferences:

- If 60 or more people show up, it's nicer to be at home.
- If fewer than 60 people show up, it's nicer to be at the bar.

Question: *What are the pure Nash equilibria here?*

Example: The El Farol Bar Problem

100 people consider visiting the *El Farol* Bar on a Thursday night. They all have identical preferences:

- If 60 or more people show up, it's nicer to be at home.
- If fewer than 60 people show up, it's nicer to be at the bar.

Question: *What are the pure Nash equilibria here?*

Example: The El Farol Bar Problem

100 people consider visiting the *El Farol* Bar on a Thursday night. They all have identical preferences:

- If 60 or more people show up, it's nicer to be at home.
- If fewer than 60 people show up, it's nicer to be at the bar.

Question: *What are the pure Nash equilibria here?* 59 people at the bar.

Example: Congestion Game

A congestion game is a tuple $\langle N, R, \boldsymbol{A}, \boldsymbol{d} \rangle$, where

- $N = \{1, \ldots, n\}$ is a finite set of players
- $R = \{1, \ldots, m\}$ is a finite set of resources
- $\boldsymbol{A} = A_1 \times \cdots \times A_n$ is a finite set of action profiles $\boldsymbol{a} = (a_1, \ldots, a_n)$, with $A_i \subseteq 2^R \setminus \{\emptyset\}$ being the set of actions available to player i
- $\boldsymbol{d} = (d_1, \ldots, d_m)$ is a vector of delay functions $d_r : \mathbb{N} \to \mathbb{R}$.

Goal: player $i \in N$ chooses a subset of resources $a_i \in A_i$

Given an action profile $\boldsymbol{a} = (a_1, \ldots, a_n)$, the cost of player i is

$$c_i(\boldsymbol{a}) = \sum_{r \in a_i} d_r(n_r(\boldsymbol{a})) \quad \text{where} \quad n_r(\boldsymbol{a}) = |\{i \in N \,:\, r \in a_i\}|.$$

Note: $u_i(\boldsymbol{a}) = -c_i(\boldsymbol{a})$ for every i here.

Example: Congestion Game

A congestion game is a tuple $\langle N, R, \boldsymbol{A}, \boldsymbol{d} \rangle$, where

- $N = \{1, \dots, n\}$ is a finite set of players
- $R = \{1, \dots, m\}$ is a finite set of resources
- $\boldsymbol{A} = A_1 \times \cdots \times A_n$ is a finite set of action profiles $\boldsymbol{a} = (a_1, \dots, a_n)$, with $A_i \subseteq 2^R \setminus \{\emptyset\}$ being the set of actions available to player i
- $\boldsymbol{d} = (d_1, \dots, d_m)$ is a vector of delay functions $d_r : \mathbb{N} \to \mathbb{R}$.

Goal: player $i \in N$ chooses a subset of resources $a_i \in A_i$

Given an action profile $\boldsymbol{a} = (a_1, \dots, a_n)$, the cost of player i is

$$c_i(\boldsymbol{a}) = \sum_{r \in a_i} d_r(n_r(\boldsymbol{a})) \quad \text{where} \quad n_r(\boldsymbol{a}) = |\{i \in N \ : \ r \in a_i\}|.$$

$c_1(a) = d_{r(1)}(\#\text{users})$

Note: $u_i(\boldsymbol{a}) = -c_i(\boldsymbol{a})$ for every i here.

Example: Congestion Game

A congestion game is a tuple $\langle N, R, \boldsymbol{A}, \boldsymbol{d} \rangle$, where

- $N = \{1, \ldots, n\}$ is a finite set of players
- $R = \{1, \ldots, m\}$ is a finite set of resources
- $\boldsymbol{A} = A_1 \times \cdots \times A_n$ is a finite set of action profiles $\boldsymbol{a} = (a_1, \ldots, a_n)$, with $A_i \subseteq 2^R \setminus \{\emptyset\}$ being the set of actions available to player i
- $\boldsymbol{d} = (d_1, \ldots, d_m)$ is a vector of delay functions $d_r : \mathbb{N} \to \mathbb{R}$.

Goal: player $i \in N$ chooses a subset of resources $a_i \in A_i$

Given an action profile $\boldsymbol{a} = (a_1, \ldots, a_n)$, the cost of player i is

$$c_i(\boldsymbol{a}) = \sum_{r \in a_i} d_r(n_r(\boldsymbol{a})) \quad \text{where} \quad n_r(\boldsymbol{a}) = |\{i \in N \ : \ r \in a_i\}|.$$

Note: $u_i(\boldsymbol{a}) = -c_i(\boldsymbol{a})$ for every i here.

Modelling the Examples

Congestion Game:

- players $N = \{1, 2, \ldots, 10\}$
- resources $R = \{\uparrow, \downarrow\}$
- action spaces $A_i = \{\{\uparrow\}, \{\downarrow\}\}$ representing the two routes
- delay functions $d_\uparrow : x \mapsto x$ and $d_\downarrow : x \mapsto 10$

El Farol Bar Problem:

- players $N = \{1, 2, \ldots, 100\}$
- resources $R = \{\text{🍸}, \text{🏠}_1, \text{🏠}_2, \ldots, \text{🏠}_{100}\}$
- action spaces $A_i = \{\{\text{🍸}\}, \{\text{🏠}_i\}\}$
- delay functions $d_{\text{🍸}} : x \mapsto \mathbb{1}_{x \geqslant 60}$ and $d_{\text{🏠}_i} : x \mapsto \frac{1}{2}$

Remark: neither example uses the full generality of congestion games (actions correspond to singleton resource sets only)

Modelling the Examples

Congestion Game:

- players $N = \{1, 2, \ldots, 10\}$
- resources $R = \{\uparrow, \downarrow\}$
- action spaces $A_i = \{\{\uparrow\}, \{\downarrow\}\}$ representing the two routes
- delay functions $d_\uparrow : x \mapsto x$ and $d_\downarrow : x \mapsto 10$

El Farol Bar Problem:

- players $N = \{1, 2, \ldots, 100\}$
- resources $R = \{🍸, 🏠_1, 🏠_2, \ldots, 🏠_{100}\}$
- action spaces $A_i = \{\{🍸\}, \{🏠_i\}\}$
- delay functions $d_{🍸} : x \mapsto \mathbb{1}_{x \geqslant 60}$ and $d_{🏠_i} : x \mapsto \frac{1}{2}$

Remark: neither example uses the full generality of congestion games (actions correspond to singleton resource sets only)

Modelling the Examples

Congestion Game:

- players $N = \{1, 2, \ldots, 10\}$
- resources $R = \{\uparrow, \downarrow\}$
- action spaces $A_i = \{\{\uparrow\}, \{\downarrow\}\}$ representing the two routes
- delay functions $d_\uparrow : x \mapsto x$ and $d_\downarrow : x \mapsto 10$

El Farol Bar Problem:

- players $N = \{1, 2, \ldots, 100\}$
- resources $R = \{\faGlass, \faHome_1, \faHome_2, \ldots, \faHome_{100}\}$
- action spaces $A_i = \{\{\faGlass\}, \{\faHome_i\}\}$
- delay functions $d_{\faGlass} : x \mapsto \mathbb{1}_{x \geqslant 60}$ and $d_{\faHome_i} : x \mapsto \frac{1}{2}$

Remark: neither example uses the full generality of congestion games (actions correspond to singleton resource sets only)

Existence of Pure Nash Equilibria

Good news:

Theorem (Rosenthal, 1973)

Every congestion game has at least one pure Nash equilibrium.

R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

Inefficiency of Equilibria

Inefficiency of Equilibria

Prisoner's Dilemma: Nash equilibria might be suboptimal!

Question: Can we quantify how "bad" Nash equilibria are?

Social cost: define the social cost of strategy profile $\boldsymbol{a}$ as

$$SC(\boldsymbol{a}) = \sum_{i \in N} c_i(\boldsymbol{a})$$

→ let $\boldsymbol{a}^*$ be a strategy profile minimizing $SC(\cdot)$ (social optimum)
→ $\boldsymbol{a}^*$ is best possible outcome if one could coordinate the players

Note: consider social welfare $SW = \sum_i u_i$ for utility maximizing players

Idea: measure worst case loss in social cost due to lack of coordination

$$\text{POA} = \max_{\boldsymbol{a} \in \text{PNE}} \frac{SC(\boldsymbol{a})}{SC(\boldsymbol{a}^*)}$$

→ termed the price of anarchy by Koutsoupias & Papadimitriou (1999)

Inefficiency of Equilibria

Prisoner's Dilemma: Nash equilibria might be suboptimal!

Question: Can we quantify how "bad" Nash equilibria are?

Social cost: define the social cost of strategy profile $\boldsymbol{a}$ as

$$SC(\boldsymbol{a}) = \sum_{i \in N} c_i(\boldsymbol{a})$$

→ let $\boldsymbol{a}^*$ be a strategy profile minimizing $SC(\cdot)$ (social optimum)
→ $\boldsymbol{a}^*$ is best possible outcome if one could coordinate the players

Note: consider social welfare $SW = \sum_i u_i$ for utility maximizing players

Idea: measure worst case loss in social cost due to lack of coordination

$$\text{POA} = \max_{\boldsymbol{a} \in \text{PNE}} \frac{SC(\boldsymbol{a})}{SC(\boldsymbol{a}^*)}$$

→ termed the price of anarchy by Koutsoupias & Papadimitriou (1999)

Inefficiency of Equilibria

Prisoner's Dilemma: Nash equilibria might be suboptimal!

Question: Can we quantify how "bad" Nash equilibria are?

Social cost: define the social cost of strategy profile $\boldsymbol{a}$ as

$$SC(\boldsymbol{a}) = \sum_{i \in N} c_i(\boldsymbol{a})$$

→ let $\boldsymbol{a}^*$ be a strategy profile minimizing $SC(\cdot)$ (social optimum)
→ $\boldsymbol{a}^*$ is best possible outcome if one could coordinate the players

Note: consider social welfare $SW = \sum_i u_i$ for utility maximizing players

Idea: measure worst case loss in social cost due to lack of coordination

$$\text{POA} = \max_{\boldsymbol{a} \in \text{PNE}} \frac{SC(\boldsymbol{a})}{SC(\boldsymbol{a}^*)}$$

→ termed the price of anarchy by Koutsoupias & Papadimitriou (1999)

Inefficiency of Equilibria

Prisoner's Dilemma: Nash equilibria might be suboptimal!

Question: Can we quantify how "bad" Nash equilibria are?

Social cost: define the social cost of strategy profile $\boldsymbol{a}$ as

$$SC(\boldsymbol{a}) = \sum_{i \in N} c_i(\boldsymbol{a})$$

→ let $\boldsymbol{a}^*$ be a strategy profile minimizing $SC(\cdot)$ (social optimum)
→ $\boldsymbol{a}^*$ is best possible outcome if one could coordinate the players

Note: consider social welfare $SW = \sum_i u_i$ for utility maximizing players

Idea: measure worst case loss in social cost due to lack of coordination

$$\text{POA} = \max_{\boldsymbol{a} \in \text{PNE}} \frac{SC(\boldsymbol{a})}{SC(\boldsymbol{a}^*)}$$

→ termed the price of anarchy by Koutsoupias & Papadimitriou (1999)

Inefficiency of Equilibria

Prisoner's Dilemma: Nash equilibria might be suboptimal!

Question: Can we quantify how "bad" Nash equilibria are?

Social cost: define the social cost of strategy profile $\boldsymbol{a}$ as

$$SC(\boldsymbol{a}) = \sum_{i \in N} c_i(\boldsymbol{a})$$

$\rightarrow$ let $\boldsymbol{a}^*$ be a strategy profile minimizing $SC(\cdot)$ (social optimum)
$\rightarrow$ $\boldsymbol{a}^*$ is best possible outcome if one could coordinate the players

Note: consider social welfare $SW = \sum_i u_i$ for utility maximizing players

Idea: measure worst case loss in social cost due to lack of coordination

$$\text{POA} = \max_{\boldsymbol{a} \in \text{PNE}} \frac{SC(\boldsymbol{a})}{SC(\boldsymbol{a}^*)}$$

$\rightarrow$ termed the price of anarchy by Koutsoupias & Papadimitriou (1999)

Inefficiency of Equilibria

Prisoner's Dilemma: Nash equilibria might be suboptimal!

Question: Can we quantify how "bad" Nash equilibria are?

Social cost: define the social cost of strategy profile $\boldsymbol{a}$ as

$$SC(\boldsymbol{a}) = \sum_{i \in N} c_i(\boldsymbol{a})$$

$\rightarrow$ let $\boldsymbol{a}^*$ be a strategy profile minimizing $SC(\cdot)$ (social optimum)
$\rightarrow$ $\boldsymbol{a}^*$ is best possible outcome if one could coordinate the players

Note: consider social welfare $SW = \sum_i u_i$ for utility maximizing players

- **Idea:** measure worst case loss in social cost due to lack of coordination

$$\text{POA} = \max_{\boldsymbol{a} \in \text{PNE}} \frac{SC(\boldsymbol{a})}{SC(\boldsymbol{a}^*)}$$

$\rightarrow$ termed the price of anarchy by Koutsoupias & Papadimitriou (1999)

Network Routing Games

$$n = 10$$

$$S_i = \{\uparrow, \downarrow\}$$

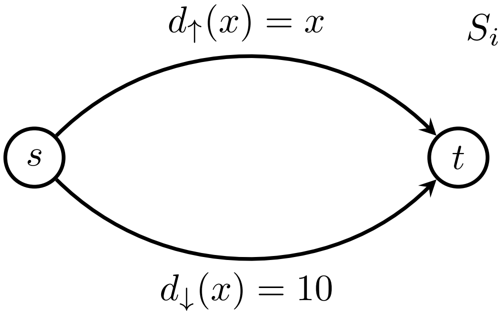

Network Routing Games

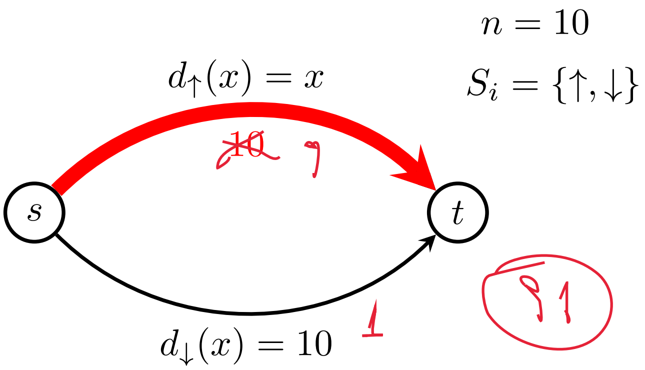

Nash equilibrium: $SC(s) = 10 \cdot 10 = 100$

Network Routing Games

$n = 10$

$S_i = \{\uparrow, \downarrow\}$

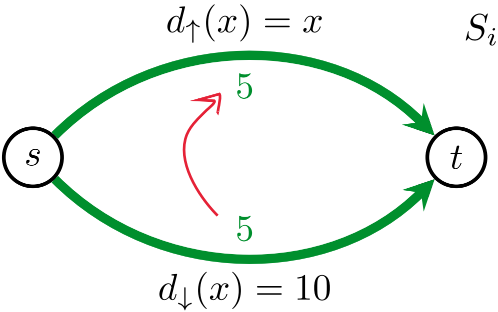

social optimum: $SC(s^*) = 5 \cdot 5 + 5 \cdot 10 = 75$

Network Routing Games

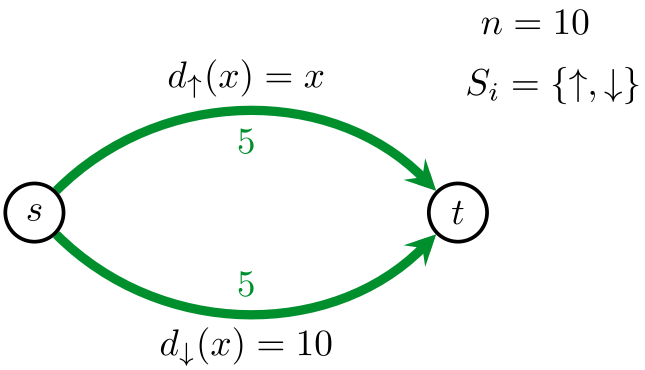

price of anarchy: $\frac{SC(s)}{SC(s^*)} = \frac{100}{75} = \frac{4}{3}$

Good news:

- POA is independent of the network structure
- POA depends on the class of delay functions
  - → POA $= \frac{5}{2}$ for affine functions (proof on next slide)
  - → POA $= O(1)$ for quadratic, cubic, . . . functions

$d_e(x) = ax + b$

Bad news:

- POA increases with the "steepness" of delay functions
- in general: unbounded!
- even worse: unbounded for practically relevant delay functions

Price of Anarchy for Congestion Games

Good news:

- POA is independent of the network structure
- POA depends on the class of delay functions
  - → POA $= \frac{5}{2}$ for affine functions (proof on next slide)
  - → POA $= O(1)$ for quadratic, cubic, … functions

Bad news:

- POA increases with the "steepness" of delay functions
- in general: unbounded!
- even worse: unbounded for practically relevant delay functions

Theorem

The price of anarchy of congestion games with affine delay functions is $\frac{5}{2}$.

Proof: All delay functions are of the form $d_r(x) = px + q$ with $p, q \in \mathbb{Z}_{\geq 0}$.
Can assume wlog that $d_r(x) = x \; \forall r$ (think about it!).
Let $\boldsymbol{a}$ be a PNE and let $\boldsymbol{a}^*$ be a social optimum. We have

See proof at the end!

$$SC(\boldsymbol{a}) = \sum_i c_i(a_i, \boldsymbol{a}_{-i}) \leq \sum_i c_i(a_i^*, \boldsymbol{a}_{-i}) = \sum_i \sum_{r \in a_i^*} d_r(n_r(a_i^*, \boldsymbol{a}_{-i}))$$

$$= \sum_i \sum_{r \in a_i^*} n_r(a_i^*, \boldsymbol{a}_{-i}) \leq \sum_i \sum_{r \in a_i^*} (n_r(\boldsymbol{a}) + 1)$$

$$= \sum_{r \in R} (n_r(\boldsymbol{a}) + 1) \sum_{i : r \in a_i^*} 1 = \sum_{r \in R} n_r(\boldsymbol{a}^*)(n_r(\boldsymbol{a}) + 1)$$

$$\leq \frac{5}{3} \sum_{r \in R} (n_r(\boldsymbol{a}^*))^2 + \frac{1}{3} \sum_{r \in R} (n_r(\boldsymbol{a}))^2 = \frac{5}{3} SC(\boldsymbol{a}^*) + \frac{1}{3} SC(\boldsymbol{a}) \quad \square$$

Tight Example

Instance:

- $N = [3]$
- $R = R_1 \cup R_2$, where $R_1 = \{h_1, h_2, h_3\}$ and $R_2 = \{g_1, g_2, g_3\}$
- delay function $d_r(x) = x$ for every $r \in R$
- each player i has two strategies: $\{h_i, g_i\}$ and $\{h_{i-1}, h_{i+1}, g_{i+1}\}$ (modulo 3).

Social optimum: every player selects his first strategy: $SC(\boldsymbol{a}^*) = 6$

Nash equilibrium: every player chooses his second strategy: $SC(\boldsymbol{a}) = \sum_i c_i(\boldsymbol{a}) = 3 \cdot 5 = 15$

Tight Example

Instance:

- $N = [3]$
- $R = R_1 \cup R_2$, where $R_1 = \{h_1, h_2, h_3\}$ and $R_2 = \{g_1, g_2, g_3\}$
- delay function $d_r(x) = x$ for every $r \in R$
- each player i has two strategies: $\{h_i, g_i\}$ and $\{h_{i-1}, h_{i+1}, g_{i+1}\}$ (modulo 3).

Social optimum: every player selects his first strategy: $SC(\boldsymbol{a}^*) = 6$

Nash equilibrium: every player chooses his second strategy: $SC(\boldsymbol{a}) = \sum_i c_i(\boldsymbol{a}) = 3 \cdot 5 = 15$

Tight Example

Instance:

- $N = [3]$
- $R = R_1 \cup R_2$, where $R_1 = \{h_1, h_2, h_3\}$ and $R_2 = \{g_1, g_2, g_3\}$
- delay function $d_r(x) = x$ for every $r \in R$
- each player i has two strategies: $\{h_i, g_i\}$ and $\{h_{i-1}, h_{i+1}, g_{i+1}\}$ (modulo 3).

Social optimum: every player selects his first strategy: $SC(\boldsymbol{a}^*) = 6$

Nash equilibrium: every player chooses his second strategy:
$SC(\boldsymbol{a}) = \sum_i c_i(\boldsymbol{a}) = 3 \cdot 5 = 15$

Unbounded POA

In general, POA for congestion games is unbounded.

Question: Can we improve the POA?
→ **natural idea:** infrastructure improvement (add new edges)

Unbounded POA

In general, POA for congestion games is unbounded.

Question: Can we improve the POA?
→ **natural idea:** infrastructure improvement (add new edges)

Unbounded POA

In general, POA for congestion games is unbounded.

Question: Can we improve the POA?
$\rightarrow$ **natural idea:** infrastructure improvement (add new edges)

Braess Paradox

$n = 10$

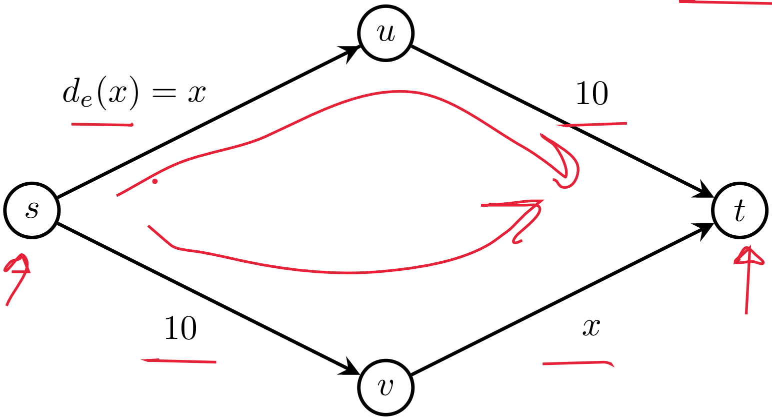

Braess Paradox

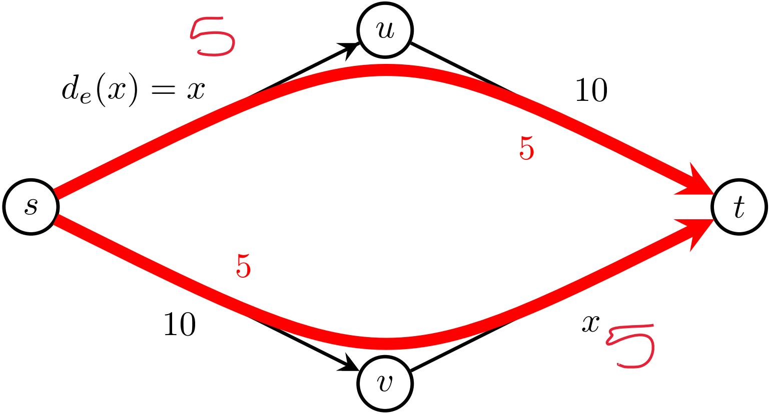

Braess Paradox

$n = 10$

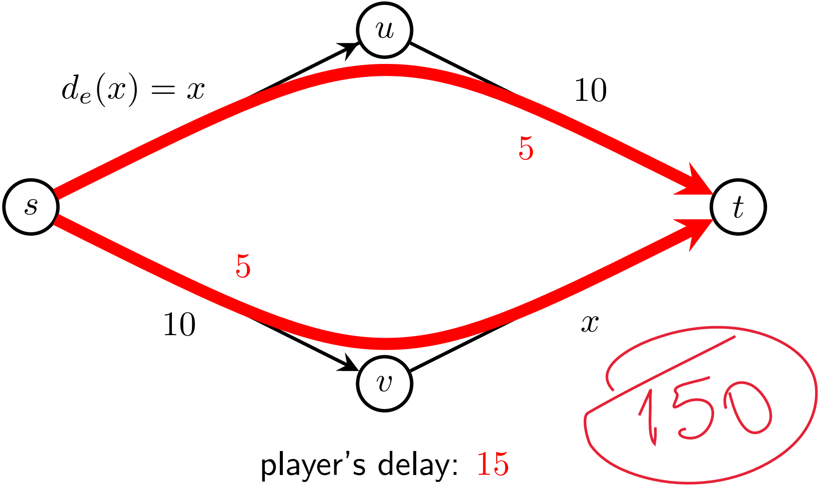

player's delay: 15

Braess Paradox

$n = 10$

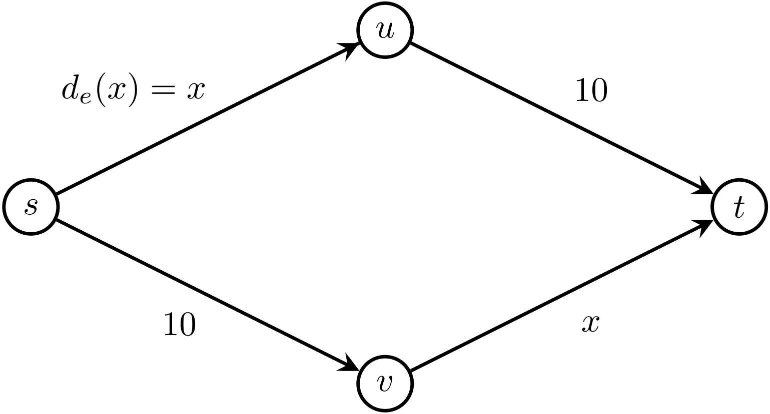

Braess Paradox

$n = 10$

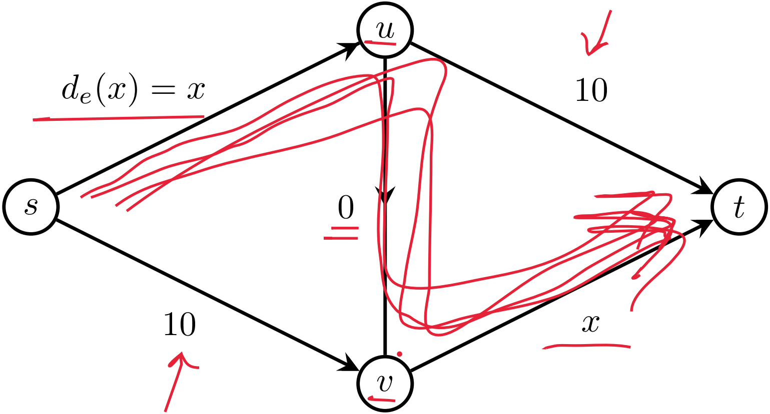

Braess Paradox

•

$n = 10$

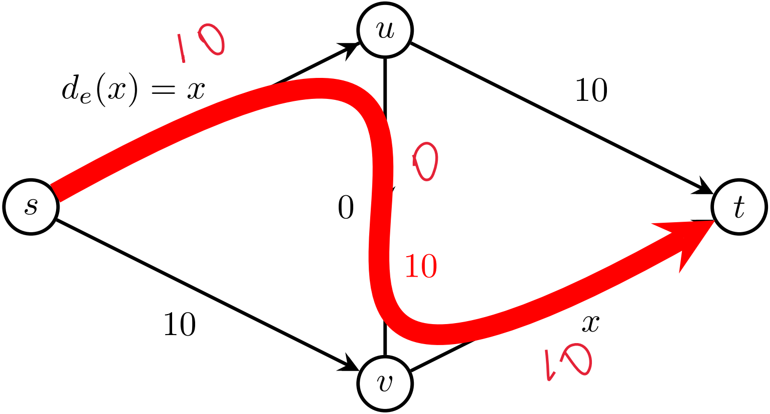

Braess Paradox

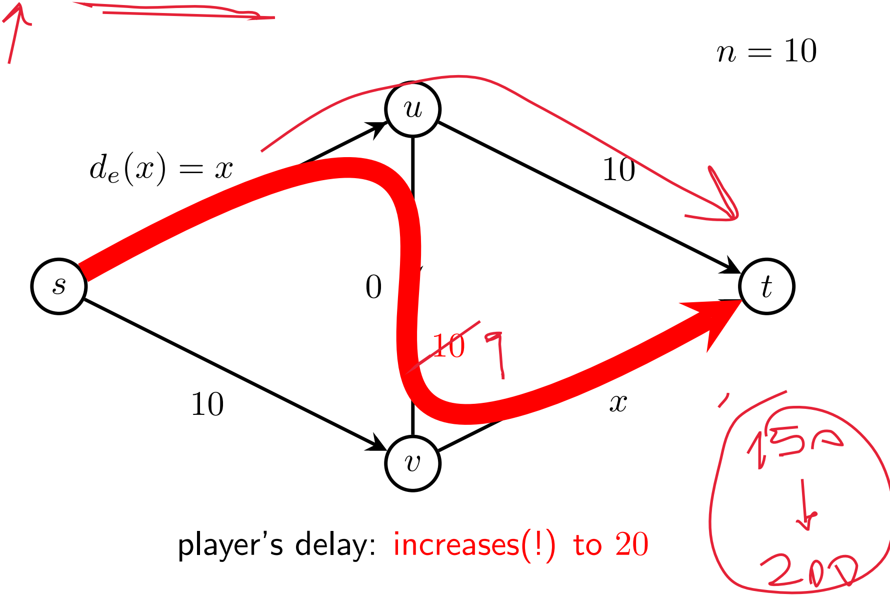

player's delay: increases(!) to 20

The New York Times
nytimes.com

December 25, 1990

What if They Closed 42d Street and Nobody Noticed?

By GINA KOLATA

ON Earth Day this year, New York City's Transportation Commissioner decided to close 42d Street, which as every New Yorker knows is always congested. "Many predicted it would be doomsday," said the Commissioner, Lucius J. Riccio. "You didn't need to be a rocket scientist or have a sophisticated computer queuing model to see that this could have been a major problem."

But to everyone's surprise, Earth Day generated no historic traffic jam. Traffic flow actually improved when 42d Street was closed.

To mathematicians, this may be a real-world example of Braess's paradox, a statistical theorem that holds that when a network of streets is already jammed with vehicles, adding a new street can make traffic flow even more slowly. Seeking Out a New Street

The reason is that in crowded conditions, drivers will pile into a new street, clogging both it and the streets that provide access to it. By the same token, removing a major thoroughfare may actually ease congestion on the streets that normally provide access to it. And because other major streets are already overcrowded, diverting still more traffic to them may not make much difference.

The New York Times
nytimes.com

December 25, 1990

What if They Closed 42d Street and Nobody Noticed?

By GINA KOLATA

ON Earth Day this year, New York City's Transportation Commissioner decided to close 42d Street, which as every New Yorker knows is always congested. "Many predicted it would be doomsday," said the Commissioner, Lucius J. Riccio. "You didn't need to be a rocket scientist or have a sophisticated computer queuing model to see that this could have been a major problem."

But to everyone's surprise, Earth Day generated no historic traffic jam. Traffic flow actually improved when 42d Street was closed.

To mathematicians, this may be a real-world example of Braess's paradox, a statistical theorem that holds that when a network of streets is already jammed with vehicles, adding a new street can make traffic flow even more slowly. Seeking Out a New Street

The reason is that in crowded conditions, drivers will pile into a new street, clogging both it and the streets that provide access to it. By the same token, removing a major thoroughfare may actually ease congestion on the streets that normally provide access to it. And because other major streets are already overcrowded, diverting still more traffic to them may not make much difference.

The New York Times
nytimes.com

December 25, 1990

Dr. Joel E. Cohen, a mathematician at Rockefeller University in New York, says the paradox does not always hold; each traffic network must be analyzed on its own. When a network is not congested, adding a new street will indeed make things better. But in the case of congested networks, adding a new street probably makes things worse at least half the time, mathematicians say.

Dr. Cohen and Dr. Frank P. Kelly of the University of Cambridge in England published the most recent analysis of the traffic paradox in the current issue of The Journal of Applied Probability. In their paper, they show that the paradox occurs when the traffic is described by a sophisticated statistical model. Previous work had used what Dr. Cohen describes as an overly simple and less realistic model.

The traffic paradox was first described in 1968 by Dr. Dietrich Braess of the Institute for Numerical and Applied Mathematics in Munster, Germany. He found that when one street was added to a simple four-street network, all the vehicles took longer to get through.

Dr. Braess's result was "very surprising," said Dr. Richard Steinberg of A.T.&T.'s Bell Laboratories in Holmdel, N.J. Dr. Steinberg and colleagues studied how often the paradox would hold true, and determined in 1983 that "it is just as likely to occur as not."

He and his colleagues also turned up a paradox of their own: that in some situations, "when you add more delays along a route, more people use it." Honk, Honk

Dr. Cohen and Dr. Kelly have now examined traffic networks with a sophisticated analytic method known as queuing theory, which describes traffic jams in terms of vehicles lining up on the streets. They found a simple traffic network in which adding a street increased travel time.

The New York Times
nytimes.com

December 25, 1990

Dr. Joel E. Cohen, a mathematician at Rockefeller University in New York, says the paradox does not always hold; each traffic network must be analyzed on its own. When a network is not congested, adding a new street will indeed make things better. But in the case of congested networks, adding a new street probably makes things worse at least half the time, mathematicians say.

Dr. Cohen and Dr. Frank P. Kelly of the University of Cambridge in England published the most recent analysis of the traffic paradox in the current issue of The Journal of Applied Probability. In their paper, they show that the paradox occurs when the traffic is described by a sophisticated statistical model. Previous work had used what Dr. Cohen describes as an overly simple and less realistic model.

The traffic paradox was first described in 1968 by Dr. Dietrich Braess of the Institute for Numerical and Applied Mathematics in Munster, Germany. He found that when one street was added to a simple four-street network, all the vehicles took longer to get through.

Dr. Braess's result was "very surprising," said Dr. Richard Steinberg of A.T.&T.'s Bell Laboratories in Holmdel, N.J. Dr. Steinberg and colleagues studied how often the paradox would hold true, and determined in 1983 that "it is just as likely to occur as not."

He and his colleagues also turned up a paradox of their own: that in some situations, "when you add more delays along a route, more people use it." Honk, Honk

Dr. Cohen and Dr. Kelly have now examined traffic networks with a sophisticated analytic method known as queuing theory, which describes traffic jams in terms of vehicles lining up on the streets. They found a simple traffic network in which adding a street increased travel time.

December 25, 1990

added, "My intuition is that it must happen a lot." He finds support for his intuition from game theory, which says that if everyone in a game acts selfishly, everyone suffers. The traffic jam is analogous to such a game, Dr. Cohen said.

He cited a German paper, published in 1969, reporting that the City of Stuttgart had tried to ease downtown traffic by adding a new street. But congestion only got worse, and so, in desperation, the authorities closed the street. Traffic flow improved.

New York's Transportation Commissioner, Mr. Riccio, has a doctoral degree himself (in engineering, from Lehigh University), and he said he favored using mathematical models to try to improve traffic flow. "I believe in these models," he said, and added that he would welcome a call from Dr. Cohen to discuss how his work could apply to New York City's daunting traffic problems.

The New York Times
nytimes.com

December 25, 1990

added, "My intuition is that it must happen a lot." He finds support for his intuition from game theory, which says that if everyone in a game acts selfishly, everyone suffers. The traffic jam is analogous to such a game, Dr. Cohen said.

He cited a German paper, published in 1969, reporting that the City of Stuttgart had tried to ease downtown traffic by adding a new street. But congestion only got worse, and so, in desperation, the authorities closed the street. Traffic flow improved.

New York's Transportation Commissioner, Mr. Riccio, has a doctoral degree himself (in engineering, from Lehigh University), and he said he favored using mathematical models to try to improve traffic flow. "I believe in these models," he said, and added that he would welcome a call from Dr. Cohen to discuss how his work could apply to New York City's daunting traffic problems.

Price of Anarchy

So: In a congestion game, the natural better-response dynamics will always lead us to a pure NE. Nice. But: How good is that equilibrium?
Recall our traffic congestion example:

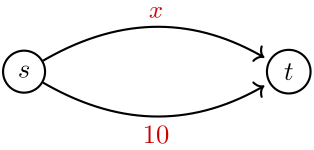

10 people overall
top delay = # on route
bottom delay = 10 minutes

If $x \leqslant 10$ players use top route, social welfare (sum of utilities) is:

$$sw(x) \;=\; -[x \cdot x + (10 - x) \cdot 10] \;=\; -[x^2 - 10x + 100]$$

This function is maximal for $x = 5$ and minimal for $x = 0$ and $x = 10$.
In equilibrium, 9 or 10 people will use the bottom route (10 is worse).
The so-called price of anarchy of this game is: $\frac{sw(10)}{sw(5)} = \frac{-100}{-75} = \frac{4}{3}$.
Thus: not perfect, but not too bad either (for this example).

Braess' Paradox

Something to think about. 10 people have to get from s to t:

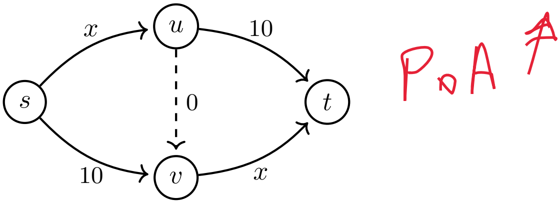

If the delay-free link from u to v is ***not*** present:

- In equilibrium, 5 people will use the top route s–u–t and 5 people the bottom route s–v–t. Everyone will take 15 minutes.

Now, if we add the delay-free link from u to v, this happens:

- In the worst equilibrium, everyone will take the route s–u–v–t and take 20 minutes! (Other equilibria are only slightly better.)

Price of Anarchy for Linear/Affine Congestion Games

Based on slides by A. Voudouris

A general technique for PoA bounds

↳ w/ pure NE

- Recall that a state $\boldsymbol{s} = (s_1, \dots, s_n)$ is an equilibrium if for each player i the strategy s_i minimizes her personal cost, given the strategies of the other players
- $\boldsymbol{s}_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ $\quad s = (s_i, s_{-i})$
- $\boldsymbol{s}$ is an equilibrium if for each player i, the strategy s_i is such that $\text{cost}_i(y, \boldsymbol{s}_{-i})$ is minimized for $y = s_i$
- Alternatively, for every possible strategy y of player i:

$$\text{cost}_i(s_i, \boldsymbol{s}_{-i}) \leq \text{cost}_i(y, \boldsymbol{s}_{-i})$$

- We have one such inequality for every player

A general technique for PoA bounds

- By adding these inequalities, we get

$$\mathrm{SC}(\boldsymbol{s}) = \sum_{i \in N} \mathrm{cost}_i(s_i, \boldsymbol{s}_{-i}) \leq \sum_{i \in N} \mathrm{cost}_i(y_i, \boldsymbol{s}_{-i})$$

- We can get an upper bound of λ on the price of anarchy if there exists a strategy y_i for every player i such that

$$\sum_{i \in N} \mathrm{cost}_i(y_i, \boldsymbol{s}_{-i}) \leq \lambda \cdot \mathrm{SC}(\boldsymbol{s}_{OPT})$$

- The goal is to pinpoint the strategy y_i for each player i, which will allow us to prove an inequality like this

Linear congestion games: PoA

Theorem

The price of anarchy of linear congestion games is at most 5/2

- $\boldsymbol{s} = (s_1, \dots s_n)$ is an equilibrium state
- $\boldsymbol{y} = (y_1, \dots, y_n)$ is an arbitrary state

$$\begin{aligned}
\mathrm{SC}(\boldsymbol{s}) &= \sum_{i \in N} \mathrm{cost}_i(s_i, \boldsymbol{s}_{-i}) \\
&\leq \sum_{i \in N} \mathrm{cost}_i(y_i, \boldsymbol{s}_{-i}) \\
&= \sum_{e \in E} \sum_{i \in N : e \in y_i} \left(a_e \cdot n_e(y_i, \boldsymbol{s}_{-i}) + b_e\right)
\end{aligned}$$

Linear congestion games: PoA

- $(y_i, \boldsymbol{s}_{-i})$ differs from $\boldsymbol{s} = (s_i, \boldsymbol{s}_{-i})$ only in the strategy of player i

  $\Rightarrow n_e(y_i, \boldsymbol{s}_{-i}) \leq n_e(\boldsymbol{s}) + 1$ for every resource $e \in E$

$$\mathrm{SC}(\boldsymbol{s}) \leq \sum_{e \in E} \sum_{i \in N: e \in y_i} (a_e \cdot n_e(y_i, \boldsymbol{s}_{-i}) + b_e)$$

$$\leq \sum_{e \in E} \sum_{i \in N: e \in y_i} (a_e \cdot (n_e(\boldsymbol{s}) + 1) + b_e)$$

appears as many times as players that use e, i.e., $n_e(y)$

Linear congestion games: PoA

- $(y_i, \boldsymbol{s}_{-i})$ differs from $\boldsymbol{s} = (s_i, \boldsymbol{s}_{-i})$ only in the strategy of player i

  ⇨ $n_e(y_i, \boldsymbol{s}_{-i}) \leq n_e(\boldsymbol{s}) + 1$ for every resource $e \in E$

.

$$\begin{aligned}
\mathrm{SC}(\boldsymbol{s}) &\leq \sum_{e \in E} \sum_{i \in N: e \in y_i} (a_e \cdot n_e(y_i, \boldsymbol{s}_{-i}) + b_e) \\
&\leq \sum_{e \in E} \sum_{i \in N: e \in y_i} (a_e \cdot (n_e(\boldsymbol{s}) + 1) + b_e) \\
&= \sum_{e \in E} n_e(\boldsymbol{y})(a_e \cdot (n_e(\boldsymbol{s}) + 1) + b_e) \\
&= \sum_{e \in E} \big(a_e \cdot n_e(\boldsymbol{y}) \cdot (n_e(\boldsymbol{s}) + 1) + b_e n_e(\boldsymbol{y})\big)
\end{aligned}$$

Linear congestion games: PoA

$n_e(y)\,(n_e(s)+1) \le \frac{1}{3}\,(5 \ldots$

- For every pair of integers $\gamma, \delta \geq 0$: $\gamma\,(\delta + 1) \leq \frac{1}{3}\,(5\gamma^2 + \delta^2)$
- Set $\gamma = n_e(\boldsymbol{y})$ and $\delta = n_e(\boldsymbol{s})$

$$\mathrm{SC}(\boldsymbol{s}) \leq \sum_{e \in E} \left(a_e \cdot n_e(\boldsymbol{y})(n_e(\boldsymbol{s}) + 1) + b_e n_e(\boldsymbol{y})\right)$$

$$\leq \sum_{e \in E} \left(a_e \cdot \frac{1}{3}(5 n_e(\boldsymbol{y})^2 + n_e(\boldsymbol{s})^2) + b_e n_e(\boldsymbol{y})\right)$$

Linear congestion games: PoA

- For every pair of integers $\gamma, \delta \geq 0$: $\gamma\,(\delta + 1) \leq \frac{1}{3}\,(5\gamma^2 + \delta^2)$
- Set $\gamma = n_e(\boldsymbol{y})$ and $\delta = n_e(\boldsymbol{s})$

$$\mathrm{SC}(\boldsymbol{s}) \leq \sum_{e \in E} \left(a_e \cdot n_e(\boldsymbol{y})(n_e(\boldsymbol{s}) + 1) + b_e n_e(\boldsymbol{y})\right)$$

$$\leq \sum_{e \in E} \left(a_e \cdot \frac{1}{3}(5 n_e(\boldsymbol{y})^2 + n_e(\boldsymbol{s})^2) + b_e n_e(\boldsymbol{y})\right)$$

$$= \sum_{e \in E} \left(\frac{5}{3} a_e n_e(\boldsymbol{y})^2 + b_e n_e(\boldsymbol{y})\right) + \frac{1}{3} \sum_{e \in E} a_e n_e(\boldsymbol{s})^2$$

? $+ b_e n_e(s)$

Linear congestion games: PoA

- For every pair of integers $\gamma, \delta \geq 0$: $\gamma\,(\delta + 1) \leq \frac{1}{3}\,(5\gamma^2 \;+ \delta^2)$
- Set $\gamma = n_e(\boldsymbol{y})$ and $\delta = n_e(\boldsymbol{s})$

$$\mathrm{SC}(\boldsymbol{s}) \leq \sum_{e \in E} \left(a_e \cdot n_e(\boldsymbol{y})(n_e(\boldsymbol{s}) + 1) + b_e n_e(\boldsymbol{y})\right)$$

$$\leq \sum_{e \in E} \left(a_e \cdot \frac{1}{3}(5 n_e(\boldsymbol{y})^2 + n_e(\boldsymbol{s})^2) + b_e n_e(\boldsymbol{y})\right)$$

$$= \sum_{e \in E} \left(\frac{5}{3} a_e n_e(\boldsymbol{y})^2 + b_e n_e(\boldsymbol{y})\right) + \frac{1}{3} \sum_{e \in E} a_e n_e(\boldsymbol{s})^2$$

SC(y) SC(s)

$$\leq \frac{5}{3} \sum_{e \in E} \left(a_e n_e(\boldsymbol{y})^2 + b_e n_e(\boldsymbol{y})\right) + \frac{1}{3} \sum_{e \in E} \left(a_e n_e(\boldsymbol{s})^2 + b_e n_e(\boldsymbol{s})\right)$$

$n_e(y)\left[a_e n_e(y) + b_e\right] = n_e(y) \cdot \mathrm{cost}_e(y)$

Linear congestion games: PoA

- Since

$$\mathrm{SC}(\boldsymbol{y}) = \sum_{e \in E} \left(a_e n_e(\boldsymbol{y})^2 + b_e n_e(\boldsymbol{y}) \right)$$

we obtain

$$\mathrm{SC}(\boldsymbol{s}) \leq \frac{5}{3}\mathrm{SC}(\boldsymbol{y}) + \frac{1}{3}\mathrm{SC}(\boldsymbol{s})$$

$$\Rightarrow \frac{\mathrm{SC}(\boldsymbol{s})}{\mathrm{SC}(\boldsymbol{y})} \leq \frac{5}{2}$$

How do we even find this??

- Since this holds for any $\boldsymbol{y}$, it also holds for $\boldsymbol{s}_{OPT}$ □

Can we do any better?

Theorem
The price of anarchy of linear congestion games is at least 5/2

- To show a lower bound, it suffices to construct a specific instance and prove that the social cost of the equilibrium is 5/2 times the optimal social cost

Can we do any better?

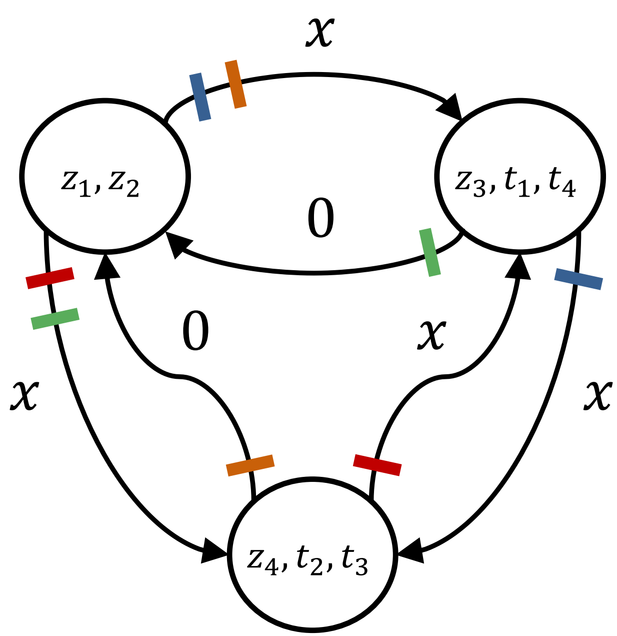

- Equilibrium: each player i uses two edges to connect z_i to t_i
- Players 1 and 2 (red, blue) have cost 3, while players 3 and 4 (green, orange) have cost 2
- By changing to the direct edge, all players would still have the same cost, so there is no reason for them to deviate

Can we do any better?

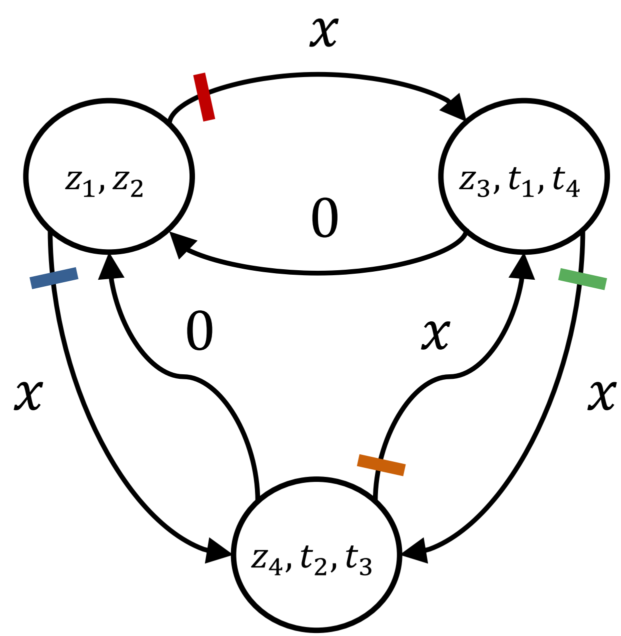

- Optimal: each player i uses the direct edge between z_i and t_i
- All players have cost 1
- SC(equilibrium) $= 10$ vs. SC(optimal) $= 4$ ⇨ PoA = 5/2

□