Chapter 4

Congestion Games

In this section, we introduce several classes of games that model the competition of players over a finite set of resources whose cost change with the demand. Traditionally, these games are phrased as cost minimization games. Minimization games are triplets $G = (N, \mathbf{S}, \boldsymbol{\pi})$ where $\boldsymbol{\pi} = (\pi_i)_{i \in N}$ and $\pi_i(\mathbf{s})$ is the private cost of player i in strategy profile $\mathbf{s}$.

In a congestion game, each player i strives to solve a combinatorial optimization problem of the form

$$\min \sum_{e \in A} c_e(\mathbf{s}) \qquad \text{s.t.} \qquad A \in \mathcal{A},$$

where $\mathcal{A} \subseteq 2^E$ is the set of feasible allocations for player i. The players' optimization problems are coupled via the cost functions of the resources as the cost c_e of resource e is a function of the strategy profile. In the most basic model, one assumes that the cost of a resource is a function of its users, i.e., $c_e(\mathbf{s}) = c_e(\mathbf{t})$ whenever $|\{j \in N : e \in s_j\}| = |\{j \in N : e \in t_j\}|$. Such coupled optimization problems are captured by the class of unweighted congestion games.

4.1 Unweighted Congestion Games

Definition 4.1 (Unweighted Congestion Game). For a finite set of resources E with cost functions $(c_e)_{e \in E}$, and an allocation vector $\boldsymbol{\mathcal{A}} = (\mathcal{A}_i)_{i \in N}$, the corresponding unweighted congestion game is the strategic game $G(\boldsymbol{\mathcal{A}}) = (N, \mathbf{S}, \boldsymbol{\pi})$, where

- N is a finite and non-empty set of n players,
- $S_i = \mathcal{A}_i$ for all players i,
- $\pi_i(\mathbf{s}) = \sum_{e \in s_i} c_e(x_e(\mathbf{s}))$ for all players i, where $x_e(\mathbf{s}) = |\{j \in N : e \in s_j\}|$.

Example 4.2. Consider the game in Figure 4.1. There are four resources e_1,e_2, e_3 and e_4 with cost functions defined as $c_{e_1}(x) = c_{e_4}(x) = 2x^3$ and $c_{e_2}(x) = c_{e_3}(x) = (x+1)^3$ for all $x \geq 0$. The feasible allocations of player 1 are $A = \{e_1, e_2\}$ and $A' = \{e_3, e_4\}$; for player 2 the feasible allocations are $B = \{e_1, e_3\}$ and $B' = \{e_2, e_4\}$.

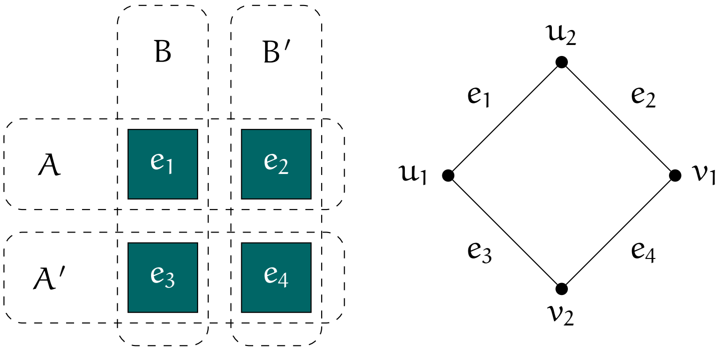

	B	B′
A	$(24, 24)$	$(29, 29)$
A′	$(29, 29)$	$(24, 24)$

(a) Resource representation **(b)** Graph representation **(c)** Strategic representation

Figure 4.1: Congestion Game

Setting $\mathcal{A} = (\mathcal{A}_i)_{i\in\{1,2\}}$ with $\mathcal{A}_1 = \{A, A'\}$ and $\mathcal{A}_2 = \{B, B'\}$ this defines a congestion game. Inspecting the strategic representation of the game in Figure 4.4c, we observe that the game has two pure Nash equilibria, i.e., (A, B) and (A', B').

Congestion games in which the feasible allocations of each player i can be represented as the set of paths between two designated vertices u_i and v_i of a given graph $D = (V, E)$ are of particular interest.

Definition 4.3 (Network Congestion Game). A congestion game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ is called a network congestion game if there is a graph $D = (V, E)$ and, for each player i, two designated vertices $u_i, v_i \in V$ such that

$$S_i = \{P \subseteq E : P \text{ is a } u_i, v_i\text{-path in } D.\}$$

for all players i.

The game from Example 4.2 obviously is a network congestion game as the graph representation in Figure 4.4b shows.

Example 4.4 (Braess' Paradox). A famous network congestion game is Braess paradox. There are 100 players striving to travel along a simple (u, v)-path in the network in Figure 4.2a. In the unique equilibrium, there are 50 players on each path resulting in a cost of 151 per player. After adding an additional edge with travel time 0, in the unique equilibrium the cost of all players is 200, see Figure 4.2b.

Theorem 4.5 (Rosenthal, 1973). Every unweighted congestion game has a pure Nash equilibrium.

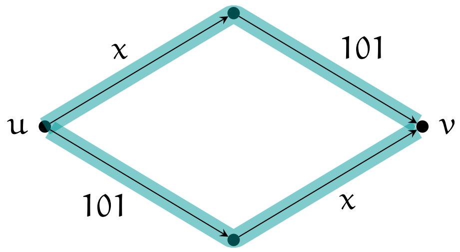

(a) Original graph

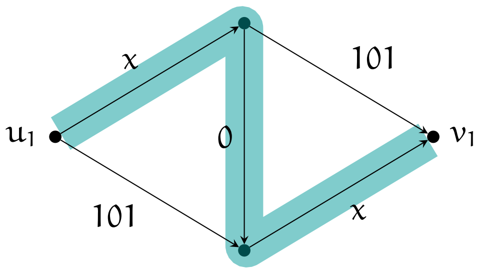

(b) Augmented graph

Figure 4.2: Breass' paradox

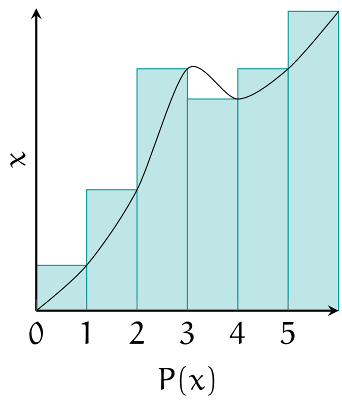

Figure 4.3: Potential function

Proof. Consider the function $P : \mathbf{S} \to \mathbb{R}$ defined as

$$P(\mathbf{s}) = \sum_{e \in E} \sum_{k=1}^{x_e(\mathbf{s})} c_e(k).$$

Intuitively, the function P corresponds to an approximation of the integral over the cost functions of the resources, see Figure 4.3.

Let $\mathbf{s} \in \mathbf{S}$ be arbitrary and let $i \in N$ and $t_i \in S_i$. We then obtain

$$\begin{aligned} P(t_i, \mathbf{s}_{-i}) &= \sum_{e \in E} \sum_{k=1}^{x_e(t_i, \mathbf{s}_{-i})} c_e(k) \\ &= \sum_{e \in E} \sum_{k=1}^{x_e(\mathbf{s})} c_e(k) + \sum_{e \in t_i \setminus s_i} c_e(x_e(\mathbf{s}) + 1) - \sum_{e \in s_i \setminus t_i} c_e(x_e(\mathbf{s})) \\ &= P(\mathbf{s}) + u_i(t_i, \mathbf{s}_{-i}) - u_i(\mathbf{s}) \end{aligned}$$

Thus, $P(t_i, \mathbf{s}_{-i}) - P(\mathbf{s}) = \pi_i(t_i, \mathbf{s}_{-i}) - \pi_i(\mathbf{s})$ for all $\mathbf{s} \in \mathbf{S}$, $i \in N$ and $t_i \in S_i$.
In particular, $P(t_i, \mathbf{s}_{-i}) < P(\mathbf{s})$ when $\pi_i(t_i, \mathbf{s}_{-i}) < \pi_i(\mathbf{s})$. Consider a sequence

$$\mathbf{s}^0, \mathbf{s}^1, \mathbf{s}^2, \ldots \tag{4.1}$$

where for all $l = 0, 1, \ldots$ there is a player i_l such that $s^{l+1} = (t_{i_l}, s^l_{-i_l})$ for some $t_{i_l} \in S_i$ and $\pi_{i_l}(t_{i_l}, s_{-i_l}) < \pi(s)$. Then, P is decreasing along the sequence and using that $\mathbf{S}$ is finite, every such sequence must be finite. The endpoint of such sequence is a pure Nash equilibrium. □

4.2 Exact Potential Games

Definition 4.6 (Exact Potential Game). A strategic game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ is called an exact potential game if there is a function $P : \mathbf{S} \to \mathbb{R}$ such that

$$P(t_i, s_{-i}) - P(s) = \pi_i(t_i, s_{-i}) - \pi_i(s)$$

for all $s \in \mathbf{S}$, $i \in N$ and $t_i \in S_i$.

Definition 4.7 (Coordination Game). A strategic game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ is a coordination game if $\pi_i(s) = \pi_j(s)$ for all $s \in \mathbf{S}$.

Definition 4.8 (Dummy Game). A strategic game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ is a dummy game if $pi_i(s_i, s_{-i}) = \pi_i(t_i, s_{-i})$ for all $i \in N$, $s \in \mathbf{S}$ and $t_i \in S_i$.

Theorem 4.9. A strategic game $G = (N, \mathbf{S}, \pi)$ is an exact potential game if and only if there are functions $\boldsymbol{\pi}^c : \mathbf{S} \to \mathbb{R}$ and $\boldsymbol{\pi}^d : \mathbf{S} \to \mathbb{R}$ such that $(N, \mathbf{S}, \boldsymbol{\pi}^c)$ is a coordination game, $(N, \mathbf{S}, \boldsymbol{\pi}^d)$ is a dummy game and $\pi(s) = \pi^c(s) + \pi^d(s)$ for all $s \in \mathbf{S}$.

Proof. "⇐": We claim that $P = \pi_i^c = \pi_j^C$ is an exact potential function. To see this, note that

$$\begin{aligned}\pi_i(t_i, s_{-i}) - \pi_i(s) &= \pi_i^c(t_i, s_{-i}) + \pi_i^d(t_i, s_{-i}) - \pi_i^c(s) - \pi_i^d(s)\\ &= \pi_i^c(t_i, s_{-i}) - \pi_i^c(s)\end{aligned}$$

for all $s \in S$, $i \in N$ and $t_i \in S_i$.
" ⇒": Let P be an exact potential function for G. For all players i and all strategy profiles s, let $\pi_i^c(s) = P(s)$ and let $\pi_i^c(s) = u_i(s) - P(s)$. Then, $G = (N, \mathbf{S}, \boldsymbol{\pi}_i^c)$ is a coordination game. Moreover,

$$\pi_i(s) - \pi_i(t_i, s_{-i}) = P(s) - P(t_i, s_{-i})$$

which is equivalent to

$$\pi_i(s) - P(s) = \pi_i(t_i, s_{-i}) - P(t_i, s_{-i}).$$

Thus, $G = (N, \mathbf{S}, \boldsymbol{\pi}^D)$ is a dummy game. □

Definition 4.10 (Isomorphism Between Games). Two games $G = (N, \mathbf{S}, \boldsymbol{\pi})$ and $G' = (N, \mathbf{S}', \boldsymbol{\pi}')$ are isomorphic if, for all $i \in N$, there is a bijection $\phi : S_i \to S_i'$ such that

$$\pi_i(\times_{i \in N} s_i) = \pi_i'(\times_{i \in N} \phi_i(s_i))$$

for all players i and strategy vectors $\mathbf{s} \in \mathbf{S}$.

Theorem 4.11. Every coordination game is isomorphic to a congestion game.

Proof. Let $G = (N, \mathbf{S}, \boldsymbol{\pi})$ be an n-player coordination game. We proceed to define a congestion game $G = (N\mathbf{S}', \boldsymbol{\pi}')$ isomorphic to G. For all $\mathbf{s} \in \mathbf{S}$, introduce a resource $e(\mathbf{s})$ with cost function

$$c_{e(\mathbf{s})}(x) = \begin{cases} \pi_1(\mathbf{s}), & \text{if } x = n \\ 0, & \text{else.} \end{cases}$$

Further, for a player i and a strategy $s_i \in S_i$, let

$$\phi_i(s_i) = \bigcup_{\mathbf{s}_{-i} \in \mathbf{S}_{-i}} \{e(s_i, \mathbf{s}_{-i})\}$$

and let

$$S_i' = \left\{ \phi_i(s_i) : s_i \in S_i \right\}$$

For all $\mathbf{s} \in \mathbf{S}$, let $\mathbf{s}' = \times_{i \in N} \phi_i(s_i)$. We obtain

$$\pi_i'(\mathbf{s}') = \sum_{e \in \phi(s_i)} c_e\big(x_e(\mathbf{s}')\big) = \sum_{e \in s_i' : x_e(\mathbf{s}') = n} c_e\big(x_e(\mathbf{s}')\big) = \pi_1(\mathbf{s}) = \pi_i(\mathbf{s}),$$

where the last equality used that G is a coordination game. To obtain an isomorphic congestion game we introduce a resource $e(\mathbf{s})$ for all strategy profiles $\mathbf{s} \in \mathbf{S}$, see Figure 4.4b. □

Example 4.12 (Congestion Game Isomorphic to a Coordination Game). Consider the coordination game G in Figure 4.4a. To obtain an isomorphic congestion game, we introduce a resource for each strategy profile $\mathbf{s}$ of G and define the players' strategies as in Figure 4.4b.

Theorem 4.13. Every dummy game is isomorphic to a congestion game.

Proof. Let $G = (N, \mathbf{S}, \boldsymbol{\pi})$ be a dummy game. For all players i and all $\mathbf{s}_{-i} \in \mathbf{S}_{-i}$, we introduce a resource $e(\mathbf{s}_{-i})$ with cost function

$$c_{e(\mathbf{s}_{-i})}(x) = \begin{cases} \pi_i(s_i, \mathbf{s}_{-i}) & \text{if } x = 1 \\ 0, & \text{else,} \end{cases}$$

	s_2	s_2'
s_1	$(0,0)$	$(1,1)$
s_1'	$(2,2)$	$(3,3)$

(a) Coordination game

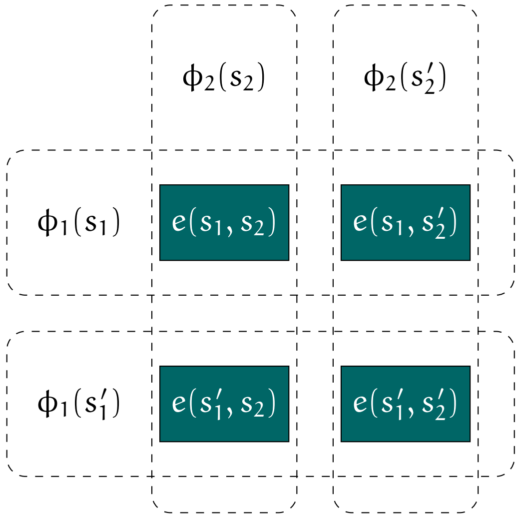

(b) Isomorphic congestion game

where $s_i \in S_i$ is arbitrary. Let E be the set of resources introduced this way. Further, for a player i with strategy $s_i \in S_i$, let

$$\phi_i(s_i) = \bigcup\nolimits_{\mathbf{t}_{-i} \in \mathbf{S}_{-i}} \left\{e(\mathbf{t}_{-i})\right\} \cup \bigcup\nolimits_{j \in N \setminus \{i\}} \bigcup\nolimits_{\mathbf{t}_{-j} \in \mathbf{S}_{-j} : t_i \neq s_i} \left\{e(\mathbf{t}_{-j})\right\}$$

and let

$$S_i' = \left\{\phi(s_i) : s_i \in S_i\right\}$$

Let $i \in N$, $s_i \in S_i$ and $\mathbf{s}_{-i} \in \mathbf{S}_{-i}$ be arbitrary and consider the strategy profile $\mathbf{s} = (s_i, \mathbf{s}_{-i})$. We claim that for any resource $e \in E$ we have

$$\begin{aligned} \left\{j \in N : e \in \phi_j(s_j)\right\} &= \{i\}, && \text{if } e = e(\mathbf{s}_{-i}) \text{ for some player } i, \\ \left|\left\{j \in N : e \in \phi_j(s_j)\right\}\right| &\geq 2, && \text{else.} \end{aligned}$$

First, note that $e(\mathbf{s}_{-i}) \in \phi_i(s_i)$. Moreover, for any player $j \neq i$, we have that

$$e(\mathbf{s}_{-i}) \notin \bigcup\nolimits_{\mathbf{t}_{-i} \in \mathbf{S}_{-i} : t_j \neq s_j} e(\mathbf{t}_{-i})$$

and, thus, also $e(\mathbf{s}_{-i}) \notin \phi_j(s_j)$. On the other hand, for any other resource $e \in E \setminus \left(\bigcup_{i \in N} e(\mathbf{s}_{-i})\right)$, we have that $e = e(\mathbf{t}_{-i})$ for some $j \in N$ and $\mathbf{t}_{-j} \in \mathbf{S}_{-j}$ and, thus, $e \in \phi_j(s_j)$. Using that $\mathbf{t}_{-j} \neq \mathbf{s}_{-j}$, we further derive the existence of a player $k \in N \setminus \{j\}$ with $t_k \neq s_k$ implying that $e(\mathbf{t}_{-j}) \in \phi_k(s_k)$.
For all $\mathbf{s} \in \mathbf{S}$, let $\mathbf{s}' = \times_{i \in N} \phi_i(s_i)$. We then obtain

$$\begin{aligned} \pi_i'(\mathbf{s}') &= \sum_{e \in \phi(s_i)} c_e\big(x_e(\mathbf{s}')\big) \\ &= \sum\nolimits_{\mathbf{t}_{-i} \in \mathbf{S}_{-i}} c_{e(\mathbf{t}_{-i})}\big(x_{e(\mathbf{t}_{-i})}(\mathbf{s}')\big) + \sum\nolimits_{j \in N \setminus \{i\}} \sum\nolimits_{\mathbf{t}_{-j} \in \mathbf{S}_{-j} : t_i \neq s_i} c_{e(\mathbf{t}_{-i})}\big(x_{e(\mathbf{t}_{-i})}(\mathbf{s}')\big) \\ &= \pi_i(s_i, \mathbf{s}_{-i}), \end{aligned}$$

	s_2	s_2'
s_1	$(0, 2)$	$(1, 2)$
s_1'	$(0, 3)$	$(1, 3)$

(a) Coordination game

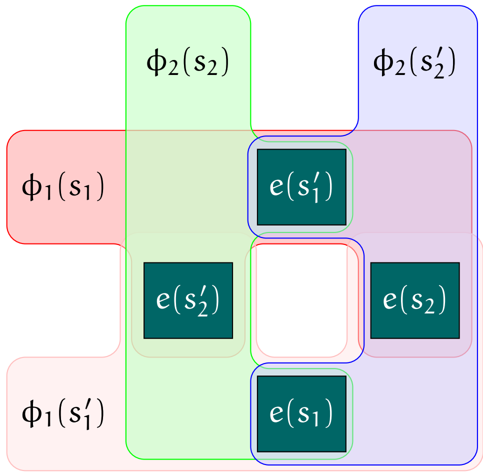

(b) Isomorphic congestion game

which concludes the proof. □

Example 4.14 (Congestion Game Isomorphic to a Dummy Game). Consider the dummy game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ in Figure 4.4a. For each $i \in N$ and each strategy partial strategy profile $\mathbf{s}_{-i} \in \mathbf{S}_{-i}$ we introduce a resource $e(\mathbf{s}_{-i})$, which leads to the strategies shown in Figure 4.4b.

We are now in position to state the main isomorphism result for congestion games.

Theorem 4.15. Every exact potential game is isomorphic to a unweighted congestion game.

Proof. By Theorem 4.9, we can write each exact potential game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ as the sum of a coordination game $G^c = (N, \mathbf{S}, \boldsymbol{\pi}^c)$ and a dummy game $G^d = (N, \mathbf{S}, \boldsymbol{\pi}^d)$. By Theorem 4.11 and Theorem 4.13, there are congestion games $G' = (N, \mathbf{S}', \boldsymbol{\pi}')$ and $G'' = (N, \mathbf{S}'', \boldsymbol{\pi}'')$ with resources E' and E'' that are isomorphic to G^c and G^d respectively, i.e., there are bijections $\phi_i' : S_i \to S_i'$ and $\phi_i'' : S_i \to S_i''$ for all players i such that

$$\pi_i^C(\times_{i \in N} s_i) = \pi_i'(\times_{i \in N} \phi_i'(s_i)) \qquad \text{and} \qquad \pi_i^D(\times_{i \in N} s_i) = \pi_i''(\times_{i \in N} \phi_i''(s_i)).$$

By renaming resources, we may assume that E and E' are disjoint. Consider the congestion game $\tilde{G} = (N, \tilde{\mathbf{S}}, \tilde{\boldsymbol{\pi}})$ with resources $E = E' \cup E''$ and cost functions as in G' and G'' and consider the bijection $\phi_i : S_i \to 2^E$ defined as $\phi_i(s_i) =$

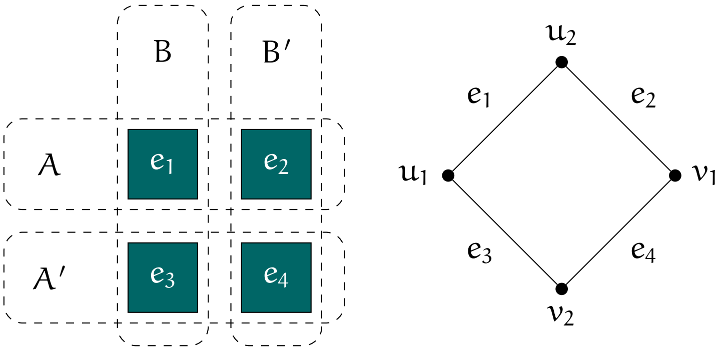

	B	B′
A	(62, 162)	(66, 160)
A′	(66, 160)	(62, 162)

(a) Resource representation **(b)** Graph representation **(c)** Strategic representation

Figure 4.4: Weighted congestion Game

$\phi_i'(s_i) \cup \phi_i''(s_i)$. We obtain

$$\begin{aligned}\tilde{\pi}_i\big(\times_{i\in N}\phi_i(s_i)\big) &= \pi_i(\times_{i\in N}\phi'(s_i)) + \pi_i(\times_{i\in N}\phi''(s_i))\\ &= \pi_i^C(\times_{i\in N}s_i) + \pi_i^D(\times_{i\in N}s_i)\\ &= \pi_i(\mathbf{s}),\end{aligned}$$

where the first equation uses that E' and E'' are disjoint. We conclude that $\tilde{G}$ is isomorphic to G. □

4.3 Weighted Congestion Games

Weighted congestion games are a straightforward generalization of unweighted congestion in which the player contribute differently to the congestion on the resources.

Definition 4.16 (Weighted Congestion Game). For a finite set of resources E with cost functions $(c_e)_{e\in E}$, and allocation vector $\mathcal{A} = (\mathcal{A}_i)_{i\in N}$ and a demand vector $\mathbf{d} = (d_i)_{i\in N} > 0$, the corresponding unweighted congestion game is the strategic game $G(\mathcal{A}) = (N, \mathbf{S}, \boldsymbol{\pi})$, where

- N is a finite and non-empty set of n players,
- $S_i = \mathcal{A}_i$ for all players i,
- $\pi_i(\mathbf{s}) = \sum_{e\in s_i} d_i c_e\big(x_e(\mathbf{s})\big)$ for all players i, where $x_e(\mathbf{s}) = \sum_{j\in N: e\in s_j} d_j$.

Example 4.17. Consider a weighted variant of the congestion game from Example 4.2 the game in Figure 4.4. Cost functions are defined as $c_{e_1}(x) = c_{e_4}(x) = 2x^3$ and $c_{e_2}(x) = c_{e_3}(x) = (x+1)^3$ for all $x \geq 0$. The players' demands are $d_1 = 1$ and $d_2 = 2$. This game does not have a pure Nash equilibrium.

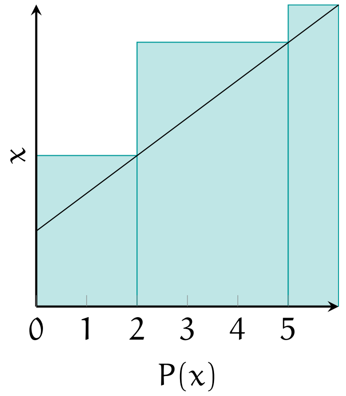

Figure 4.5: For an affine function the potential function is independent of the ordering of the players.

> **Theorem 4.18.** Weighted congestion games with affine cost functions have a pure Nash equilibrium.

Proof. Consider the function $P : \mathbf{S} \to \mathbb{R}$ defined as

$$P(\mathbf{s}) = \sum_{e \in E} \sum_{i \in N : e \in s_i} d_i c_e\Big(\sum\nolimits_{j \in \{1,\dots,i\} : e \in s_j} d_j\Big).$$

For players of unit weight this function is equal to the potential function given for unweighted congestion games. Again, the contribution of a resource to the potential function corresponds to a discrete approximation of the integral. However, this approximation is not uniform and, thus, depends on the indices of the players (respectively, their demands), in general. The crucial observation is that for affine function, the potential is independent of the ordering of the players, see Figure 4.5.
In fact, we obtain for cost functions of the form $c_e(x) = a_e x + b_e$ with $a_e, b_e \in \mathbb{R}$ that

$$\begin{aligned}
P(\mathbf{s}) &= \sum_{e \in E} \sum_{i \in N : e \in s_i} d_i c_e\Big(\sum\nolimits_{j \in \{1,\dots,i\} : e \in s_j} d_j\Big) \\
&= \sum_{e \in E} \left(b_e x_e(\mathbf{s}) + \sum\nolimits_{i,j \in N : e \in s_i \cap s_j, i \leq j} a_e d_i d_j \right) \\
&= \sum_{e \in E} \left(b_e x_e(\mathbf{s}) + \frac{1}{2} \sum\nolimits_{i,j \in N : e \in s_i \cap s_j} a_e d_i d_j + \frac{1}{2} \sum_{i \in N : e \in s_i} a_e d_i^2 \right) \\
&= \frac{1}{2} \sum_{e \in E} \Big(c_e(x_e(\mathbf{s})) + c_e(0) \Big) x_e(\mathbf{s}) + \frac{1}{2} \sum_{i \in N} \sum_{e \in s_i} \Big(c_e(d_i) - c_e(0) \Big) d_i.
\end{aligned}$$

We claim that P is an exact potential, i.e., $P(t_i, \mathbf{s}_{-i}) - P(\mathbf{s}) = \pi_i(t_i, \mathbf{s}_{-i}) - \pi_i(\mathbf{s})$ for all $i \in N$, $\mathbf{s} \in \mathbf{S}$ and $t_i \in S_i$. Since the potential function is independent of the ordering of the players, it is without loss of generality to assume that $i = n$. We then calculate

$$P(t_n, \mathbf{s}_{-n}) = P(\mathbf{s}) + d_n \sum_{e \in t_n \setminus s_n} c_e\big(x_e(t_n, \mathbf{s}_{-n})\big) - d_n \sum_{e \in s_n \setminus t_n} c_e\big(x_e(\mathbf{s})\big)$$
$$= P(\mathbf{s}) + \pi_i(t_n, \mathbf{s}_{-i}) - \pi_i(\mathbf{s}).$$

We conclude that P is a potential function and, thus, the game has a pure Nash equilibrium. □

For $\phi \in R$, let

$$\mathcal{C}_{\exp}(\phi) = \{f : \mathbb{R}_{\geq 0} \to \mathbb{R} \mid f(x) = a e^{\phi x} + b \text{ with } a, b \in \mathbb{R}\}.$$

Theorem 4.19. For every $\phi \in \mathbb{R}$, every weighted congestion game with cost functions in $\mathcal{C}_{\exp}(\phi)$ has a pure Nash equilibrium.

Proof. For $\phi = 0$, the theorem follows from Theorem 4.18 so we may assume that $\phi \neq 0$. Consider the function $P : \mathbf{S} \to \mathbb{R}$ defined as

$$P(\mathbf{s}) = \sum_{e \in E} \sum_{i \in N : e \in s_i} \operatorname{sgn}(\phi)(1 - e^{-\phi d_i}) c_e\Big(\sum\nolimits_{j \in \{1,\dots,i\} : e \in s_j} d_j\Big).$$

For cost functions of the form $c_e(x) = a_e e^{\phi x} + b_e$, this give

$$P(\mathbf{s}) = \sum_{e \in E} \sum_{i \in N : e \in s_i} \Big(\operatorname{sgn}(\phi)(1 - e^{-\phi d_i}) a_e \exp\Big(\phi \sum\nolimits_{j \in \{1,\dots,i\} : e \in s_j} d_j\Big) + b_e\Big)$$
$$= \sum_{e \in E} \sum_{i \in N : e \in s_i} \Big(\operatorname{sgn}(\phi) a_e e^{\phi \sum_{j \in \{1,\dots,i\} : e \in s_j} d_j} - \operatorname{sgn}(\phi) a_e e^{\phi \sum_{j \in \{1,\dots,i-1\} : e \in s_j} d_j}$$
$$+ b_e \operatorname{sgn}(\phi)(1 - e^{\phi d_i})\Big).$$

Collapsing the telescoping sum, we obtain

$$P(\mathbf{s}) = \operatorname{sgn}(\phi) \sum_{e \in E} \Big(a_e\big(e^{\phi x(\mathbf{s})} - 1\big) + b_e \sum\nolimits_{i \in N : e \in s_i} (1 - e^{-\phi d_i})\Big),$$

which is independent of the ordering of the players. For $\lambda_i = \operatorname{sgn}(\phi)\big(1 - e^{-\phi d_i}\big)$ we then obtain

$$P(t_n, \mathbf{s}_{-n}) = P(\mathbf{s}) + \lambda_n \sum_{e \in t_n \setminus s_n} c_e\big(x_e(t_n, \mathbf{s}_{-n})\big) - \lambda_n \sum_{e \in s_n \setminus t_n} c_e\big(x_e(\mathbf{s})\big)$$
$$= P(\mathbf{s}) + \frac{d_n}{\lambda_n}\big(\pi_i(t_n, \mathbf{s}_{-i}) - \pi_i(\mathbf{s})\big).$$

Using that $d_n/\lambda_n > 0$, we conclude that every sequence of profitable unilateral improvements is finite and, thus, the game has a pure Nash equilibrium. □

Definition 4.20 (Weighted Potential Game). A strategic game $G = (N, \mathbf{S}, \boldsymbol{\pi})$ is called weighted potential game if there is vector $(w_i)_{i\in N} > 0$ and a function $P : \mathbf{S} \to \mathbb{R}$ such that

$$P(t_i, \mathbf{s}_{-i}) - P(\mathbf{s}) = w_i\Big(\pi_i(t_i, \mathbf{s}_{-i}) - \pi_i(\mathbf{s})\Big)$$

for all $\mathbf{s} \in \mathbf{S}$, $i \in N$ and $t_i \in S_i$.

In the following, we show that the set of affine cost functions and the set of exponential cost functions are the only sets of cost functions for which general existence results as in Theorem 4.18 and Theorem 4.19 are possible. In order to make the statement precise, we introduce the notion of consistency of a set of cost functions.

Definition 4.21 (Consistent Set of Cost Functions). A set $\mathcal{C}$ of cost functions $c : \mathbb{R}_{\geq 0} \to \mathbb{R}$ is called consistent (for weighted congestion games) if every weighted congestion game with the property that $c_e \in \mathcal{C}$ for all $e \in E$ has a pure Nash equilibrium.

The following lemma states that when characterizing consistent sets of cost functions, it is without loss of generality to assume that the set is closed under positive integer scalar multiplication.

Lemma 4.22. A set $\mathcal{C}$ of cost functions is consistent if and only if the set

$$\bar{\mathcal{C}} := \{g : \mathbb{R}_{\geq 0} \to \mathbb{R} \mid f(x) = \lambda f(x) \text{ for some } \lambda \in \mathbb{N}, f \in \mathcal{C}\}$$

is consistent.

Proof Sketch. Every resource e with cost function λf where $\lambda \in \mathbb{N}$ and $f \in \mathcal{C}$ can be replaced by λ resources with cost function f. □

In order to characterize the sets of consistent cost functions, the following lemma will be useful.

Lemma 4.23. Let $\mathcal{C}$ be a consistent set of strictly increasing functions. Then,

$$\frac{f(x+y) - f(x)}{f(x+y) - f(y)} = \frac{g(x+y+t) - g(x+t)}{g(x+y+t) - g(y+t)} \tag{4.2}$$

for all $f, g \in \mathcal{C}$ and $x, y, t \in \mathbb{R}_{\geq 0}$.

Proof. By Lemma 4.22, it is without loss of generality to assume that $\mathcal{C}$ is closed under scalar multiplication. Let $\kappa, \lambda \in \mathbb{N}$ be arbitrary and consider a weighted

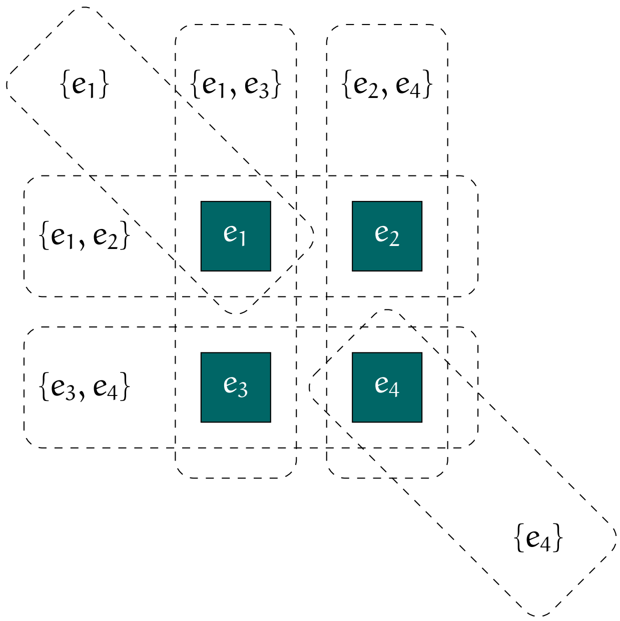

Figure 4.6: Congestion game constructed in the proof of Lemma 4.23.

congestion game with four resources $E = \{e_1, e_2, e_3, e_4\}$ with cost functions $c_{e_1} = c_{e_4} = \kappa g$, $c_{e_2} = c_{e_3} = \lambda f$ and four players i, j, k, k' with demands $d_i = x$, $d_j = y$, $d_k = d_{k'} = t$ and strategy sets $S_i = \big\{\{e_1, e_2\}, \{e_3, e_4\}\big\}$, $S_j = \big\{\{e_1, e_3\}, \{e_2, e_4\}\big\}$, $S_k = \big\{\{e_1\}\big\}$, $S_{k'} = \big\{\{e_4\}\big\}$.
Players i and j have two strategies each and players k and k' have one strategy each resulting in four distinct strategy profiles. There are only two types of strategy profiles: in the first type players i and j are together on a resource with cost function κg, but are not together on a resource with cost function λf. For the second type, it is the other way around. In a strategy profile $\mathbf{s}^1$ of the first type we have

$$\pi_i(\mathbf{s}^1) = \kappa g(x + y + t) + \lambda f(x) \qquad \pi_j(\mathbf{s}^1) = \kappa g(x + y + t) + \lambda f(y) \tag{4.3a}$$

while in a strategy profile $\mathbf{s}^2$ of the second type we have

$$\pi_i(\mathbf{s}^2) = \kappa g(x + t) + \lambda f(x + y) \qquad \pi_j(\mathbf{s}^2) = \kappa g(y + t) + \lambda f(x + y). \tag{4.3b}$$

In order to have a pure Nash equilibrium, one of these two types must be beneficial to both players, i.e., at least one of the following two cases holds:

$$\pi_\ell(\mathbf{s}^1) \le \pi_\ell(\mathbf{s}^2) \text{ for } \ell \in \{i, j\} \qquad \text{or} \qquad \pi_\ell(\mathbf{s}^1) \ge \pi_\ell(\mathbf{s}^2) \text{ for } \ell \in \{i, j\}. \tag{4.4}$$

Combining (4.3) and (4.4), we obtain that at least one of the following cases holds:

$$\frac{\kappa}{\lambda} \le \min\left\{\frac{f(x+y) - f(x)}{g(x+y+t) - g(x+t)}, \frac{f(x+y) - f(y)}{g(x+y+t) - g(y+t)}\right\} \qquad \text{or}$$

$$\frac{\kappa}{\lambda} \ge \max\left\{\frac{f(x+y) - f(x)}{g(x+y+t) - g(x+t)}, \frac{f(x+y) - f(y)}{g(x+y+t) - g(y+t)}\right\}.$$

This yields

$$\frac{f(x+y)-f(x)}{g(x+y+t)-g(x+t)} = \frac{f(x+y)-f(y)}{g(x+y+t)-g(y+t)}.$$

Multiplying this equation with $\frac{g(x+y+t)-g(x+t)}{f(x+y)-f(y)}$ gives the claimed result. □

Lemma 4.24. Let f be a continuous and strictly increasing function such that

$$\frac{f(x+y)-f(x)}{f(x+y)-f(y)} = \frac{f(x+y+t)-f(x+t)}{f(x+y+t)-f(y+t)} \tag{4.5}$$

for all $x, y, t \in \mathbb{R}_{\geq 0}$. Then, f is one of the two functional forms:

$$f(x) = ax + b, \qquad \text{or} \qquad f(x) = ae^{\phi x} + b.$$

Proof. Let $\epsilon > 0$ be arbitrary and let $x = \epsilon$, $y = 2\epsilon$, and $t = m\epsilon$ for some $m \in \mathbb{N}$. Using (4.5), we obtain

$$\gamma := \frac{f(3\epsilon)-f(1\epsilon)}{f(3\epsilon)-f(2\epsilon)} = \frac{f((m+3)\epsilon)-f((m+1)\epsilon)}{f((m+3)\epsilon)-f((m+2)\epsilon)}$$

for all $m \in \mathbb{N}$. We note that $\gamma \neq 1$ since $f(1\epsilon) \neq f(2\epsilon)$. Setting $a_m := f(m\epsilon)$ for all $m \in \mathbb{N}$ and rearranging terms, we derive that the sequence $(a_m)_{m\in\mathbb{N}}$ obeys the recursive formula

$$0 = a_{m+3} - \frac{\gamma}{\gamma-1}a_{m+2} + \frac{1}{\gamma-1}a_{m+1}$$

for all $m \in \mathbb{N}$. The characteristic equation of this recurrence relation equals

$$x^2 - \frac{\gamma}{\gamma-1}x + \frac{1}{\gamma-1} = (x-1)\Big(x - \frac{1}{\gamma-1}\Big).$$

If $\gamma \neq 2$, the characteristic equation has two distinct roots and a_m can be calculated explicitly and uniquely as

$$a_m = a\Big(\frac{1}{\gamma-1}\Big)^m + b \tag{4.6}$$

for some constants $a, b \in \mathbb{R}$. If, on the other hand, $\gamma = 2$, we can calculate a_m as

$$a_m = am + b \tag{4.7}$$

for some constants $a, b \in \mathbb{R}$. Thus, f is affine or exponential for every integer multiple of ϵ. As ϵ was arbitrary we may conclude that f is affine or exponential on a dense subset of $\mathbb{R}_{\geq 0}$. Using that f is continuous, the claimed result follows. □

Theorem 4.25. A set $\mathcal{C}$ of strictly increasing and continuous functions is consistent for weighted congestion games if and only if one of the following two cases holds:

1. $\mathcal{C}$ contains only affine functions of the form $c(x) = ax + b$.
2. $\mathcal{C}$ contains only exponential functions of the form $c(x) = ae^{\phi x} + b$, where ϕ is equal for all functions in $\mathcal{C}$.

Proof. The "if"-part of the statement follows from Theorem 4.18 and Theorem 4.19, so it suffices to show the "only if"-part.
By Lemma 4.24, all functions in $\mathcal{C}$ are either linear or exponential. It is left to show that $\mathcal{C}$ neither contains both affine and exponential functions nor $\mathcal{C}$ contains exponential functions with different ϕ in the exponent. To this end, recall that, by Lemma 4.23 it is necessary that

$$\frac{f(x+y)-f(x)}{f(x+y)-f(y)} = \frac{g(x+y+t)-g(x+t)}{g(x+y+t)-g(y+t)}$$

for all $f, g \in \mathcal{C}$ and $x, y, t \in \mathbb{R}_{\geq 0}$. For all $x, y, t \in \mathbb{R}_{\geq 0}$, the ratio $\frac{f(x+y+t)-f(x)}{f(x+y+t)-f(y+t)}$ is equal to x/y if f is linear, and is equal to $e^{\phi(x-y)}$ if f is of the form $ae^{\phi x} + b$. Thus, (4.5) is satisfied for all x, y only if $\mathcal{C}$ is as claimed. □

4.4 Computation of Equilibria

As congestion games are potential games, every sequence of unilateral improvements is finite. That is, in principle a pure Nash equilibrium can be computed by following such a sequence of better replies, see Algorithm 4.1.

```
Input: Congestion game G
Output: Nash equilibrium s of G
Choose s^0 ∈ S arbitrarily
k = 0
while s^k is not a Nash equilibrium do
    find player i and t_i ∈ S_i with π_i(t_i, s_{-i}) < π_i(s)
    s^{k+1} = (t_i, s_{-i})
    k = k + 1
end
return s^k
```

Algorithm 4.1: Better-reply algorithm

Usually, we assume that we can check in polynomial time for a player i whether there is a strategy $t_i \in S_i$ with $\pi_i(t_i, \mathbf{s}_{-i}) < \pi_i(\mathbf{s})$. This is trivial, if the set $S_i \subseteq 2^E$ is given explicitly. If S_i is given implicitly, e.g., the set S_i corresponds to the set of s_i, t_i-paths in a given graph, we can still solve the problem via a shortest-path-computation.

However, it is not possible, to give a polynomial bound on the number of iterations of said algorithm, even if all cost functions are affine-linear with integer coefficients, i.e., $c(x) = ax + b$ with $a, b \in \mathbb{N}$. To see this, note that $P(\mathbf{s}^0) \leq \sum_{e\in E} \sum_{k=1}^{n} c_e(k) = \sum_{e\in E} a_e \frac{n(n+1)}{2} + \sum_{e\in E} nb_e$. Since a_e and b_e are encoded in binary, this is not polynomially bounded in the input.

4.4.1 The Class PLS

In fact, one can show that to compute the Nash equilibrium of a congestion game is as hard as any other local search problem. The fact that a local search problem has a solution can be established via the following simple graph theoretic statement.

> Every finite acyclic digraph has a sink.

In a local search problem, the acyclicity of the graph is established via a potential function.

Let us denote the subclass of problems in TFNP whose solutions are the locally optimal solutions of an objective function by PLS. The acronym PLS stands for polynomial local search.

Definition 4.26. A problem Π_R is in PLS if, for all instances I and all solutions $x \in S_I$, $S_i = \Sigma^{p(|I|)}$, there is a neighborhood $\mathcal{N}_I(x) \subseteq S_I$ and polynomial algorithms $C_I : S_I \to \mathbb{N}$, $N_I : S_I \to \mathcal{N}_I \cup \{\emptyset\}$, $S_i = \Sigma^{p(|I|)}$ such that

- $C_I : S_I \to \mathbb{N}$ computes the cost of solution s for instance I
- $N_I : S_I \to \mathcal{N}_I(x) \cup \{\emptyset\}$ computes s' with $C(s') < C(s)$, or $\emptyset$, when s is locally optimal.
- $(I, x) \in R \Leftrightarrow N_I(x) = \emptyset$

CircuitFlip
Input: Boolean circuit with n input bits and m output bits.
Output: locally optimal input $s \in \{0, 1\}$.

Local Max Cut
Input: Graph $G = (V, A)$ with integer weights w_a on each arc a.
Output: partition $V_1, V_2 \in 2^V$ such that $V_1 \cap V_2 = \emptyset$, $V_1 \cup V_2 = V$ such that $\delta(V_1) = \sum_{a=(v,w)\in A: v\in V_1, w\notin V_1} w_a \geq \delta(V_1')$ for all V_1' with $|V_1 \Delta V_1'| = 1$.

Between problems in PLS there is the usual notion of reducability in TFNP, i.e., Π_R reduces to problem Π_S if there are polynomial time computable functions f and g such that $(I, g(x)) \in R \Leftrightarrow (f(I), x) \in S$.

Theorem 4.27. Local Max Cut is PLS-complete.

Theorem 4.28. Calculating a Nash equilbrium for a congestion game (where the $S_i \subseteq 2^E$ are given explicilty) is PLS-complete.

Proof. Calculating a Nash equilibrium is clearly in PLS so we only have to argue bout completeness. For the PLS-hardness, we reduce from the PLS-complete problem local max cut. Let $G = (V, A)$ be an instance of Local Max Cut. Let W be the sum of of the weights of all arcs in A, i.e., $W = \sum_{e \in E} w(e)$. Observe, that for every cut (V_1, V_2) we have $w(\delta(V_1)) = W - w(E(V_1)) - w(E(V_2))$, where $E(U) \subseteq E$ is the set of all arcs having both endpoints in $U \subseteq V$. Instead of maximizing the weight of arcs that traverse the cut, we can also minimize the weight of the arcs the remain inside the two sets of the partition.
The reduction is now as follows: For every arc $a \in A$ we add two resources $e_{a,1}$ and $e_{a,2}$. For $j \in \{1, 2\}$, we define

$$c_{e_{a,j}}(x) = \begin{cases} 0, & \text{if } x \leq 1, \\ w_a/2, & \text{else.} \end{cases}$$

Furthermore, for every node $v \in V$ we have a player $v \in N$. Player v has two strategies, $A_v := \{\mathcal{A}_{v,1}, \mathcal{A}_{v,2}\}$, where $\mathcal{A}_{v,j} := \{e_{a,1} : e \in \delta(v)\}$ for $j \in \{1, 2\}$. The strategies can be interpreted as follows. If player v chooses aloocation $\mathcal{A}_{v,1}$, then the node $v \in V$ belongs to the set V_1 of the partition (V_1, V_2); if the player chooses $\mathcal{A}_{v,2}$, then the node v belongs to the set V_2. An arc $a = \{u, v\}$ contributes the cost $w(e)$ to the cut if and only if player u and v both belong to the same part of the partition. This transformation can be done in polynomial time.
Using the above construction, we can translate a strategy profile of the congestion game into a feasible cut of Local Max Cut and vice versa. Let $\mathbf{s}$ be a Nash equilibrium of the congestion games and let (V_1, V_2) be the induced cut in G.
As $\mathbf{s}$ is a Nash equilibrium, we have for any v with $s_v = \mathcal{A}_{v,1}$ that

$$\sum_{a=(v,w), w \in V_2} \frac{w_a}{2} \geq \sum_{a=(v,w), w \in V_1} \frac{w_a}{2},$$

The same holds for a player v with $s_v = \mathcal{A}_{v,2}$. Thus, the cut is locally optimal. □

Theorem 4.28 shows that an efficient algorithm for calculating Nash equilibria in symmetric congestion games would imply the existence of efficient algorithms for many local search problems. This is very unlikely but not impossible. However, it is known that efficient algorithms cannot use the local search paradigm: There exist instances of Max Cut, such that every path from a node to a sink in the transition graph has exponential length.
One can also show that computing a Nash equilbrium is PLS-complete for network congestion games with affine linear costs, but the proof is more involved and beyond the scope of this lecture.

Theorem 4.29 (Ackermann et al., 2008). Calculating a Nash equilibrium for a network congestion games with affine-linear costs is PLS-complete.

4.4.2 Symmetric Network Congestion Games

Theorem 4.30. In symmetric network congestion games with non-decreasing costs, a Nash equilibrium can be computed in polynomial time.

Proof. Let G be a symmetric network congestion game with resource set E. We proceed to show that the minimum of Rosenthals potential function can be computed in polynomial time by a min-cost flow computation.
To this end, we replace each edge e in the graph by n parallel edges $e_1, \dots, e_n$ between the same nodes. Edge e_i is assigned cost $c_e(i)$, for $i \in \{1, \dots, n\}$. Since the edge costs functions are non-decreasing, we have $c_{e(1)} \leq c_{e(2)} \leq \cdots \leq c_{e(n)}$. All edges have capacity 1.
By standard minimum cost flow algorithms, we can compute an integral minimum cost flow with flow value n in polynomial time. By the monotonicity of the cost functions, it is without loss of generality to assume that the flow sends the flow only along the cheapest k copies of each edge. Thus, the cost for sending the flow along these edges is $c_{e(1)} + \cdots + c_{e(k)}$, which corresponds to the potential that Rosenthal's potential function assigns to edge e if k players use this edge. Consequently, we can translate the optimal solution of the min-cost flow problem into a state of the congestion game whose potential corresponds to the cost of the flow. Hence, the min-cost flow solution corresponds to a Nash equilibrium that globally minimizes Rosenthal's potential function. □

4.4.3 Matroid Congestion Games

For a ground set E, let $\mathcal{I} \subseteq 2^E$ be a set system over E.

Definition 4.31 (Matroid). A pair $(E, \mathcal{I})$ with $\mathcal{I} \subseteq 2^E$ is a matroid if
1. $\emptyset \in \mathcal{I}$.
2. $J \in \mathcal{I}$ and $I \subset J \Rightarrow I \in \mathcal{I}$ for all $I, J \in 2^E$.
3. $I, J \in \mathcal{I}$ and $|I| < |J| \Rightarrow$ there is $x \in J \setminus I$ with $I + x \in \mathcal{I}$.

Systems that only satisfy (1) and (2) are called independence system. Every set $I \in \mathcal{I}$ is called independent and sets $I \in 2^E \setminus \mathcal{I}$ are called dependent. Set-wise minimal dependent sets are called circuits. Set-wise maximal independent sets are called basis. It follows from (3), that all bases of a matroid have the same cardinality. The cardinality of a basis of a matroid is called the rank of a matroid.

Example 4.32 (Graphical matroid). Let G be a graph. Then the subset $\mathcal{I} \subset 2^E$ of forests of G is a matroid. If G is connected, the bases are the maximal cycle-free graphs, i.e., the spanning trees of the graph.

Example 4.33 (Uniform matroid). Let E be a set of n elements. For an integer k, the set $\{I \subseteq E : |I| \leq k\}$ is called a uniform matroid.

Example 4.34 (Partition matroid). Let E be a set of n elements and let $E_1, \ldots, E_k$ be a partition of E, i.e., $\bigcup_{i=1}^{k} E_i = E$ and $E_i \cap E_j = \emptyset$ for all $i \neq j$. Further, let $b_1, \ldots, b_k$ be non-negative integers. Then

$$\{I \subseteq E : |I \cap E_i| \leq b_i \text{ for all } i = 1, \ldots, k\}$$

is called a partition matroid.

Lemma 4.35. Let E be a finite set, $\mathcal{B}$ be the set of bases of a matroid $(E, \mathcal{I})$ and $(c_e)_{e \in E}$ a vector of real-values weights. Then, the minimization problem

$$\min\Big\{\sum_{e \in B} c_e : B \in \mathcal{B}\Big\}$$

can be solved in polynomial time by the standard greedy algorithm, see Algorithm 4.2.
Moreover, every optimal solution of this minimization problem can be obtained by the greedy algorithm for an appropriate tie-breaking rule for edges with the same weight.

Input: ground set E with weights $(c_e)_{e \in E}$, matroid $(E, \mathcal{I})$, given by an independence oracle
Output: basis $B \in \mathcal{I}$ minimizing $\sum_{e \in B} c_e$
sort elements such that $c_1 \leq c_2 \leq \cdots \leq c_n$
$I := \emptyset$
for $e = 1, \ldots, n$ **do**
  if $I \cup \{e\} \in \mathcal{I}$ **then** $I := I \cup \{e\}$
end
return I

Algorithm 4.2: Greedy algorithm for matroids

A congestion game is a matroid congestion game, if, for every player i, there is a matroid $(E_i, \mathcal{I}_i)$ such that $\mathcal{A}_i$ is equal to the set of bases of $\mathcal{I}_i$.

Theorem 4.36. Let G be a matroid congestion game. Then, every best-reply sequence converges in polynomial time to a Nash equilibrium.

Proof. The game has m resources with cost functions c_e, each of which can be used by x players with $x \in \{1, \dots, n\}$. Thus, there are at most $n \cdot m$ different cost values that can occur on the resources. Let L be a list of these cost values sorted non-decreasingly such that each value appears only once. For each resource E, consider the alternative cost function

$$c'_e : \mathbb{N} \to \{1, \dots, n \cdot m\},$$

where $c'_e(x)$ is equal to the index of the entry in the list L. Moreover, let π' denote the private cost function based on the alternative cost functions c'_e.
Fix an arbitrary strategy profile $\mathbf{s}$. Let i be a player such that there is a $s'_i \in S_i$ with $\pi_i(s'_i, \mathbf{s}_{-i}) < \pi_i(\mathbf{s})$ and let

$$t_i \in \arg\min_{s'_i \in S_i} \pi_i(s'_i, \mathbf{s}_{-i}). \tag{4.8}$$

be computed by the greedy algorithm. Then,

$$\pi'_i(t_i, \mathbf{s}_{-i}) < \pi'_i(\mathbf{s}),$$

i.e., the private cost decreases also with respect to the alternative costs after a best reply. By Lemma 4.35, it is without loss of generality to assume that the minimum in (4.8) is computed by the greedy algorithm. It is easy to verify that this greedy algorithm computes the same output for the original costs and the alternative costs (assuming the same tie-breaking rule). This implies $t_i \in \arg\min_{s'_i \in S_i} \pi'_i(s'_i, \mathbf{s}_{-i})$, and, thus, $\pi'_i(t_i, \mathbf{s}_{-i}) \leq \pi'_i(\mathbf{s})$. If $\pi'_i(t_i, \mathbf{s}_{-i}) = \pi'_i(\mathbf{s})$, then there is a tie-breaking rule such that s_i is the output of the greedy algorithm for the alternative costs, which implies $s_i \in \arg\min_{s'_i \in S_i} \pi_i(s'_i, \mathbf{s}_{-i})$, a contradiction.
Next, consider Rosenthal's potential function P' for the alternative costs. Using that $c'_e(x) \leq nm$ for all $e \in E$ and $x \in \{1, \dots, n\}$, we obtain

$$P'(\mathbf{s}) = \sum_{e \in E} \sum_{k=1}^{x_e(\mathbf{s})} \leq n^2 m^2.$$

Using that every best-response reduces the potential by at least 1, the claimed result follows. □

4.5 Non-Atomic Players

In this section, we consider a slightly different model with an infinite number of infinitesimally small players. This model is conveniently represented as a flow.

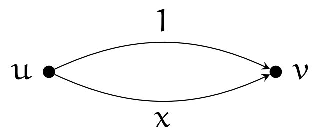

Figure 4.7: Pigou's example

Definition 4.37 (Flow). Let $G = (V, E)$ be a graph and let N be a set of commodities $(u_i, v_i, d_i) \in V \times V \times \mathbb{R}_{>0}$. For each commodity i, a flow is non-negative a vector $f^i = (f^i_P)_{P \in \mathcal{P}_i}$, where $\mathcal{P}_i$ is the set of all paths from u_i to v_i. The flow is feasible if $\sum_{P \in \mathcal{P}_i} f_P = d_i$. A multicommodity flow is a vector of flows $(f^i)_{i \in N}$, one for each commodity.

An alternative formulation of flows is the edge representation. An edge flow for commodity i is a non-negative vector $f^i = (f^i_e)_{e \in E}$. An edge flow is feasible if

$$\sum_{e \in \delta+(v)} f^i_e - \sum_{e \in \delta^-(v)} f^i_e = \begin{cases} d_i, & \text{if } v = u_i \\ -d_i, & \text{if } v = v_i \\ 0, & \text{else.} \end{cases}$$

It is a basic fact of the theory of flows that for each feasible edge flow, there is a corresponding feasible path flow and vice versa. In the following, we use path flows and edge flows interchangeably and we write $f_e = \sum_{i \in N} f^i_e$.
The equilibrium concept used in this section is that of a Wardrop flow.

Definition 4.38 (Wardrop's First Principle). Let G be a graph with flow-dependent edge costs $c_e : \mathbb{R}_{\geq 0} \to \mathbb{R}$ and commodities (u_i, v_i, d_i). A multi-commodity flow f satisfies Wardrop's first principle if for every commodity i

$$\sum_{e \in P} c_e(f_e) \leq \sum_{e \in Q} c_e(f_e) \text{ for all } P, Q \in \mathcal{P}_i \text{ with } f_P > 0.$$

A flow satisfying Wardrop's first principle is called Wardrop equilibrium.

Example 4.39 (Pigou's example). Consider the parallel link network in Fig. 4.7. The upper edge has a constant cost of 1; the lower link has a cost equal to the flow on that edge. One unit wants to travel from u to v.
In the unique Wardrop equilibrium, all flow is routed along the bottom edge.
To compute the System optimum, assume that we send $(1 - p)$ units along the top edge and p units along the bottom edge. The minimum is attained for

$$\min_{p \in [0,1]} (1 - p) + p^2,$$

which is minimized for $p = 1/2$ and gives a total cost of $3/4$.

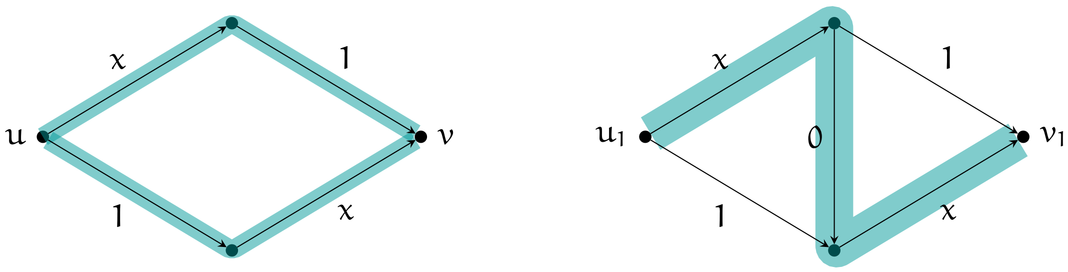

(a) Original graph

(b) Augmented graph

Figure 4.8: Breass' paradox

Example 4.40 (Braess' Paradox reconsidered). Reconsider Braess' paradox in a non-atomic context. There is 1 unit of flow travelling from u to v in the network in Figure 4.8a. In the unique equilibrium, there is $1/2$ units of flow on each path resulting in a total cost of $3/2$. After adding an additional edge with travel time 0, in the unique equilibrium the cost of all players is 2, see Figure 4.8b.

Theorem 4.41. In every network with continuous edge-cost functions c_e, a Wardrop flow exists. Moreover, this flow appears as the limit of a sequence of pure Nash equilibria of atomic congestion games played on the same graph.

Proof. We show the result only for rational demands d_i. By scaling demands and cost functions, we may assume that the demand d_i of each commodity (u_i, v_i, d_i) is integral. For all $n \in N$, let G^n the atomic congestion game, where, for each commodity i, there are $d_i \cdot n$ players with demand $1/n$, whose set of strategies is equal to all simple (u_i, v_i)-paths. By Rosenthal's Theorem, each of these games has a pure Nash equilibrium. For each game G^n, choose a Nash equilibrium $\mathbf{s}^n$ arbitrarily.
For each such game G^n, consider the path flow f^n defined as

$$f^n_{P_i} = \frac{|\{\text{ players of commodity } i \text{ that use } P_i \text{ in } \mathbf{s}^n\}|}{n}$$

for all commodities i and (u_i, v_i)-paths P_i.
The sequence $(f^n)_{n\in\mathbb{N}}$, viewed as path flows for each commodity, is clearly bounded. Thus, by the Theorem of Bolzano and Weierstrass, there is a convergent subsequence $(f^{n(\ell)})_{\ell\in\mathbb{N}}$. Let f^* be the limit of the convergent subsequence. We will argue that this path flow f^* satisfies Wardrop's first principle.
First note that since for each flow f^n, we have for each commodity i that the sum of the flow values of the (u_i, v_i)-paths is d_i, and the set of such flows is bounded, also f^* is a valid multi-commodity path flow that satisfies all commodities' demands.

In remains to show the defining inequality of Wardrop's first principle.
Suppose the inequality is not true, i.e., there is a commodity i and (u_i, v_i)-paths P_i, P_i' with $f_{P_i} > 0$, but there is $\epsilon > 0$ s.t.

$$\sum_{e \in P_i} c_e(f_e^*) > \epsilon + \sum_{e \in P_i'} c_e(f_e^*). \tag{4.9}$$

Let m denote the number of edges in G. First, note that since the cost functions c_e are continuous, they are uniformly continuous on $[0, d]$, where $d = \sum_i d_i$. Thus, there is $\delta > 0$ such that $|c_e(x) - c_e(x - \delta)| < \frac{\epsilon}{2m}$ for all $x \in [0, d]$ and $e \in E$.
Let k be the number of commodities and let $\tilde{l} \in \mathbb{N}$ be chosen such that the following hold:

1. $\left|f_{P_i}^* - f_{P_i}^{n(\ell)}\right| \leq \frac{\delta}{2k}$ for all commodities i, (u_i, v_i)-paths P_i and $\ell \geq \tilde{\ell}$,
2. $\frac{1}{n(\tilde{\ell})} \leq \frac{\delta}{2k}$, and
3. $f_{P_i}^{n(\ell)} > 0$ for all $\ell \geq \tilde{\ell}$.

We obtain for $\ell \geq \tilde{\ell}$ that

$$\sum_{e \in P_i} c_e(f_e^*) \leq \sum_{e \in P_i} c_e(f_e^{n(\ell)}) + \frac{\epsilon}{2} \leq \sum_{e \in P_i'} c_e\left(f_e^{n(\ell)} + \frac{1}{n(\ell)}\right) + \frac{\epsilon}{2} \leq \sum_{e \in P_i'} c_e(f_e^*) + \epsilon.$$

The first and the last inequality are satisfied since the cost functions c_e are uniformly continuous and f_e^* and $f_e^{n(\ell)}$ differ by at most δ. The second inequality uses that the flow $f^{n(\ell)}$ is a Nash equilibrium, in the sense that no atomic user (with demand $1/n(\ell)$) may improve when switching to path P_i'. This contrdicts (4.9). □

Lemma 4.42. A feasible flow f is a Wardrop flow if and only if it satisfies the following variational inequality:

$$\sum_{e \in E} c_e(f_e)(f_e - g_e) \leq 0 \tag{4.10}$$

for all feasible multi-commodity flows g.

Proof. "$\Leftarrow$": Let f be a flow satisfying (4.10), $i \in N$ and $P, Q \in \mathcal{P}_i$ be two paths with $f_P^i > 0$. Let $\lambda = f_P^i$ and consider the flow g defined as $g^j = g^i$ for all $j \in N \setminus i$ and

$$g_e^i = \begin{cases} f_e^i - \lambda, & \text{if } e \in P \setminus Q \\ f_e^i + \lambda, & \text{if } e \in Q \setminus P \\ f_e^i, & \text{else.} \end{cases}$$

By construction, g is feasible and, thus, (4.10) implies

$$\sum_{e \in E} c_e(f_e)(f_e - g_e) = \sum_{e \in P} c_e(f_e)\lambda - \sum_{e \in Q} c_e(f_e)\lambda \leq 0.$$

This implies $\sum_{e\in P} c_e(f_e) \leq \sum_{e\in Q} c_e(f_e)$.
"$\Rightarrow$": Let f be a Wardrop flow. Then, there is a constant k_i for every commodity i such that $\sum_{e\in P} c_e(f_e) = k_i$ for all $P \in \mathcal{P}_i$ with $f_P > 0$. In addition, $\sum_{e\in Q} c_e(f_e) \geq k_i$ for all $Q \in \mathcal{P}_i$ with $f_Q = 0$. For an arbitrary feasible flow g, we obtain

$$\begin{aligned}\sum_{e\in E} c_e(f_e)f_e &= \sum_{i\in N}\sum_{P\in\mathcal{P}_i} k_i f_P^i = \sum_{i\in N} k_i \sum_{P\in\mathcal{P}_i} f_P^i \\ &= \sum_{i\in N} k_i d_i = \sum_{i\in N} k_i \sum_{P\in\mathcal{P}_i} g_P^i = \sum_{i\in N}\sum_{P\in\mathcal{P}_i} k_i g_P^i \\ &\leq \sum_{i\in N}\sum_{P\in\mathcal{P}_i} g_P^i \sum_{e\in P} c_e(f_e) = \sum_{e\in E} c_e(f_e) g_e,\end{aligned}$$

which concludes the proof. □

Using this variational inequality, it can be shown that a Wardrop flow f can be computed by solving a convex optimization problem.

Theorem 4.43. For continuous and non-decreasing cost functions, a feasible flow is a Wardrop flow if and only if solves the following convex optimization problem:

$$\text{minimize} \quad \sum_{e\in E}\int_0^{f_e} c_e(t)\,dt \tag{4.11}$$

$$\text{subject to} \quad \sum_{e\in\delta^+(v)} f_e^i - \sum_{e\in\delta^-(v)} f_e^i = \begin{cases} d_i, & \text{if } v = u_i \\ -d_i, & \text{if } v = v_i \\ 0, & \text{else.}\end{cases} \quad \text{for all } v\in V, i\in N$$

$$f_e^i \geq 0 \qquad \text{for all } e\in E, i\in N.$$

Proof. We only show that a Wardrop equilibrium f is an optimal solution to the above problem. Using that the cost functions are non-decreasing, the objective functions is clearly convex. Let $P(f) = \sum_{e\in E}\int_0^{f_e} c_e(t)dt$. Recall that the first order Taylor polynomial of a function $h : \mathbb{R}^n \to \mathbb{R}$ in $a \in \mathbb{R}^n$ has the form

$$T_h(x; a) = h(a) + (x - a)^T \nabla h(a).$$

Consider the linear approximation (i.e., the first two terms of the Taylor series) of P in f, i.e.,

$$T_P(x; f) = P(f) + \sum_{e\in E} c_e(f_e)(x_e - f_e).$$

By the convexity of $P(x)$, the tangent plane is below the function, i.e.,

$$T_P(x; f) = P(f) + \sum_{e\in E} c_e(f_e)(x_e - f_e) \leq P(x)$$

for all feasible flows x. This implies

$$\begin{aligned} P(f) &\le P(x) + \sum_{e\in E} c_e(f_e)(f_e - x_e). \\ &\le P(x) \end{aligned}$$

where we used variational inequality. We conclude that f minimizes the convex program as claimed. □

With a similar approach, also the system optimal flow can be computed efficiently, we here require that the cost functions are semi-convex, i.e., that the function $x \mapsto xc_e(x)$ is convex for all $e \in E$.

Theorem 4.44. For continuous and semi-convex cost functions, the system optimal flow can be computed solving the following convex problem:

$$\begin{aligned} &\text{minimize} \sum_{e\in E} c_e(f_e)f_e && (4.12) \\ &\text{subject to} \sum_{e\in\delta^+(v)} f_e^i - \sum_{e\in\delta^-(v)} f_e^i = \begin{cases} d_i, & \text{if } v = u_i \\ -d_i, & \text{if } v = v_i \\ 0, & \text{else.} \end{cases} && \text{for all } v \in V, i \in N \\ &f_e^i \ge 0 && \text{for all } e \in E, i \in N. \end{aligned}$$

Note that the system (4.12) is similar to that of (4.11) except for the fact that for the system optimum, we minimize $C(f) = \sum_e c_e(f_e)f_e$ while for the equilibrium we minimize $\sum_e \int_0^{f_e} c_e(t)dt$. Next observe that

$$c_e(f_e)f_e = \int_0^{f_e} \tilde{c}_e(t)dt$$

if

$$c_e'(f_e)f_e + f_e = \tilde{c}_e,$$

the system optimum is a Wardrop flow with respect to the modified cost functions $\tilde{c}_e$ defined as $\tilde{c}_e(x) = c_e'(x)x + c_e(x)$.
Thus, we obtain the following direct corollary from Lemma 4.42.

Corollary 4.45. A feasible flow f solves (4.12) if and only if

$$\sum_{e\in E} \Big(c_e(f_e) + c_e'(f_e)f_e\Big)(f_e - g_e) \le 0$$

for all feasible flows g.

4.6 Efficiency of Equilibria

In light of Pigou's and Braess' network in which a Wardrop flow has higher total cost than the system optimium, it is a natural question to ask how bad an equilibrium can be compared to an optimal flow.

Definition 4.46 (Price of anarchy). For a non-atomic routing game G, let f be a Wardrop equilibrium and let f^* be a system optimal flow. Then, the price of anarchy is defined as

$$\rho(G) = \frac{C(f)}{C(f^*)}.$$

Moreover, for a set $\mathcal{G}$ of instances, let $\rho(\mathcal{G}) = \sup_{G \in \mathcal{G}} \rho(G)$.

4.6.1 Nonatomic Games with Affine Cost Functions

Theorem 4.47. Let $\mathcal{G}$ be the class of non-atomic routing games with affine cost functions. Then, $\rho(\mathcal{G}) = 4/3$.

Proof. The inequality $\rho(\mathcal{G}) \geq 4/3$ follows from Pigou's example (Example 4.39. We proceed to show $\rho(\mathcal{G}) \leq 4/3$. Let f be a Wardrop equilibrium and let g be a feasible flow. We obtain

$$C(f) = \sum_{e \in E} c_e(f_e) f_e \leq \sum_{e \in E} c_e(f_e) g_e$$

by the variational inequality. Thus,

$$C(f) \leq \sum_{e \in E} \Big(c_e(g_e) g_e + \big(c_e(f_e) - c_e(g_e) \big) g_e \Big)$$

Let us set $W_e(f_e, g_e) = \big(c_e(f_e) - c_e(g_e)\big) g_e$. Then, we obtain

$$C(f) \leq \sum_{e \in E} c_e(g_e) g_e + \sum_{e \in E} W_e(f_e, g_e).$$

We proceed to bound $W_e(f_e, g_e)$ in terms of $c_e(f_e) f_e$. To this end let

$$\mu = \sup_{f_e, g_e \in \mathbb{R}_{\geq 0}} \frac{\big(c_e(f_e) - c_e(g_e)\big) g_e}{c_e(f_e) f_e}.$$

First, note that for $g_e > f_e$, we have $\mu \leq 0$ as c_e is non-decreasing, so the supremum is clearly attained for $g_e \leq f_e$. For affine functions, this ratio is upper bounded by

1/4, see Figure . For a more formal proof, note that for affine functions of type $c(x) = ax + b$ with $a, b \in \mathbb{R}_{\geq 0}$, we have

$$\begin{aligned}\mu &\leq \sup_{f,g,a,b\in\mathbb{R}_{\geq 0}} \frac{(af+b-ag-b)g}{(af+b)f}\\ &= \sup_{f,g,a,b\in\mathbb{R}_{\geq 0}} \frac{(af-ag)g}{(af+b)f}\\ &= \sup_{f,g,a\in\mathbb{R}_{\geq 0}} \frac{(af-ag)g}{af^2}\\ &= \sup_{f,g\in\mathbb{R}_{\geq 0}} \frac{(f-g)g}{f^2}\\ &= \sup_{f,g\in\mathbb{R}_{\geq 0}} \left(1-\frac{g}{f}\right)\frac{g}{f}\\ &= \sup_{x\in\mathbb{R}_{>0}} (1-x)x\\ &= \frac{1}{4}\end{aligned}$$

Finally, we obtain

$$C(f) \leq C(g) + \frac{1}{4}C(f),$$

which implies the result. □

4.6.2 Nonatomic Games with Arbitrary Cost Functions

For a cost function $c : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, let

$$\mu(c) = \sup_{f,g\in\mathbb{R}_{\geq 0}} \frac{\big(c(f)-c(g)\big)g}{c(f)f}.$$

For a class $\mathcal{C}$ of functions, let $\mu(\mathcal{C}) = \sup_{c\in\mathcal{C}} \mu(c)$.

Theorem 4.48. Let G be a non-atomic routing game with cost functions in $\mathcal{C}$ where $\mu(\mathcal{C}) < 1$, then $\rho(G) \leq \frac{1}{1-\mu(\mathcal{C})}$.

Proof. As in the proof of Theorem 4.47, we obtain

$$\begin{aligned}C(f) = \sum_{e\in E} c_e(f_e)f_e \leq \sum_{e\in E} c_e(f_e)g_e = \sum_{e\in E}\Big(c_e(f_e)g_e + c_e(g_e)g_e - c_e(g_e)g_e\Big)\\ \geq C(g) + \mu(\mathcal{C})C(f),\end{aligned}$$

which implies the result. □

Under a mild assumption on $\mathcal{C}$, we can also give a matching lower bound.

Theorem 4.49. Let $\mathcal{C}$ be a class of functions that include all constant functions and let $\mathcal{G}$ be the set of all non-atomic routing games with cost functions in $\mathcal{C}$. Then $\rho(\mathcal{C}) \geq \frac{1}{1-\mu(\mathcal{C})}$.

Proof. Consider a Pigou-style instance with two cost functions c_1 and c_2 where a total demand of d has to be routed from the origin to the destination. Let $c_1 = c$ where $c \in \mathcal{C}$ is arbitrary and let c_2 be the function that is constant to $c(d)$. Then, sending a flow of d on the first edge is a Wardrop equilibrium f with cost $C(f) = dc(d)$. The cost of the system optimum f^* equals

$$C(f^*) = \min_{x\in[0,d]}\{xc(x) + c(d)(d-x)\} = dc(d) - \max_{x\in[0,d]}\big\{\big(c(d)-c(x)\big)x\big\}.$$

Thus,

$$\frac{C(f)}{C(f^*)} = \left(1 - \frac{\max_{x\in[0,d]}\{(c(d)-c(x))x\}}{dc(d)}\right)^{-1}.$$

Taking the supremum over $d \in \mathbb{R}_{\geq 0}$ and $c \in \mathcal{C}$, we obtain the claimed result. □

For a parameter $m > 0$, let $\mathcal{C}_m$ be the class of continuous, non-decreasing and semi-convex functions such that $c(\lambda x) \geq \lambda^m c(x)$ for all $c \in \mathcal{C}$ and $x \geq 0$ and $\lambda \in [0,1]$. Note that for all $m \in \mathbb{N}$, the class $\mathcal{C}_m$ contains all polynomials with non-negative coefficients and maximal degree at most m since

$$\sum_{k=0}^{m} a_k(\lambda x)^k \geq \lambda^k \sum_{k=0}^{m} a_k x^k.$$

In addition, $\mathcal{C}_1$ contains all semi-convex and concave functions.

Lemma 4.50. $\mu(\mathcal{C}_m) \leq \frac{m}{(m+1)^{1+1/m}}$.

Proof. Recall that

$$\mu(\mathcal{C}_m) = \sup_{c\in\mathcal{C}}\ \sup_{f,g\in\mathbb{R}_{\geq 0}} \frac{\big(c(f)-c(g)\big)g}{c(f)f}.$$

The supremum is clearly attained for $g \leq f$. Substituting $x = g/f$, we obtain

$$\begin{aligned}
\mu(\mathcal{C}_m) &= \sup_{c\in\mathcal{C}}\ \sup_{f,x\in\mathbb{R}_{\geq 0}} \frac{\big(c(f)-c(xf)\big)xf}{c(f)f}\\
&\leq \sup_{c\in\mathcal{C}}\ \sup_{f,x\in\mathbb{R}_{\geq 0}} \frac{\big(c(f)-x^m c(f)\big)xf}{c(f)f}\\
&= \sup_{x\in\mathbb{R}_{\geq 0}} \big(1-x^m\big)x
\end{aligned}$$

The unique solution to this maximization problem is attained for $x = \sqrt[m]{\frac{1}{m+1}}$ which yields the claimed bound. □

Corollary 4.51. Let G be a non-atomic routing games with cost functions in $\mathcal{C}_m$. Then, $\rho(G) \leq \left(1 - \frac{m}{(m+1)^{1+1/m}}\right)^{-1}$.