

Prices of Anarchy and Stability

Algorithmic Game Theory

Winter 2022/23

Wardrop Games

Price of Anarchy

Smoothness

Price of Stability

Wardrop Traffic Model

A **Wardrop game** is given by

- ▶ Directed graph $G = (V, E)$
- ▶ k kinds of traffic (termed “commodities”) with source-sink pair (s_i, t_i) and rate $r_i > 0$, für $i \in [k] = \{1, \dots, k\}$
- ▶ Commodity i composed of infinitely many, infinitesimally small players with total mass r_i .
- ▶ Latency functions $d_e : [0, 1] \rightarrow \mathbb{R}$, for every $e \in E$.

Wardrop games are a model for traffic routing (or, more generally, resource allocation) scenarios with a large number of users, in which the influence of a single user is very small (e.g., in large traffic or computer networks).

Single-Commodity Games

Simplifying Assumptions:

- ▶ Latency functions are non-negative, non-decreasing, and convex.
- ▶ We consider **single-commodity** games, with source s and sink t .
- ▶ Rate normalized to $r_1 = 1$.

Flows and Latencies in Wardrop Games:

- ▶ \mathcal{P} is the set of all paths P from s to t .
- ▶ A **flow** gives a flow value $f_P \in [0, 1]$, for every $P \in \mathcal{P}$, with $\sum_{P \in \mathcal{P}} f_P = 1$.
- ▶ The **edge flow** on $e \in E$ is $f_e = \sum_{P \ni e} f_P$.
- ▶ The **edge latency** for flow f is $d_e(f_e)$, for every $e \in E$.
- ▶ The **path latency** for flow f is $d_P(f) = \sum_{e \in P} d_e(f_e)$, for every $P \in \mathcal{P}$.

Wardrop Equilibrium

Definition (Wardrop Equilibrium)

A flow f is a **Wardrop equilibrium** if for every pair of paths $P_1, P_2 \in \mathcal{P}$ with $f_{P_1} > 0$ we have $d_{P_1}(f) \leq d_{P_2}(f)$.

Observations:

- ▶ In a Wardrop equilibrium the flow values must satisfy a condition. We do not need to give explicit s - t -paths for every one of the infinitely many players.
- ▶ If we distribute players to paths according to flow values, then every player gets a path with a latency that is currently the best among all s - t -paths.
- ▶ However, in a Wardrop equilibrium there also could be **countably infinitely many players** that choose suboptimal paths. Since these players have no mass (Lebesgue measure 0), they do not harm the equilibrium flow condition.


Social Cost

We consider as social cost the average latency of the flow.

Definition (Social Cost)

The **social cost** of a flow f is the weighted average of all player costs/path latencies

$$C(f) = \sum_{P \in \mathcal{P}} d_P(f) \cdot f_P = \sum_{e \in E} d_e(f_e) f_e .$$




Social Cost

We consider as social cost the average latency of the flow.


Definition (Social Cost)

The **social cost** of a flow f is the weighted average of all player costs/path latencies


$$C(f) = \sum_{P \in \mathcal{P}} d_P(f) \cdot f_P = \sum_{e \in E} d_e(f_e) f_e .$$



Wardrop Equilibria – Examples




Wardrop Equilibria – Examples




Social Cost: 2

Wardrop Equilibria – Examples



Social Cost: 2

Wardrop Equilibria – Examples



Social Cost: 2

Social Cost: 1.5

Braess Paradox: Destroying a fast connection **improves** the cost in equilibrium **for every single player**.

The Price of Anarchy

How good or desirable is a Wardrop equilibrium in terms of social cost? How much worse is the social cost of an equilibrium than the social cost of an optimal flow?

Which parameters of the game influence the social cost at equilibrium?

$$\text{Price of Anarchy} \stackrel{\text{def}}{=} \frac{\text{Worst social cost of any equilibrium}}{\text{Optimal social cost}} .$$

$$\text{Price of Stability} \stackrel{\text{def}}{=} \frac{\text{Best social cost of any equilibrium}}{\text{Optimal social cost}} .$$

Theorem (Roughgarden, Tardos, 2002)

The price of anarchy in Wardrop games with affine latency functions is at most $4/3$.

Price of Anarchy – Proof

(Correa, Schulz, Stier-Moses, 2008)

Let f be an equilibrium flow. Let g be any arbitrary flow.

$$\begin{aligned} C(f) &= \sum_{P \in \mathcal{P}} f_P d_P(f) \\ &\leq \sum_{P \in \mathcal{P}} g_P d_P(f) \\ &= \sum_{e \in E} g_e d_e(f_e) \\ &= \sum_{e \in E} g_e (d_e(g_e) + d_e(f_e) - d_e(g_e)) \\ &= C(g) + \sum_{e \in E} g_e (d_e(f_e) - d_e(g_e)) . \end{aligned}$$

Price of Anarchy – Proof

(Correa, Schulz, Stier-Moses, 2008)

Lemma

For every edge $e \in E$,

$$g_e (d_e(f_e) - d_e(g_e)) \leq \frac{1}{4} f_e d_e(f_e) .$$


Proof:

The lemma obviously holds for $f_e < g_e$, since $d_e(f_e) \leq d_e(g_e)$. Then the left side of the inequality is at most 0.

Price of Anarchy – Proof

(Correa, Schulz, Stier-Moses, 2008)

Hence, suppose $f_e \geq g_e$ and consider the following picture of the function d_e .



Comparing the two areas in the picture we see that

$$g_e (d_e(f_e) - d_e(g_e)) \leq \frac{1}{4} f_e d_e(f_e) .$$

This proves the lemma. □

Price of Anarchy – Proof

(Correa, Schulz, Stier-Moses, 2008)

Application of the lemma implies

$$\begin{aligned} C(f) &= C(g) + \sum_{e \in E} g_e (d_e(f) - d_e(g)) \\ &\leq C(g) + \frac{1}{4} \sum_{e \in E} f_e d_e(f) \\ &= C(g) + \frac{1}{4} C(f) . \end{aligned}$$

This proves $\frac{3}{4}C(f) \leq C(g)$ for every flow g , and the theorem follows. \square

Theorem (Roughgarden, Tardos, 2002)

The price of anarchy in Wardrop games with latency functions given by positive polynomials of degree d is at most $d + 1$.

Existence and Uniqueness of Wardrop Equilibria

Theorem (Beckmann, McGuire, Winsten, 1956)

Every Wardrop game with continuous latency functions has at least one Wardrop equilibrium. For continuous and strictly increasing functions all edge flows in equilibrium are unique.

Proof Idea:

Wardrop games have a **potential function**

$$\Phi(f) = \sum_{e \in E} \int_{x=0}^{f_e} d_e(x) dx .$$

Consider a flow f^* that minimizes this potential. We formulate finding f^* as an optimization problem:

Existence and Uniqueness

$$\begin{array}{ll}
 \text{Minimize} & \sum_{e \in E} \int_{x=0}^{f_e} d_e(x) dx \\
 \text{subject to} & f_e = \sum_{P \in \mathcal{P}} f_P \quad \text{for every } e \in E \\
 & \sum_{P \in \mathcal{P}} f_P = 1 \\
 & f_P \geq 0 \quad \text{for every } P \in \mathcal{P}
 \end{array}$$

Every latency function d_e is continuous. Hence,

$$D_e(f_e) = \int_{x=0}^{f_e} d_e(x) dx$$

is continuous and differentiable. Thus, the problem has an optimal solution f^* .

Existence and Uniqueness

Since $f_e = \sum_{P \in \mathcal{P}} f_P$, the derivative yields

$$\frac{d}{df_P} \sum_{e \in E} \int_{x=0}^{f_e} d_e(x) dx = \sum_{e \in P} d_e(f) = d_P(f_P) .$$

We must distribute the rate of 1 optimally onto the paths. We decrease $\Phi(f)$ by decreasing f_P with high derivative and increasing $f_{P'}$ with small derivative. Hence, for any optimal flow one cannot move flow from paths with higher latency to paths with smaller latency.

Hence, for every optimal solution f^* , every path P with $f_P^* > 0$ and every arbitrary other path P' it holds that $d_P(f^*) \leq d_{P'}(f^*)$. An optimal solution f^* is a Wardrop equilibrium.

Existence and Uniqueness

If d_e is strictly increasing in f_e , then so is d_P in f_P , and $\Phi(f)$ becomes strictly convex in f_P . If $\Phi(f)$ is strictly convex, then for all optimal solutions the value f_e^* is unique, for every $e \in E$. \square

Remarks:

- ▶ Unique edge flows \Rightarrow **unique social cost** in equilibrium.
Hence, in these games: Price of anarchy = Price of stability.
- ▶ Uniqueness does **not hold** for the decomposition into path flows f_P .
- ▶ For minimization of $\Phi(f)$ we can use the **Ellipsoid method** for convex minimization. The algorithm takes **polynomial time** to compute the Wardrop flow (more precisely: a Wardrop flow wrt. any fixed numerical precision).

Corollary

A Wardrop equilibrium can be computed in polynomial time.

Wardrop Games

Price of Anarchy

Smoothness

Price of Stability

Price of Anarchy in Finite Games

Price of Anarchy for Nash equilibria:

- ▶ Strategic game Γ , **social cost** $\text{cost}(s)$ for every state s of Γ
- ▶ Consider Σ^{PNE} as the set of pure Nash equilibria of Γ
- ▶ **Price of Anarchy** is the ratio:

$$PoA = \frac{\max_{s' \in \Sigma^{PNE}} \text{cost}(s')}{\min_{s \in \Sigma} \text{cost}(s)}$$

PoA is a worst-case ratio and measures how much the worst PNE costs in comparison to an optimal state of the game.

Assumption


We here choose $\text{cost}(s) = \sum_{i \in \mathcal{N}} c_i(s)$ throughout.

Is there a technique to bound the price of anarchy in many games?


Congestion Games with Linear Delays

PoA in congestion games with linear delays $d_r(x) = a_r \cdot x$, for $a_r \geq 0$:


The following game has 4 players going from (1) u to w , (2) w to v , (3) v to u and (4) u to v . Each player has a short (direct edge) and a long (along the 3rd vertex) strategy:



Congestion Games with Linear Delays

Optimum s^* 

$$\text{cost}(s^*) = 1 + 1 + 1 + 1 = 4$$

A bad PNE s 

$$\text{cost}(s) = 3 + 2 + 2 + 3 = 10$$

PoA in this game at least 2.5. **Is this the worst case?**

Congestion Games with Linear Delays

We prove that the PoA is at most 2.5:

Let s be the worst PNE. If player i in s deviates to another strategy, her cost will not decrease. Let s^* be an optimal state. It holds, in particular, that $c_i(s) \leq c_i(s_i^*, s_{-i})$.

This allows to bound the social cost by

$$\text{cost}(s) = \sum_{i \in \mathcal{N}} c_i(s) \leq \sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) . \quad (1)$$

This is an entangled sum – for every player we consider her *individual cost* if *she alone* would deviate to her strategy in s^* . How can we relate this term to $\text{cost}(s)$ and $\text{cost}(s^*)$?

Congestion Games with Linear Delays

Let's consider this term for congestion games. We use $n_r = n_r(s)$ for the load of resource r in state s and $n_r^* = n_r(s^*)$ for the load in s^* .

If player i deviates to s_i^* , then she experiences on every resource $r \in s_i^*$ a load of at most $n_r + 1$ (possibly only n_r if $r \in s_i \cap s_i^*$)

$$\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \leq \sum_{i \in \mathcal{N}} \sum_{r \in s_i^*} d_r(n_r + 1) .$$

Since exactly n_r^* players consider resource r in their deviation, it holds:

$$\sum_{i \in \mathcal{N}} \sum_{r \in s_i^*} d_r(n_r + 1) = \sum_{r \in \mathcal{R}} n_r^* d_r(n_r + 1) = \sum_{r \in \mathcal{R}} n_r^* a_r \cdot (n_r + 1) .$$

We use the following lemma without proof:

Lemma (Christodoulou, Koutsoupias, 2005)

For all non-negative integer numbers $y, z \in \{0, 1, 2, 3, \dots\}$

$$y(z + 1) \leq \frac{5}{3} \cdot y^2 + \frac{1}{3} \cdot z^2 .$$

Congestion Games with Linear Delays

Using $y = n_r^*$ and $z = n_r + 1$ we obtain:

$$\begin{aligned}
 \sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) &\leq \sum_{r \in \mathcal{R}} a_r n_r^* (n_r + 1) \\
 &\leq \sum_{r \in \mathcal{R}} a_r \left(\frac{5}{3} (n_r^*)^2 + \frac{1}{3} n_r^2 \right) \\
 &= \frac{5}{3} \sum_{r \in \mathcal{R}} n_r^* (a_r n_r^*) + \frac{1}{3} \sum_{r \in \mathcal{R}} n_r (a_r n_r) \\
 &= \frac{5}{3} \cdot \text{cost}(s^*) + \frac{1}{3} \cdot \text{cost}(s)
 \end{aligned}$$

and, hence,

$$\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \leq \frac{5}{3} \cdot \text{cost}(s^*) + \frac{1}{3} \cdot \text{cost}(s) \tag{2}$$

Congestion Games with Linear Delays

This allows to bound the price of anarchy as follows:

$$\begin{aligned} \text{cost}(s) &\leq \frac{5}{3} \cdot \text{cost}(s^*) + \frac{1}{3} \cdot \text{cost}(s) \\ \Rightarrow \text{PoA} = \frac{\text{cost}(s)}{\text{cost}(s^*)} &\leq \frac{5/3}{1 - 1/3} = 2.5 \end{aligned} \quad (3)$$

Based on this proof, we can formulate a **framework**:

1. Set up inequality (1). It relies only on the definition of $\text{cost}(s)$ and on the fact that s is a PNE.
2. Derive inequality (2) with a pair of numbers (λ, μ) . The pair $(5/3, 1/3)$ was specific for linear congestion games.
3. Conclude the calculation with inequality (3). The bound on the PoA depends only on the pair of numbers in inequality (2).

Wardrop Games

Price of Anarchy

Smoothness

Price of Stability

Smoothness

The only game-specific information in this framework are the numbers in inequality (2). For a given game, if we can show an inequality of this kind, then we can directly obtain a bound on the PoA using the framework.

Definition

A game is called (λ, μ) -smooth for $\lambda > 0$ and $\mu \leq 1$ if, for every pair of states $s, s' \in \Sigma$, we have

$$\sum_{i \in \mathcal{N}} c_i(s'_i, s_{-i}) \leq \lambda \cdot \text{cost}(s') + \mu \cdot \text{cost}(s) . \quad (4)$$

The framework implies directly:

Theorem

In a (λ, μ) -smooth game, the PoA for pure Nash equilibria is at most

$$\frac{\lambda}{1 - \mu} .$$

Beispiel: Wardropspiele

The proof for the PoA in Wardrop games was also using the framework.

In the first step

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

we establish an inequality equivalent to (1). The term on the right is essentially an “entangled sum” assuming that in equilibrium f every infinitesimally small player deviates individually to her strategy in an optimal flow g .

Using the lemma and the picture proof we see that for every pair f, g of flows

$$\sum_e g_e d_e(f_e) \leq \sum_{e \in E} g_e d_e(f_e) + \frac{1}{4} \sum_{e \in E} f_e d_e(f_e) = 1 \cdot C(g) + \frac{1}{4} \cdot C(f) .$$

Thus: Wardrop games with affine delays are $(1, 1/4)$ -smooth.

A conclusion as in (3) yields $C(f)/C(g) \leq 1/(1 - 1/4) = 4/3$.

Beyond PNE?

Many games do not have PNE. What is the price of anarchy for more general equilibrium concepts?


What is the deterioration in social cost when all players play with no-regret algorithms?

Definition

The **price of anarchy for coarse-correlated equilibria** or **price of total anarchy** is the smallest $\rho \geq 1$ such that, for every coarse-correlated equilibrium \mathcal{V} and every state s' of the game we have

$$\mathbb{E}_{s \sim \mathcal{V}}[\text{cost}(s)] \leq \rho \cdot \text{cost}(s') .$$

A Hierarchy of Equilibrium Concepts



In every game, coarse-correlated, correlated, mixed, and pure Nash equilibria form a hierarchy of inclusion. As such, the PoA for more general equilibria can only be higher.

Smoothness can do more...

In the framework, we only used that (2) holds for a *pure Nash equilibrium* s and a *social optimum* s^* . In the definition of smoothness, however, we require that it holds for *every pair of states*. This stronger property allows to prove

Theorem

In a (λ, μ) -smooth game, the PoA for coarse-correlated equilibria is at most

$$\frac{\lambda}{1 - \mu} .$$

Proof:

Let s^* be an optimal state. We obtain step (1) as follows:

$$\begin{aligned} \mathbb{E}_{s \sim \mathcal{V}}[\text{cost}(s)] &= \mathbb{E}_{s \sim \mathcal{V}} \left[\sum_{i \in \mathcal{N}} c_i(s) \right] = \sum_{i \in \mathcal{N}} \mathbb{E}_{s \sim \mathcal{V}}[c_i(s)] \\ &\leq \sum_{i \in \mathcal{N}} \mathbb{E}_{s \sim \mathcal{V}}[c_i(s_i^*, s_{-i})] = \mathbb{E}_{s \sim \mathcal{V}} \left[\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \right] \end{aligned}$$

Price of Anarchy for Coarse-Correlated Equilibria

We now use the smoothness property pointwise for s^* and all possible states in the support of the distribution \mathcal{V}

$$\begin{aligned}\mathbb{E}_{s \sim \mathcal{V}} \left[\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \right] &\leq \mathbb{E}_{s \sim \mathcal{V}} [\lambda \cdot \text{cost}(s^*) + \mu \cdot \text{cost}(s)] \\ &\leq \lambda \cdot \text{cost}(s^*) + \mu \cdot \mathbb{E}_{s \sim \mathcal{V}}[\text{cost}(s)]\end{aligned}$$

This corresponds to step (2).

The remaining calculation is exactly the same as (3):

$$\begin{aligned}\mathbb{E}_{s \sim \mathcal{V}}[\text{cost}(s)] &\leq \lambda \text{cost}(s^*) + \mu \cdot \mathbb{E}_{s \sim \mathcal{V}}[\text{cost}(s)] \\ \Rightarrow \frac{\mathbb{E}_{s \sim \mathcal{V}}[\text{cost}(s)]}{\text{cost}(s^*)} &\leq \frac{\lambda}{1 - \mu}\end{aligned}$$



Smoothness Examples

Theorem

Every Wardrop routing game with affine delay functions is $(1, \frac{1}{4})$ -smooth. Thus, the PoA for coarse-correlated equilibria is upper bounded by $4/3$.


Theorem

Every congestion game with affine delay functions is $(\frac{5}{3}, \frac{1}{3})$ -smooth. Thus, the PoA for coarse-correlated equilibria is upper bounded by 2.5.

In both cases, the upper bound applies to coarse-correlated equilibria. Also, for both classes of games there are cases, in which this bound is attained by a pure Nash equilibrium. We call game classes with this property **tight**.

The Price of Anarchy for (λ, μ) -smooth Games

Since they apply to coarse-correlated equilibria, smoothness proofs yields an upper bound for a variety of equilibrium concepts. Also, as a matter of fact, the smoothness bound $\lambda/(1 - \mu)$ applies under conditions slightly more general than coarse-correlated equilibria:




The Price of Anarchy for **tight** (λ, μ) -smooth Games

In contrast, for **tight** classes of games, there is a game with a pure NE that yields price of anarchy of $\lambda/(1 - \mu)$. Hence, in this case the hierarchy collapses:

Optimum

1

Worst PNE, Worst MNE,
Worst Correlated, Worst
Coarse-Correlated, $\frac{\lambda}{1-\mu}$



Tightness in Congestion Games

Theorem (Roughgarden, 2003, Informal)

For a large class of non-decreasing, non-negative latency functions, the PoA for pure Nash equilibria in Wardrop games is $\lambda/(1 - \mu)$, and it is achieved on a two-node, two-link network (like Pigou's example).

Theorem (Roughgarden, 2009, Informal)

For a large class of non-decreasing, non-negative delay functions, the PoA for pure NE in congestion games is $\lambda/(1 - \mu)$, and it is achieved on an instance consisting of two cycles with possibly many nodes (like the example for affine delays above).

Thus, we have tightness and universal worst-case network structures in large classes of Wardrop and congestion games.

Limits of Smoothness

Smoothness yields informative bounds if good social cost in a game depends only mildly on coordinated behavior of players.

Consider the following game, in which small social cost depends strongly on the coordination of the strategy choices of both players:

	A	B
A	1	100
B	100	1

Limits of Smoothness

	A	B
A	1	100
B	100	1

- ▶ Optimal states are PNE with social cost 2.
- ▶ PoA for PNE is 1.
- ▶ Worst mixed NE is $x_1 = x_2 = (0.5, 0.5)$.
- ▶ Expected social cost 101. PoA for mixed NE is 50.5

Limits of Smoothness

	A	B
A	1	100
B	100	1

- ▶ Every coarse-correlated equilibrium \mathcal{V} must fulfill

$$\begin{aligned} & \max(\Pr[s = (A, B)], \Pr[s = (B, A)]) \\ & \leq \min(\Pr[s = (A, A)], \Pr[s = (B, B)]) \end{aligned}$$

- ▶ Worst coarse-correlated equilibrium is uniform distribution (hence, the worst mixed NE).
- ▶ PoA for coarse-correlated equilibria is 50.5

Limits of Smoothness

	A	B
A	1	100
B	100	1

For bounds on λ and μ consider PNE $s = (A, A)$ and optimum $s^* = (B, B)$:

$$200 = c_1(B, A) + c_2(A, B) = \sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \leq \lambda \cdot \text{cost}(s^*) + \mu \cdot \text{cost}(s) = 2\lambda + 2\mu$$

Then we have $\lambda \geq 100 - \mu$ and $PoA \leq (100 - \mu)(1 - \mu) = 1 + 99/(1 - \mu)$.
With $\mu = 0$ we obtain 100 as smallest upper bound.

Limits of Smoothness

	A	B
A	1	100
B	100	1

Even if we only consider the smoothness condition w.r.t. PNE and optimal states, we obtain a bound of 100 for the PoA.

This is **100 times larger** than the real PoA for PNE.

Also, it's almost $|\mathcal{N}|$ **times larger** than the PoA for coarse-correlated equilibria.

In this game we need **coordination for PNE and small social cost**. In $\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i})$ we only have states with a **cost much larger** than in a PNE or an optimum. This implies that **λ must be increased substantially** and the resulting bound becomes unrealistic.

Wardrop Games

Price of Anarchy

Smoothness

Price of Stability

Decreasing Delays: Equal-Sharing Games

Equal-Sharing Game

- ▶ Set \mathcal{N} of n players, set \mathcal{R} of m resources
- ▶ Player i allocates subset of resources, strategy set $\Sigma_i \subseteq 2^{\mathcal{R}}$
- ▶ Resource $r \in \mathcal{R}$ has fixed cost $c_r \geq 0$.
- ▶ Cost c_r is assigned in equal shares to the players allocating r (if any).

Equal-sharing games are **congestion games with delays** $d_r(x) = c_r/x$.

Social cost turns out to be the sum of costs of resources allocated by at least one player:

$$\text{cost}(S) = \sum_{i \in \mathcal{N}} c_i(S) = \sum_{i \in \mathcal{N}} \sum_{r \in S_i} d_r(n_r) = \sum_{\substack{r \in \mathcal{R} \\ n_r \geq 1}} n_r \cdot c_r / n_r = \sum_{\substack{r \in \mathcal{R} \\ n_r \geq 1}} c_r .$$

Price of Stability

Theorem

Every equal-sharing game is $(n, 0)$ -smooth. Thus, the PoA for coarse-correlated equilibria is at most n . The class of equal-sharing games is tight, i.e., there are games in which the PoA for pure Nash equilibria is exactly n .

The PoA is large, but PNE are not necessarily unique. What do other PNE cost, what about the **best one**?

Price of Stability for pure Nash equilibria:

- ▶ Consider Σ^{PNE} as the set of pure Nash equilibria of a game Γ
- ▶ **Price of Stability** is the ratio:

$$PoS = \frac{\min_{S' \in \Sigma^{PNE}} \text{cost}(S')}{\text{cost}(S^*)} .$$

PoS is a best-case ratio. It measures the social cost of the cheapest PNE and expresses by how much it is worse than the optimal social cost in the game.

Price of Stability in Equal-Sharing Games

Theorem

For every equal-sharing game the price of stability for pure Nash equilibria is at most $\mathcal{H}_n = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} = O(\log n)$.

Proof:

Rosenthal's potential function in equal-sharing games is bounded by

$$\begin{aligned}
 \Phi(S) &= \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r} c_r / i &= \sum_{\substack{r \in \mathcal{R} \\ n_r \geq 1}} c_r \cdot \left(1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n_r} \right) \\
 &&\leq \sum_{\substack{r \in \mathcal{R} \\ n_r \geq 1}} c_r \cdot \mathcal{H}_n \\
 &&= \text{cost}(S) \cdot \mathcal{H}_n .
 \end{aligned}$$

Price of Stability in Equal-Sharing Games

In $\Phi(S)$ we account for each player allocating resource r a contribution of c_r/i for some $i = 1, \dots, n_r$. In her cost $c_i(S)$ we account only c_r/n_r . This shows that for every state S in the game

$$\text{cost}(S) \leq \Phi(S) \leq \text{cost}(S) \cdot \mathcal{H}_n .$$

Now suppose we start at an optimum S^* and iteratively perform improvement steps for single players. This leads to some PNE S . Every such move decreases the potential function. We thus have $\Phi(S) \leq \Phi(S^*)$ and

$$\text{cost}(S) \leq \Phi(S) \leq \Phi(S^*) \leq \text{cost}(S^*) \cdot \mathcal{H}_n .$$

Hence, there is at least one PNE that is only a factor of \mathcal{H}_n more costly than the optimum S^* . □

A similar proof technique (called *potential function method*) can be applied to arbitrary potential games and gives upper bounds on the price of stability.


A tight lower bound on the PoS

Theorem

There is an equal-sharing game with price of stability arbitrarily close to $\mathcal{H}_n = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n}$.

Proof:

For an arbitrarily small $\epsilon > 0$ consider the following example.



For every $i = 1, \dots, n$ the strategy set Σ_i of player i is the set of paths from s_i to t .

Consider any PNE. There can be no player on the edge with cost $1 + \epsilon$. Thus, the price of stability is arbitrarily close to \mathcal{H}_n . \square

Arbitrary Sharing

Equal sharing games are congestion games and have nice structural properties. However, in many cases there is no central authority that dictates cost shares. Instead, players themselves have to come up with a way to share the cost of an investment. This motivates the study of cost sharing games with **arbitrary sharing**.

Resource Buying Games

- ▶ Set \mathcal{N} of n players, set \mathcal{R} of m resources
- ▶ Player i allocates subset of resources, has **action set** $A_i \subseteq 2^{\mathcal{R}}$
- ▶ Resource $r \in \mathcal{R}$ has **fixed cost** $c_r \geq 0$.
- ▶ Cost c_r must be **shared (somehow)** by players allocating r (if any).

Arbitrary Sharing

Strategies and Costs:

- ▶ i picks subset of resources $S_i \in A_i$, and **payment** $p_{ir} \geq 0$ for each $r \in S_i$.
- ▶ Strategy is pair (S_i, p_i) , resource set and vector of payments
- ▶ Resource r is **bought** if

$$\sum_{i:r \in S_i} p_{ir} \geq c_r$$

- ▶ Cost of player i in state (S, p) :

$$c_i(S, p) = \begin{cases} \infty & \text{if some } r \in S_i \text{ is not bought} \\ \sum_{r \in S_i} p_{ir} & \text{otherwise} \end{cases}$$

Preliminaries

In state (S, p) we define the **set of allocated resources** $R_S = \{r \in R \mid r \in S_i\}$.

Consider a **PNE** (S, p) . Then...

- ▶ each $r \in R_S$ is **bought exactly**,
 - ▶ each $r \in R \setminus R_S$ is **not paid at all**, so
- $$\Rightarrow \text{cost}(S, p) = \sum_i c_i(S, p) = \sum_{r \in R_S} c_r$$

Consider an **optimal state** (S^*, p^*) that minimizes $\text{cost}(S, p)$. Then...

- ▶ each $r \in R_{S^*}$ is **bought exactly**,
 - ▶ each $r \in R \setminus R_{S^*}$ is **not paid at all**, so
- $$\Rightarrow \text{cost}(S^*, p^*) = \sum_i c_i(S^*, p^*) = \sum_{r \in R_{S^*}} c_r.$$

Preliminaries

In state (S, p) we define the **set of allocated resources** $R_S = \{r \in R \mid r \in S_i\}$.

Consider a **PNE** (S, p) . Then...

- ▶ each $r \in R_S$ is **bought exactly**,
 - ▶ each $r \in R \setminus R_S$ is **not paid at all**, so
- $$\Rightarrow \text{cost}(S, p) = \sum_i c_i(S, p) = \sum_{r \in R_S} c_r$$

Consider an **optimal state** (S^*, p^*) that minimizes $\text{cost}(S, p)$. Then...

- ▶ each $r \in R_{S^*}$ is **bought exactly**,
 - ▶ each $r \in R \setminus R_{S^*}$ is **not paid at all**, so
- $$\Rightarrow \text{cost}(S^*, p^*) = \sum_i c_i(S^*, p^*) = \sum_{r \in R_{S^*}} c_r.$$

Why?

Existence, Complexity, and Cost of Equilibria

In every PNE and every optimal state, the players exactly share the cost of the allocated subset of resources. The **sharing rule**, however, is not fixed and subject to strategic choice by the players.

Some questions:

- ▶ Do PNE always exist?
- ▶ Can we compute them in polynomial time?
- ▶ How large are prices of anarchy and stability?

Existence, Complexity, and Cost of Equilibria

In every PNE and every optimal state, the players exactly share the cost of the allocated subset of resources. The **sharing rule**, however, is not fixed and subject to strategic choice by the players.

Some questions:

- ▶ Do PNE always exist?
- ▶ Can we compute them in polynomial time?
- ▶ How large are prices of anarchy and stability?


Proposition

There is a resource buying game with two players that does not have any PNE.

Existence

Proof:

Consider the following network game. Resources are edges, edge labels represent costs. Strategies of player i are s_i - t_i -paths, for $i = 1, 2$.



Suppose (S, p) is a PNE. Exactly three edges are allocated and must be bought.

W.l.o.g. assume $e = (s_1, t_2)$ is not bought. Then (s_1, s_2) is bought entirely by 1 and (t_2, t_1) by 2.

If some player i pays $\varepsilon > 0$ to (s_2, t_1) , she has a cost of at least $1 + \varepsilon$. She can deviate to the other path and buy only e at a cost of 1 instead. Hence, (s_2, t_1) is not bought, so (S, p) cannot be a PNE. \square

Results

Games with arbitrary sharing are not congestion games.

Moreover, there are classes of games with n players, such that...

- ▶ Deciding existence of a PNE is NP-hard
- ▶ A PNE exists, and the price of anarchy is as large as n .
- ▶ A PNE exists, and the price of stability is as large as $n - 2$.

Interesting special case: **Single-Source Network Games**

- ▶ Directed network $G = (V, E)$
- ▶ Edges are resources, $R = E$, each with a fixed cost $c_e \geq 0$
- ▶ Each player has a source vertex $s_i \in V$
- ▶ There is a global target vertex $t \in V$
- ▶ Strategies for i are the s_i - t -paths in G

Price of Stability with Arbitrary Sharing in Single-Source Networks

Theorem

Every single-source network game has a PNE. The price of stability is 1.

Note: In **games with equal sharing**, the price of stability is H_n even for single-source networks.

Proof:

Consider an optimal state (S^*, p^*) . The allocated edges E_{S^*} form the cheapest network to connect all s_i to t . This is called an **optimal Steiner tree**. We use the short notation $T^* = E_{S^*}$.

How do we share the cost of T^* such that no player wants to deviate?

We assign payments using a **bottom-up** approach. For an edge e , we denote $N_e = \{i \mid e \in S_i\}$, i.e., these players have e in their path from s_i and t in T^* . Initially, all payments $p_{ie} = 0$ for all $e \in T^*$ and $i \in \mathcal{N}$.

Paying for an optimal Steiner tree T^*

Consider T^* rooted in t and the edges $e \in T^*$ bottom-up.
For each e in this order do the following:

- ▶ For each $i \in N_e$:
- ▶ Define subtree cost c^i for each $e' \in E \setminus \{e\}$ as follows:

$$c_{e'}^i = \begin{cases} p_{ie'} & \text{for each } e' \in S_i \text{ below } e \text{ (previously assigned payments)} \\ 0 & \text{for each } e' \in T^* \text{ not below } e \text{ (rest of } T^*) \\ c_{e'} & \text{for each } e' \notin T^* \text{ (rest of } G) \end{cases}$$

- ▶ Compute deviation path, i.e., cheapest s_i - t -path in $G - \{e\}$ w.r.t. c^i
- ▶ $\Delta_i^e =$ cost of the deviation path w.r.t. c^i
- ▶ $b_i^e = \Delta_i^e - \sum_{e' \in S_i} p_{ie'}$ (budget of i for edge e)
- ▶ Assign $p_{ie} = \min \left\{ b_i^e, c_e - \sum_{i' \neq i} p_{i'e} \right\}$

Deviation Paths

Consider a deviation path for some player i . It has the following structure:

1. It starts at s_i and departs from S_i at a deviation vertex d_i . It has to leave S_i , since it cannot use e . Note that $s_i = d_i$ is possible.
2. Starting in d_i it has to use edges in $E \setminus T^*$ until it again reaches a node in T^* , the reentry vertex r_i . Note r_i exists since the path ends in t .
3. **By induction**, we assume i has no profitable deviation for a subpath below e . Hence, the reentry vertex of i must not be below e . Thus, the r_i - t -path in T^* has cost 0 under c^i . So the deviation path uses this cost-0-path from r_i to t inside T^* .

Clearly, b_i^e is the maximum amount i can contribute to e before the deviation path becomes cheaper than S_i under cost c^i .

Deviation Points

If the **budgets of players in N_e suffice** to pay for c_e , the algorithm advances to the next higher edge e' . Note that we maintain the inductive hypothesis that there are no profitable deviations of any player w.r.t. to payments to T^* below e' .

Otherwise, suppose the **budgets do not suffice**, i.e. $c_e > \sum_{i \in N_e} b_i^e$. Then consider the deviation points d_i for all $i \in N_e$.

Let D be the (unique) subset of **highest** deviation points, i.e., on each path from some s_j to e pick the deviation point (of potentially some other player $i \neq j$) closest to e .

For each $d_i \in D$ consider the second part of the player's deviation path connecting d_i to r_i . We denote the part by Z_i and its' cost by b_i . It is the total original cost of Z_i . b_i generates the budget for i for all edges in S_i from d_i up to (and including) e .

Improving T^*

Consider e , as well as all edges from T^* that are below e and above any deviation point $d_i \in D$. We denote this set by

$$H = \{e' \in T^* \mid e' = e, \text{ or } e' \text{ below } e \text{ and above } d_i \in D\}$$

The budgets do not suffice to pay for e , but we managed to pay exactly for all edges below e . Hence,

$$\begin{aligned} c_e &> \sum_{i \in N_e} b_i^e \\ &\geq \sum_{i: d_i \in D} b_i^e \\ &= \sum_{i: d_i \in D} b_i - \sum_{e' \in H \cap S_i} p_{ie'} \\ &\geq \sum_{i: d_i \in D} b_i - \sum_{e' \in H \setminus \{e\}} c_{e'} \end{aligned}$$

Improving T^*

Let $Z = \bigcup_{i:d_i \in D} Z_i$. Then we have

$$\sum_{e' \in H} c_{e'} > \sum_{i:d_i \in D} b_i \geq \sum_{e' \in Z} c_{e'} \quad (5)$$

Now consider the new network $(T^* \setminus H) \cup Z$. This network is feasible in the sense that it connects all s_i to t . By (5) it has strictly cheaper cost than T^* . This is a contradiction, since T^* is the optimal Steiner tree.

Hence, the budgets of players in N_e must always suffice to pay for e .

This proves that the bottom-up algorithm terminates with payments that constitute a PNE buying T^* .

As such, a PNE always exists, and the price of stability is 1. □

Polynomial Time?

The proof also works for games in undirected graphs G . Unfortunately, computing an optimal Steiner tree T^* is NP-hard, in both undirected and directed graphs.

We can apply the previous network improvement step to compute a $(1 + \varepsilon)$ -approximate PNE that buys an α -approximate Steiner tree, where α is the factor of the best-known approximation algorithm (currently 1.39 for undirected graphs).

The application requires to temporarily subsidize the edge costs in order to maintain that the improvement steps make sufficient progress in terms of cost.

More Extensions

Tree Connection Games:

- ▶ Undirected G , each $i \in \mathcal{N}$ wants to connect individual s_i and t_i
- ▶ Imagine we build a separate *requirement graph* G_C :
For every $i \in \mathcal{N}$, insert s_i, t_i and connect them with an edge.
- ▶ If G_C is connected, every PNE must be a tree.
- ▶ Then the game has an optimal PNE.

Two-Source Games:

- ▶ Each player $i \in \mathcal{N}$ has sources s_{i1} and s_{i2}
- ▶ There is a global target node t
- ▶ Action set A_i contains all edge sets that connect s_{i1}, s_{i2}, t
- ▶ PNE existence can be NP-hard to decide

Matroid Strategies

- ▶ Each action set A_i is composed of the bases of a matroid
- ▶ Then the game has an optimal PNE.

Recommended Literature

- ▶ Nisan, Roughgarden, Tardos, Vazirani. Algorithmic Game Theory, 2007. (Chapters 18 and 19.3).
- ▶ Awerbuch, Azar, Epstein. The Price of Routing Unsplittable Flow. SIAM Journal on Computing 42(1):160–177, 2013.
- ▶ Christodoulou, Koutsoupias. The Price of Anarchy of Finite Congestion Games. STOC 2005.
- ▶ Blum, Hajiaghayi, Ligett, Roth. Regret Minimization and the Price of Total Anarchy. STOC 2008.
- ▶ Roughgarden. Intrinsic Robustness of the Price of Anarchy. Journal of the ACM 62(5):32:1–32:42, 2015.
- ▶ Roughgarden. Twenty Lectures on Algorithmic Game Theory, 2016. (Chapter 14)

Recommended Literature

- ▶ Anshelevich, Dasgupta, Kleinberg, Roughgarden, Tardos, Wexler. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing* 38(4):1602–1623, 2008.
- ▶ Anshelevich, Dasgupta, Tardos, Wexler. Near-Optimal Network Design with Selfish Agents. *Theory of Computing* 4(1):77–109, 2008.
- ▶ Hoefer. Non-cooperative Tree Creation. *Algorithmica* 53(1):104–131, 2009.
- ▶ Harks, Peis. Resource Buying Games. *Algorithmica* 70(3):493–512, 2014.