Price of Anarchy

Instructor: Thomas Kesselheim

One of the main goals of algorithmic game theory is to quantify the performance of a system of selfish agents. Usually the “social cost” incurred by all players is higher than if there is a central authority taking charge to minimize social cost. We will develop tools that will allow us to (upper and lower) bound the potential increase.

Here we will define *social cost* as the sum of all players' cost; formally, for a state s let $cost(s) = \sum_{i \in \mathcal{N}} c_i(s)$ denote the social cost of s. Sometimes it makes more sense to consider the *maximum* cost incurred by any player.

1 Motivating Example

Example 6.1 (Pigou's Example, Discrete Version)**.** *Consider the following symmetric network congestion game with four players.*

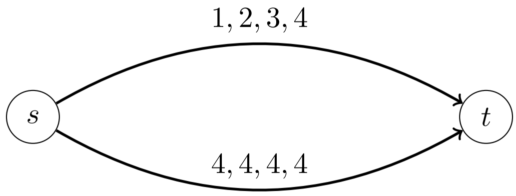

There are five kinds of states:

(a) all players use the top edge, social cost: 16

(b) three players use the top edge, one player uses the bottom edge, social cost: 13

(c) two players use the top edge, two players use the bottom edge, social cost: 12

(d) one player uses the top edge, three players use the bottom edge, social cost: 13

(e) all players use the bottom edge, social cost: 16

Observe that only states of kind (a) and (b) can be pure Nash equilibria. The social cost, however, is minimized by states of kind (c). Therefore, when considering pure Nash equilibria, due to selfish behavior, we lose up to a factor of $\frac{16}{12}$ *and at least a factor of* $\frac{13}{12}$.

More generally, we refer to the worst-case ratio between the social cost at equilibrium and the optimal social cost as the price of anarchy. The best-case ratio between these two quantities is the price of stability.

2 Definition

The price of anarchy and the price of stability of course depend on what kind of equilibria we consider. To define these two notions formally, we therefore assume that there is a set Eq of probability distributions over the set of states S, which correspond to equilibria. In the case of pure Nash equilibria, each of these distributions concentrates all its mass on a single point.

Definition 6.2. *Given a cost-minimization game, let* Eq *be a set of probability distributions over the set of states* S*. For some probability distribution* p*, let* $cost(p) = \sum_{s \in S} p(s) cost(s)$ *be the expected social cost. The* price of anarchy for Eq *is defined as*

$$PoA_{\mathsf{Eq}} = \frac{\max_{p \in \mathsf{Eq}} cost(p)}{\min_{s \in S} cost(s)} \quad .$$

The price of stability for Eq *is defined as*

$$PoS_{\mathsf{Eq}} = \frac{\min_{p \in \mathsf{Eq}} cost(p)}{\min_{s \in S} cost(s)} \quad .$$

Given the respective equilibria exist, we have

$$1 \leq PoS_{\mathsf{CCE}} \leq PoS_{\mathsf{CE}} \leq PoS_{\mathsf{MNE}} \leq PoS_{\mathsf{PNE}} \leq PoA_{\mathsf{PNE}} \leq PoA_{\mathsf{MNE}} \leq PoA_{\mathsf{CE}} \leq PoA_{\mathsf{CCE}} \quad .$$

3 Smooth Games

A very helpful technique to derive upper bounds on the price of anarchy is *smoothness*.

Definition 6.3. *A game is called* (λ, μ)-smooth *for* $\lambda > 0$ *and* $\mu < 1$ *if, for every pair of states* $s, s^* \in S$*, we have*

$$\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \leq \lambda \cdot cost(s^*) + \mu \cdot cost(s) \quad .$$

Observe that this condition needs to hold for *all* states $s, s^* \in S$, as opposed to only pure Nash equilibria or only social optima. We consider the cost that each player incurs when unilaterally deviating from s to his strategy in s^*. If the game is smooth, then we can upper-bound the sum of these costs in terms of the social cost of s and s^*.

Smoothness directly gives a bound for the price of anarchy, even for coarse correlated equilibria.

Theorem 6.4. *In a* (λ, μ)*-smooth game, the PoA for coarse correlated equilibria is at most*

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

Proof. Let s be distributed according to a coarse correlated equilibrium p, and let s^* be an optimum solution, which minimizes social cost. Note that $cost(p) = \mathbf{E}_{s \sim p}[cost(s)]$. Then:

$$\begin{aligned}
\mathbf{E}_{s \sim p}[cost(s)] &= \sum_{i \in \mathcal{N}} \mathbf{E}_{s \sim p}[c_i(s)] && \text{(by linearity of expectation)} \\
&\leq \sum_{i \in \mathcal{N}} \mathbf{E}_{s \sim p}[c_i(s_i^*, s_{-i})] && \text{(as } p \text{ is a CCE)} \\
&= \mathbf{E}_{s \sim p}\left[\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i})\right] && \text{(by linearity of expectation)} \\
&\leq \mathbf{E}_{s \sim p}[\lambda \cdot cost(s^*) + \mu \cdot cost(s)] && \text{(by smoothness)}
\end{aligned}$$

On both sides subtract $\mu \cdot \mathbf{E}_{s \sim p}[cost(s)]$, this gives

$$(1 - \mu) \cdot \mathbf{E}_{s \sim p}[cost(s)] \leq \lambda \cdot cost(s^*)$$

and rearranging yields

$$\frac{\mathbf{E}_{s \sim p}[cost(s)]}{cost(s^*)} \leq \frac{\lambda}{1 - \mu} \quad .$$

□

That is, in a (λ, μ)-smooth game, we have

$$PoA_{\mathsf{PNE}} \leq PoA_{\mathsf{MNE}} \leq PoA_{\mathsf{CE}} \leq PoA_{\mathsf{CCE}} \leq \frac{\lambda}{1-\mu} \ .$$

For many classes of games, there are choices of λ and μ such that all relations become equalities. These games are referred to as *tight*.

4 Tight Bound for Affine Delay Functions

We next provide a tight bound on the price of anarchy for (non-decreasing) *affine delay functions* of the form $d_r(n_r(S)) = a_r \cdot n_r(S) + b_r$, where $a_r, b_r \geq 0$.

Theorem 6.5. *Every congestion game with affine delay functions is* $\left(\frac{5}{3}, \frac{1}{3}\right)$*-smooth. Thus, the PoA is upper bounded by* $\frac{5}{2} = 2.5$*, even for coarse-correlated equilibria.*

We use the following lemma:

Lemma 6.6 (Christodoulou, Koutsoupias, 2005)**.** *For all integers* $y, z \in \mathbb{Z}$ *we have*

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

Proof. Consider the case $y = 1$. Note that, as z is an integer, we have $(z-1)(z-2) \geq 0$. Therefore, we have

$$z^2 - 3z + 2 = (z-1)(z-2) \geq 0 \ ,$$

which implies

$$z \leq \frac{2}{3} + \frac{1}{3}z^2 \ ,$$

and therefore

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

Now consider the case $y > 1$. We now use

$$0 \leq \left(\sqrt{\frac{3}{4}}y - \sqrt{\frac{1}{3}}z\right)^2 = \frac{3}{4}y^2 + \frac{1}{3}z^2 - yz \ .$$

Using $y \leq \frac{y^2}{2}$, we get

$$y(z+1) = yz + y \leq \frac{3}{4}y^2 + \frac{1}{3}z^2 + \frac{1}{2}y^2 \leq \frac{5}{3}y^2 + \frac{1}{3}z^2 \ .$$

Finally, for the case $y \leq 0$, observe that the claim is trivial for $y = 0$ or $z \geq -1$, because $y(z+1) \leq 0 \leq \frac{5}{3} \cdot y^2 + \frac{1}{3} \cdot z^2$. In the case that $y < 0$ and $z < -1$, we use that $y(z+1) \leq |y|(|z|+1)$ and apply the bound for positive y and z shown above. □

Proof of Theorem 6.5. Given two states s and s^*, we have to bound

$$\sum_{i \in \mathcal{N}} c_i(s_i^*, s_{-i}) \ .$$

We have

$$c_i(s_i^*, s_{-i}) = \sum_{r \in s_i^*} d_r(n_r(s_i^*, s_{-i})) \ .$$

Furthermore, as all d_r are non-decreasing, we have $d_r(n_r(s_i^*, s_{-i})) \leq d_r(n_r(s) + 1)$. This way, we get

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

By exchanging the sums, we have

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

To simplify notation, we write n_r for $n_r(s)$ and n_r^* for $n_r(s^*)$. Recall that delays are $d_r(n_r) = a_r n_r + b_r$. In combination, we get

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

Let us consider the term for a fixed $r \in \mathcal{R}$. We have

$$n_r^*(a_r(n_r + 1) + b_r) = a_r n_r^*(n_r + 1) + b_r n_r^* \ .$$

Lemma 6.6 implies that

$$n_r^*(n_r + 1) \leq \frac{5}{3}(n_r^*)^2 + \frac{1}{3} n_r^2 \ .$$

Thus, we get

$$\begin{aligned} n_r^*(a_r(n_r + 1) + b_r) &\leq \frac{5}{3} a_r (n_r^*)^2 + \frac{1}{3} a_r n_r^2 + b_r n_r^* \\ &\leq \frac{5}{3} a_r (n_r^*)^2 + \frac{5}{3} b_r n_r^* + \frac{1}{3} a_r n_r^2 + \frac{1}{3} b_r n_r \\ &= \frac{5}{3} \left(a_r n_r^* + b_r\right) n_r^* + \frac{1}{3} \left(a_r n_r + b_r\right) n_r \ , \end{aligned}$$

where in the second step we used that $b_r \geq 0$. Summing up these inequalities for all resources $r \in \mathcal{R}$, we get

$$\begin{aligned} \sum_{r \in \mathcal{R}} n_r^*(a_r(n_r + 1) + b_r) \quad &\leq \quad \frac{5}{3} \sum_{r \in \mathcal{R}} n_r^*(a_r n_r^* + b_r) + \frac{1}{3} \sum_{r \in \mathcal{R}} n_r(a_r n_r + b_r) \\ &= \quad \frac{5}{3} \cdot cost(s^*) + \frac{1}{3} \cdot cost(s) \ , \end{aligned}$$

which shows $\left(\frac{5}{3}, \frac{1}{3}\right)$-smoothness. □

Theorem 6.7. *There are congestion games with affine delay functions whose price of anarchy for pure Nash equilibria is* $\frac{5}{2}$.

Proof sketch. We consider the following (asymmetric) network congestion game. Notation 0 or x on an edge means that $d_r(x) = 0$ or $d_r(x) = x$ for this edge.

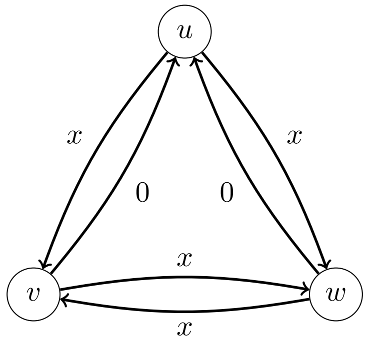

There are four players with different source sink pairs. Refer to this table for a socially optimal state of social cost 4 and a pure Nash equilibrium of social cost 10.

player	source	sink	strategy in OPT	cost in OPT	strategy in PNE	cost in PNE
1	u	v	$u \to v$	1	$u \to w \to v$	3
2	u	w	$u \to w$	1	$u \to v \to w$	3
3	v	w	$v \to w$	1	$v \to u \to w$	2
4	w	v	$w \to v$	1	$w \to u \to v$	2

□

Recommended Literature

- B. Awerbuch, Y. Azar, A. Epstein. The Price of Routing Unsplittable Flow. STOC 2005. (PoA for pure NE in congestion games)
- G. Christodoulou, E. Koutsoupias. The Price of Anarchy of finite Congestion Games. STOC 2005. (PoA for pure NE in congestion games)
- A. Blum, M. Hajiaghayi, K. Ligett, A. Roth. Regret Minimization and the Price of Total Anarchy. STOC 2008. (PoA for no-regret sequences)
- T. Roughgarden. Intrinsic Robustness of the Price of Anarchy. STOC 2009. (Smoothness Framework and unification of previous results)