Lecture #7 Notes Summary

Bounds on the Price of Anarchy for atomic and non-atomic congestion games

---

Price of Anarchy of non-atomic Games

The price of anarchy is a measure of how inefficient the equilibrium of a game can be. Consider a $N-$player game where the payoff function $\boldsymbol{\pi} : \Sigma \to \mathbb{R}^N$ is a nonnegative cost that players want to minimize (rather than *maximize*, as in the first lecture). We can define a general cost function $C : \Sigma \to \mathbb{R}_+$, such that $C(\boldsymbol{P})$ is a measure of *how bad* the strategy profile $\boldsymbol{P} = (P_1, \ldots, P_N) \in \Sigma$ is for the society. Typical choices are: $C(P_1, \ldots, P_N) = \sum_{i=1}^N \pi_i(P_1, \ldots, P_N)$ if we want to minimize the average cost, or $C(P_1, \ldots, P_N) = \max_{i=1}^N \pi_i(P_1, \ldots, P_N)$ if we want the outcome of the game to be as *fair* as possible.

Definition 1 (Price of Anarchy)**.** Consider a $N-$player game $\boldsymbol{\pi}$ with a cost function $C : \Sigma \to \mathbb{R}_+$. The *price of anarchy* (PoA) of the game is defined as the ratio between the efficiency of the worst Nash equilibrium and the best *centralized* solution (the social optimum). More precisely, let $\mathcal{E} \subseteq \Sigma$ denote the set of Nash equilibriums of the game. We have:

$$PoA(\boldsymbol{\pi}) := \frac{\max_{\boldsymbol{P} \in \mathcal{E}} C(\boldsymbol{P})}{\min_{\boldsymbol{P} \in \Sigma} C(\boldsymbol{P})} \geq 1.$$

Specifically, for non-atomic congestion games the price of anarchy compares the average travel time of the Wardrop equilibrium (which is uniquely defined) to the social optimum. Consider a non-atomic congestion game $\mathcal{G} = (E, \mathcal{S}, \boldsymbol{w}, \boldsymbol{d})$. We denote by F the set of feasible flows, by X the set of load vectors induced by a flow in F, and by T the set of pairs $\{(\boldsymbol{f}, \boldsymbol{x}) : \ \boldsymbol{f} \in F, \ \boldsymbol{x} \text{ induced by } \boldsymbol{f}\}$. For simplicity, we still call a flow-load pair $(\boldsymbol{f}, \boldsymbol{x}) \in T$ a *flow*. Recall that the total cost of a flow $(\boldsymbol{f}, \boldsymbol{x}) \in T$ only depends on the vector of loads $\boldsymbol{x}$:

$$C(\boldsymbol{x}) := \sum_{P \in \mathcal{S}} f_P c_P(\boldsymbol{x}) = \sum_{P \in \mathcal{S}} f_P \sum_{e \in P} d_e(\boldsymbol{x}) = \sum_{e \in E} d_e(\boldsymbol{x}) \sum_{P \ni e} f_P = \sum_{e \in E} x_e d_e(x_e).$$

Now, let $(\boldsymbol{f^*}, \boldsymbol{x^*})$ be a WE of the game and $(\boldsymbol{f^o}, \boldsymbol{x^o})$ be a social optimum. We have:

$$PoA(\mathcal{G}) := \frac{C(\boldsymbol{x^*})}{\min_{\boldsymbol{x} \in X} C(\boldsymbol{x})} = \frac{C(\boldsymbol{x^*})}{C(\boldsymbol{x^o})},$$

where the last expression does not depend on the choice of a social optimum $\boldsymbol{x^o}$. It does not depend on the choice of a WE $\boldsymbol{x^*}$ neither, which follows from

Proposition 1. *Let $(\boldsymbol{f}, \boldsymbol{x})$ and $(\boldsymbol{f'}, \boldsymbol{x'})$ be two WE of a non-atomic congestion game $\mathcal{G}$. Then, $C(\boldsymbol{x}) = C(\boldsymbol{x'})$.*

Proof. We know that $d_e(x_e) = d_e(x'_e)$ for all element e (by Theorem 6 of Lecture #6). So $c_P(\boldsymbol{x}) = c_P(\boldsymbol{x'})$ for all P, and

$$C(\boldsymbol{x'}) = \sum_{P \in \mathcal{S}} f'_P c_P(\boldsymbol{x'}) = \sum_{P \in \mathcal{S}} f'_P c_P(\boldsymbol{x}).$$

Now, recall that for all i, a WE puts weight only on the strategies P that minimize $c_P(\boldsymbol{x})$ over S^i. Denote by $L_i := \min_{P \in S^i} c_P(\boldsymbol{x})$ the minimal cost over S^i. We have:

$$C(\boldsymbol{x'}) = \sum_{i=1}^N \sum_{\{P \in S^i : f'_P > 0\}} L^i \ f'_P = \sum_{i=1}^N L_i \sum_{P \in S^i} f'_P = \sum_{i=1}^N L_i w_i,$$

which is also the value of $C(\boldsymbol{x})$. □

We also recall an alternative characterization of Wardrop Equilibriums (which was already encountered in Exercise 3 of Worksheet #6):

Proposition 2. *Consider a non-atomic congestion game $(E, \mathcal{S}, \boldsymbol{w}, \boldsymbol{d})$. The flow $(\boldsymbol{f}, \boldsymbol{x}) \in T$ is a Wardrop equilibrium iff*

$$\forall \boldsymbol{x'} \in X, \quad \sum_{e \in E} x_e \; d_e(x_e) \leq \sum_{e \in E} x'_e \; d_e(x_e).$$

Proof. By rearranging the terms in the sums, the condition of the proposition becomes:

$$\forall \boldsymbol{f'} \in F, \quad \sum_{P \in \mathcal{S}} c_P(\boldsymbol{x}) f_P \leq \sum_{P \in \mathcal{S}} c_P(\boldsymbol{x}) f'_P.$$

This means that $\boldsymbol{f}$ minimizes the expression $\sum_{P \in \mathcal{S}} c_P(\boldsymbol{x}) f_P$ over F (for the current load $\boldsymbol{x}$). Clearly, this is possible if and only if $\boldsymbol{f}$ puts weight only on the strategies P minimizing $c_P(\boldsymbol{x})$ over some S^i. □

We are now ready to prove our first result on the price of anarchy:

Theorem 3. *Let $\mathcal{G} = (E, \mathcal{S}, \boldsymbol{w}, \boldsymbol{d})$ be a non-atomic congestion game in which all delay functions are affine ($d_e(x) = a_e x + b_e$ for some $a_e, b_e \geq 0$). Then,*

$$PoA(\mathcal{G}) \leq \frac{4}{3}.$$

Proof. Let $(\boldsymbol{f^*}, \boldsymbol{x^*})$ be a WE and $(\boldsymbol{f^o}, \boldsymbol{x^o})$ be a social optimum of $\mathcal{G}$. Now, consider an arbitrary feasible flow $(\boldsymbol{f}, \boldsymbol{x}) \in T$. We have

$$C(\boldsymbol{x^*}) = \sum_{e \in E} d_e(x_e^*) \; x_e^* \leq \sum_{e \in E} d_e(x_e^*) \; x_e = \sum_{e \in E} (a_e x_e^* + b_e) x_e,$$

where the inequality follows from Proposition 2. Now, note that for all scalars x_e and x_e^* we have

$$0 \leq (x_e - \frac{x_e^*}{2})^2 = x_e^2 + \frac{x_e^{*2}}{4} - x_e x_e^*.$$

This implies $a_e x_e x_e^* \leq a_e x_e^2 + a_e x_e^{*2}/4$ for all nonnegative a_e, and hence:

$$C(\boldsymbol{x^*}) \leq \sum_{e \in E} (a_e x_e + b_e) x_e + \frac{a_e x_e^{*2}}{4} \leq C(\boldsymbol{x}) + \frac{C(\boldsymbol{x^*})}{4}.$$

It follows that $\frac{3}{4} C(\boldsymbol{x^*}) \leq C(\boldsymbol{x})$ for all $\boldsymbol{x} \in X$. So in particular we have

$$\frac{C(\boldsymbol{x^*})}{C(\boldsymbol{x^o})} \leq \frac{4}{3}.$$

□

Remark Note that this bound can be attained by some instances. Consider the non-atomic variant of the Braess paradox seen at the beginning of last lecture. In the WE all drivers take the $A \to B$ edge, which yields a cost of 200 per driver, while in the social optimum the drivers split 50/50 on the upper and lower road, at a cost of 150 per driver. So the PoA of this instance is $200/150 = 4/3$.

By using similar techniques, we can actually prove a much stronger result. For every delay function $d(\cdot)$ (i.e., mapping $\mathbb{R}_+$ onto itself, continuous and nondecreasing), define

$$\beta_d(v) := \frac{1}{vd(v)} \max_{x \geq 0} \Big\{ x\big(d(v) - d(x)\big) \Big\},$$

with the convention $0/0 = 0$. For a family $\mathcal{D}$ of delay functions, define further

$$\beta(\mathcal{D}) := \sup_{d \in \mathcal{D}} \sup_{v \geq 0} \beta_d(v),$$

and note that $\beta(\mathcal{D}) \leq 1$.

Theorem 4. *Let $\mathcal{G} = (E, \mathcal{S}, \boldsymbol{w}, \boldsymbol{d})$ be a non-atomic congestion game in which all delay functions are from the family $\mathcal{D}$, and assume that $\beta(\mathcal{D}) < 1$. Then,*

$$PoA(\mathcal{G}) \leq \frac{1}{1 - \beta(\mathcal{D})}.$$

Proof. Let $(\boldsymbol{f^*}, \boldsymbol{x^*})$ be a WE and $(\boldsymbol{f^o}, \boldsymbol{x^o})$ be a social optimum of $\mathcal{G}$. Now, consider an arbitrary feasible flow $(\boldsymbol{f}, \boldsymbol{x}) \in T$. As in the previous proof we have

$$C(\boldsymbol{x^*}) = \sum_{e \in E} d_e(x_e^*)\ x_e^* \leq \sum_{e \in E} d_e(x_e^*)\ x_e.$$

By definition of $\beta_d(\cdot)$, we have $vd(v)\beta_d(v) \geq x\big(d(v) - d(x)\big)$ for all $x, v \geq 0$. Hence, we have

$$\forall e \in E,\ x_e d_e(x_e^*) \leq x_e^* d_e(x_e^*) \beta_{d_e}(x_e^*) + x_e d_e(x_e).$$

This implies

$$C(\boldsymbol{x^*}) \leq \sum_{e \in E} x_e^* d_e(x_e^*) \beta(\mathcal{D}) + x_e d_e(x_e) = \beta(\mathcal{D}) C(\boldsymbol{x^*}) + C(\boldsymbol{x}).$$

It follows that $\big(1 - \beta(\mathcal{D})\big) C(\boldsymbol{x^*}) \leq C(\boldsymbol{x})$ for all $\boldsymbol{x} \in X$. So in particular we have

$$\frac{C(\boldsymbol{x^*})}{C(\boldsymbol{x^o})} \leq \frac{1}{1 - \beta(\mathcal{D})}.$$

□

Remark We will see in Exercise 3 of Worksheet #7 that this bound can be attained for a very simple type of instances (*Pigou Network*).

Price of Anarchy of atomic games

The study of the PoA in atomic games is complicated by the combinatorial nature of the problem. Moreover, several Nash equilibriums can exist. Consider an atomic congestion game $\mathcal{G} = (E, \boldsymbol{S}, \boldsymbol{w}, \boldsymbol{d})$. Similarly as for non-atomic instances, the total cost of a strategy profile $\boldsymbol{P} = (P_1, \ldots, P_N) \in \boldsymbol{S}$ inducing a load vector $\boldsymbol{x}$ is:

$$C(\boldsymbol{x}) = \sum_{i=1}^{n} w_i \sum_{e \in P_i} d_e(x_e) = \sum_{e \in E} x_e d_e(x_e).$$

Denote by X the set of loads induced by a feasible profile $\boldsymbol{P} \in \boldsymbol{S}$, and by X_{eq} the set of load vectors induced by a Nash Equilibrium $\boldsymbol{P^*}$ of $\mathcal{G}$. By definition, the PoA of the game is:

$$PoA(\mathcal{G}) := \frac{\max_{\boldsymbol{x^*} \in X_{eq}} C(\boldsymbol{x^*})}{\min_{\boldsymbol{x} \in X} C(\boldsymbol{x})}.$$

In other words, $PoA(\mathcal{G})$ is the smallest positive constant α such that $C(\boldsymbol{x^*}) \leq \alpha C(\boldsymbol{x})$ for all $(\boldsymbol{x}, \boldsymbol{x^*}) \in X \times X_{eq}$.

We have a positive result for instances with affine delay functions, but it was shown that the PoA can grow as $p^{p/2}$ for instances with polynomial delay functions of degree $\leq p$.

Theorem 5. *Let $\mathcal{G} = (E, \boldsymbol{S}, \boldsymbol{w}, \boldsymbol{d})$ be an atomic congestion game with affine (and nondecreasing) delay functions ($d_e(x) = a_e x + b_e$ for some $a_e, b_e \geq 0$). Then,*

$$PoA(\mathcal{G}) \leq \frac{3 + \sqrt{5}}{2} \simeq 2.618.$$

The proof of this result follows from the following lemma:

Lemma 6. *Let $\boldsymbol{x^*} \in X_{eq}$ and $\boldsymbol{x} \in X$. Then, $C(\boldsymbol{x^*}) \leq C(\boldsymbol{x}) + \sum_{e \in E} a_e x_e x_e^*$.*

Proof of the lemma. Denote by $\boldsymbol{P^*}$ the NE profile inducing $\boldsymbol{x^*}$ and by $\boldsymbol{P}$ the strategy profile inducing $\boldsymbol{x}$. It follows from the characterization of a Nash equilibrium that for all player i,

$$\sum_{e \in P_i^*} a_e x_e^* + b_e \leq \sum_{e \in P_i^* \cap P_i} a_e x_e^* + b_e + \sum_{e \in P_i \setminus P_i^*} a_e (x_e^* + w_i) + b_e \leq \sum_{e \in P_i} a_e (x_e^* + w_i) + b_e.$$

Now, multiplying by w_i and summing over i, we obtain

$$C(\boldsymbol{x^*}) = \sum_{i=1}^N w_i \sum_{e \in P_i} a_e (x_e^* + w_i) + b_e \leq \sum_{i=1}^N w_i \sum_{e \in P_i} a_e (x_e^* + x_e) + b_e,$$

where the inequality follows from the fact that, for an element $e \in P_i$, $x_e = \sum_{\{j : e \in P_j\}} w_j \geq w_i$. Now, we rearrange the order of the terms in the summation, and we obtain:

$$C(\boldsymbol{x^*}) \leq C(\boldsymbol{x}) + \sum_{i=1}^N w_i \sum_{e \in P_i} a_e x_e^* = C(\boldsymbol{x}) + \sum_{e \in E} a_e x_e^* \underbrace{\sum_{\{j : e \in P_j\}} w_i}_{x_e}.$$

□

We are now ready to prove the theorem:

Proof. By using the Cauchy Schwarz Inequality, we have

$$\sum_{e \in E} a_e x_e x_e^* = \sum_{e \in E} (\sqrt{a_e} x_e)(\sqrt{a_e} x_e^*) \leq \left(\sum_{e \in E} a_e x_e^2 \right)^{\frac{1}{2}} \left(\sum_{e \in E} a_e x_e^{*2} \right)^{\frac{1}{2}} \leq \sqrt{C(\boldsymbol{x}) C(\boldsymbol{x^*})}.$$

Combining this inequality with Lemma 6, we obtain:

$$C(\boldsymbol{x}^*) \leq C(\boldsymbol{x}) + \sqrt{C(\boldsymbol{x})C(\boldsymbol{x}^*)}$$
$$\iff \frac{C(\boldsymbol{x}^*)}{C(\boldsymbol{x})} \leq 1 + \sqrt{\frac{C(\boldsymbol{x}^*)}{C(\boldsymbol{x})}}.$$

Finally, set $U := \sqrt{\frac{C(\boldsymbol{x}^*)}{C(\boldsymbol{x})}}$. We have $U^2 \leq 1 + U$, which implies $U \leq \frac{1+\sqrt{5}}{2}$ and $U^2 \leq \frac{3+\sqrt{5}}{2}$.

□

Remark If the game $\mathcal{G}$ is unweighted (i.e., $w_i = 1$ for all i), then we can improve the upper bound on the price of anarchy: $PoA(\mathcal{G}) \leq \frac{5}{2}$. These bounds are tight, as will be seen in Exercise 1 of Worksheet #7 (i.e., the bounds are attained for some instances).

Remark It is questionable whether players might behave according to the *worst* Nash Equilibrium. If we are optimistic and think that the players rather converge to the *best* Nash Equilibrium, then we can use the notion of *price of stability* (cf. Exercises 2 and 6 of Worksheet #7):

$$PoS(\mathcal{G}) := \frac{\min_{\boldsymbol{x}^* \in X_{eq}} C(\boldsymbol{x}^*)}{\min_{\boldsymbol{x} \in X} C(\boldsymbol{x})}.$$