More Results on the Price of Anarchy

Instructor: Thomas Kesselheim

Today, we will wrap up our discussion of the inefficiency due to selfish behavior and show two further results. The first one is about convergence of best-response sequences in congestion games. We already know that all improvement sequences converge to pure Nash equilibria and then we can apply the price-of-anarchy bound to the equilibrium. Today we ask the question what happens on the way. In the second part of the lecture, we will consider a utility-maximization game and derive a price-of-anarchy bound there. In both cases, we use new techniques and we see known ones in a new light.

1 Guarantees for Best-Response Sequences

The price-of-anarchy results make statements about equilibria. For several games, one can also show guarantees for states that are not yet equilibria but players are still one their way there.

For example, in congestion games, pure Nash equilibria are hard to compute; probably there is no polynomial-time algorithm to find them. Improvement sequences always converge but sometimes all sequences are very long. However, for special kinds of improvement dynamics we will show that the PoA guarantee holds most of the time, long before convergence.

Consider the following *maximum-gain best-response dynamics*: While we are not in a pure Nash equilibrium, activate the player who can reduce her cost by the largest amount and let her switch to a best response. That is, from s we move to the state (s'_i, s_{-i}) by choosing i and s'_i such that $c_i(s) - c_i(s'_i, s_{-i})$ is maximized.

We start from an arbitrary state $s^{(0)}$ and get an improvement sequence $s^{(0)}, s^{(1)}, \ldots, s^{(T)}$, where T is a pure Nash equilibrium. Recall the Rosenthal potential of a state $\Phi(s) = \sum_{r \in \mathcal{R}} \sum_{k=1}^{n_r(s)} d_r(k)$. If the delay functions are non-decreasing and non-negative then $\Phi(s) \leq cost(s)$.

Theorem 10.1. *For any $0 < \gamma < 1$, the following statement holds. In a (λ, μ)-smooth congestion game with non-negative non-decreasing delay functions, for all but $\frac{n}{\gamma(1-\mu)} \ln \frac{\Phi(s^{(0)})}{\Phi(s^{(T)})}$ many states $s^{(t)}$ in the sequence generated by the maximum-gain best-response dynamics*

$$cost(s^{(t)}) \leq \frac{\lambda}{(1-\mu)(1-\gamma)} \min_{s \in S} cost(s) \ .$$

Let us understand the meaning of this theorem: The parameter γ can be arbitrary. For concreteness, think of it being set to $\frac{1}{2}$. Consider a game with positive integer delay functions. Then $\Phi(s^{(0)})$ is an upper bound on the length of the sequence that we generate. However, if the delays are large, then this might be a very large number. In comparison, $\ln \frac{\Phi(s^{(0)})}{\Phi(s^{(T)})}$ is only logarithmic in this quantity. In particular, it is polynomial in the size of the input if the delay functions are represented as binary numbers.

Proof. Let s^* be a state of minimum social cost. Let $\delta_i(s) = c_i(s) - c_i(s^*_i, s_{-i})$ be the amount by which player i can decrease her cost by unilaterally switching to s^*_i. If s is a pure Nash equilibrium, then $\delta_i(s) \leq 0$ for all i. Otherwise, it can have any value.

Let us call a state s *good* if $\sum_{i=1}^n \delta_i(s) \leq \gamma(1-\mu) cost(s)$, otherwise call it *bad*. We will upper-bound the number of bad states in the sequence and show that good states fulfill the claimed guarantee.

As, by definition, $c_i(s) = c_i(s^*_i, s_{-i}) + \delta_i(s)$, we can rewrite $cost(s)$ for any t as

$$cost(s) = \sum_{i=1}^n c_i(s) = \sum_{i=1}^n c_i(s^*_i, s_{-i}) + \sum_{i=1}^n \delta_i(s) \ .$$

By applying the smoothness inequality $\sum_{i=1}^n c_i(s_i^*, s_{-i}) \leq \lambda cost(s^*) + \mu cost(s)$, we get

$$cost(s) \leq \lambda cost(s^*) + \mu cost(s) + \sum_{i=1}^n \delta_i(s) \enspace .$$

If the state s is good, then this implies

$$cost(s) \leq \lambda cost(s^*) + \mu cost(s) + (1-\mu)\gamma cost(s) \enspace .$$

or equivalently

$$cost(s) \leq \frac{\lambda}{(1-\mu)(1-\gamma)} cost(s^*) \enspace ,$$

so the bound holds.

The definition of bad means that if $s^{(t)}$ is bad then the players can gain a lot by unilaterally switching to their strategy in s^*. We will now use the Rosenthal potential to bound the number of such steps.

Recall that the potential change matches the cost change of the switching player. So if from $s^{(t)}$ to $s^{(t+1)}$ player i changes her strategy, then

$$\Phi(s^{(t)}) - \Phi(s^{(t+1)}) = c_i(s^{(t)}) - c_i(s^{(t+1)}) \enspace .$$

We consider maximum-gain best-response improvements, which means that $c_i(s^{(t)}) - c_i(s^{(t+1)})$ is maximized among all choices of players i and deviation strategies. One particular choice would be to let the player i' with maximum $\delta_{i'}(s^{(t)})$ move to $s_{i'}^*$. So $c_i(s^{(t)}) - c_i(s^{(t+1)}) \geq \max_{i'} \delta_{i'}(s^{(t)})$. In combination, this gives us

$$\Phi(s^{(t)}) - \Phi(s^{(t+1)}) \geq \max_i \delta_i(s^{(t)}) \geq \frac{1}{n} \sum_{i=1}^n \delta_i(s^{(t)}) \enspace ,$$

where we also used that the maximum is no smaller than the average.

If $s^{(t)}$ is bad then $\frac{1}{n}\sum_{i=1}^n \delta_i(s^{(t)}) \geq \frac{\gamma(1-\mu)}{n} cost(s^{(t)})$. Furthermore, $cost(s^{(t)}) \geq \Phi(s^{(t)})$. Overall, this gives us for every bad $s^{(t)}$

$$\Phi(s^{(t)}) - \Phi(s^{(t+1)}) \geq \max_i \delta_i(s^{(t)}) \geq \frac{\gamma(1-\mu)}{n} \Phi(s^{(t)}) \enspace ,$$

and therefore

$$\Phi(s^{(t+1)}) \leq \left(1 - \frac{\gamma(1-\mu)}{n}\right) \Phi(s^{(t)}) \enspace .$$

If $s^{(t)}$ is good, then we simply use

$$\Phi(s^{(t+1)}) \leq \Phi(s^{(t)}) \enspace .$$

Let now B denote the number of bad states that we see throughout the sequence. It has to hold

$$\Phi(s^{(T)}) \leq \left(1 - \frac{\gamma(1-\mu)}{n}\right)^B \Phi(s^{(0)}) \leq \exp\left(-B\frac{\gamma(1-\mu)}{n}\right) \Phi(s^{(0)}) \enspace ,$$

where we used that $1 + x \leq \exp(x)$ for all $x \in \mathbb{R}$. Taking the logarithm on both sides, we get

$$\ln \Phi(s^{(T)}) \leq -B\frac{\gamma(1-\mu)}{n} + \ln \Phi(s^{(0)}) \enspace ,$$

or equivalently

$$B \leq \frac{n}{\gamma(1-\mu)} \ln \frac{\Phi(s^{(0)})}{\Phi(s^{(T)})} \enspace .$$

This proves the theorem. □

2 Price of Anarchy of Utility-Maximization Games

Most of our examples so far in this course were cost-minimization games. For the basic definitions there is no real difference when one turns to utility-maximization games instead. However, for the price of anarchy, the story is different, as we will seen in the following example.

Let us consider the following *market sharing game*. There are n firms, which are our players N, and m markets M. Each firm can decide to invest in one of these markets. Therefore, for player $i \in N$, the strategy set S_i is a subset of M.

Each market $j \in M$ has a total demand v_j. If k firms invest in the same market, then the market's demand is shared equally. So every firm gets a utility of $\frac{v_j}{k}$.

This way, the utility of player $i \in N$ in state $s \in S$ is

$$u_i(s) = \frac{v_{s_i}}{n_{s_i}(s)} \quad , \quad \text{where } n_j(s) = |\{i \in N \mid s_i = j\}| \enspace .$$

The *social welfare* of a state s is defined as the sum of player utilities, or equivalently, as the sum of demands that are fulfilled

$$SW(s) = \sum_{i \in N} u_i(s) = \sum_{\substack{j \in M \\ n_j(s) \geq 1}} v_j = \sum_{j \in \{s_1, \ldots, s_n\}} v_j \enspace .$$

Example 10.2. *There are n markets $1, \ldots, n$; each player can invest in every market. For some $\epsilon > 0$, the demands are $v_1 = n + \epsilon$, $v_2 = \ldots = v_n = 1$.*

The social welfare is maximized by each player investing in a different market. In this case, $SW(s) = 2n - 1 + \epsilon$. However, the only pure Nash equilibrium is that all players invest in market 1. Here, $SW(s) = n + \epsilon$.

So far, these games look a lot like the cost-sharing games from last time. And, indeed, they are. We can even interpret them as congestion games: Set $\mathcal{R} = M$, so the markets become the resources, and set $d_j(k) = -\frac{v_j}{k}$ for all $j \in M$ and all k. Now the players' cost functions in the congestion game are exactly the negative utility functions of the market sharing game: $c_i(s) = -u_i(s)$.

Observation 10.3. *Every Market Sharing Game has a pure Nash equilibrium.*

2.1 Price of Anarchy

Interestingly, despite the similarity to cost-sharing games, the price of anarchy is a lot different. Let us first define the price of anarchy for utility-maximization games. The definition is analogous to the one for cost-minimization game.

Definition 10.4. *Given a utility-maximization game, let* Eq *be a set of probability distributions over the set of states S. For some probability distribution p, let $SW(p) = \mathbf{E}_{s \sim p}[SW(p)] = \sum_{s \in S} p(s) SW(s)$ be the expected social welfare. The* price of anarchy for Eq *is defined as*

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

So, we swap minima and maxima and the social optimum is now in the numerator and the equilibrium in the denominator. This way, the PoA is still greater than 1. You will also find it defined as the reciprocal. In any case, values closer to 1 are better.

We can also adapt the smoothness definition as follows.

Definition 10.5. *A utility-maximization game is called* (λ, μ)-smooth *for $\lambda > 0$ and $\mu \geq 0$ if, for every pair of states $s, s^* \in S$, we have*

$$\sum_{i \in \mathcal{N}} u_i(s^*_i, s_{-i}) \geq \lambda \cdot SW(s^*) - \mu \cdot SW(s) \enspace .$$

There is again an analogous theorem that smoothness implies a bound on the price of anarchy.

Theorem 10.6. *In a* (λ, μ)*-smooth utility-maximization game, the PoA for coarse correlated equilibria is at most* $\frac{1+\mu}{\lambda}$*.*

Proof. For simplicity, we prove the theorem only for pure Nash equilibria. The generalization to coarse correlated equilibria works exactly as in the case of cost-minimization games.

Let s be a pure Nash equilibrium, s^* be a social optimum. Then we have

$$\begin{aligned} SW(s) &= \sum_{i \in N} u_i(s) \\ &\geq \sum_{i \in N} u_i(s_i^*, s_{-i}) \\ &\geq \lambda \cdot SW(s^*) - \mu \cdot SW(s) \ . \end{aligned}$$

So $(1 + \mu) SW(s) \geq \lambda SW(s^*)$. □

So, it only remains to give a smoothness proof.

Theorem 10.7. *The market-sharing game is* $(1, 1)$*-smooth. So* $PoA_{\mathsf{CCE}} \leq 2$*.*

Proof. Observe that

$$\begin{aligned} u_i(s_i^*, s_{-i}) &= \frac{v_{s_i^*}}{n_{s_i^*}(s_i^*, s_{-i})} \\ &\geq \begin{cases} v_{s_i^*} & \text{if } s_i^* \notin \{s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n\} \\ 0 & \text{otherwise} \end{cases} \\ &\geq \begin{cases} v_{s_i^*} & \text{if } s_i^* \notin \{s_1, \ldots, s_n, s_1^*, \ldots, s_{i-1}^*\} \\ 0 & \text{otherwise} \end{cases} \ . \end{aligned} \tag{1}$$

Denote by $T = \{s_1, \ldots, s_n\}$ all markets invested in in s and by $T^* = \{s_1^*, \ldots, s_n^*\}$ all markets invested in in s^*. We can now write

$$\sum_{i \in N} u_i(s_i^*, s_{-i}) \geq \sum_{j \in T^* \setminus T} v_j = \sum_{j \in T^* \cup T} v_j - \sum_{j \in T} v_j \ .$$

This is because if we take the sum over the terms (1) then every element of T^* exactly appears once unless it is in T.

This gives us

$$\sum_{i \in N} u_i(s_i^*, s_{-i}) \geq \sum_{j \in T^* \cup T} v_j - \sum_{j \in T} v_j \geq \sum_{j \in T^*} v_j - \sum_{j \in T} v_j = SW(s^*) - SW(s) \ .$$

This is exactly the requirement for smoothness. □

Further Reading

- Tim Roughgarden's lecture notes `http://theory.stanford.edu/~tim/f13/l/l16.pdf` and lecture video `https://youtu.be/QOCmIDzOcJA?t=1h1m33s` (There seems to be a flaw in this version of the proof.)
- B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrkoni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games, EC 2008 (Convergence result)
- A. Vetta, Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions, FOCS 2002 (Generalization of the result for market sharing game)