Algorithmic Game Theory	Summer 2015, Week 5

Price of Stability and Strong Nash Equilibria in Congestion Games

ETH Zürich	Paul Dütting

The price of anarchy concept provides a worst-case guarantee. It is a compelling concept when the price of anarchy is small. In this case we know that no matter which equilibrium the players will reach, the performance will be close to optimal. On the other hand, when it is large, it tells us that equilibria can be very inefficient.

In this lecture we will discuss a class of games with particularly bad worst-case equilibria and two orthogonal techniques to make a more refined statement about its strategic outcomes. The first takes a best-case perspective; the second sticks with the worst-case perspective but strengthens the equilibrium concept.

1 Motivating Example

We will consider *fair cost sharing games*, which are congestion games with delays $d_r(x) = c_r/x$ for constant c_r for every resource $r \in \mathcal{R}$. That is, we have a set $\mathcal{N}$ of n players and a set $\mathcal{R}$ of m resources. Player i allocates some resources, i.e., his strategy set is $\Sigma_i \subseteq 2^{\mathcal{R}}$. Each resource $r \in \mathcal{R}$ has fixed cost $c_r \geq 0$. The cost c_r is assigned in equal shares to the players allocating r (if any). That is, for strategy profile S denote by $n_r(S)$ the number of players using resource r. Then the cost $c_i(S)$ of player i is $c_i(S) = \sum_{r \in S_i} d_r(n_r(S)) = \sum_{r \in S_i} c_i/n_r(S)$

As before we will evaluate equilibria by means of the *social cost*, the sum over all players' costs, $cost(S) = \sum_{i \in \mathcal{N}} c_i(S)$, which can be rewritten as the sum of costs of resources allocated by at least one player:

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

The price of anarchy for pure Nash equilibria can be as big as the number of players n, even in a symmetric game.

Example 5.1 (Lower bound on price of anarchy)**.** *For $\epsilon > 0$, consider the following network cost sharing game, in which edge labels indicate the cost c_r of this resource:*

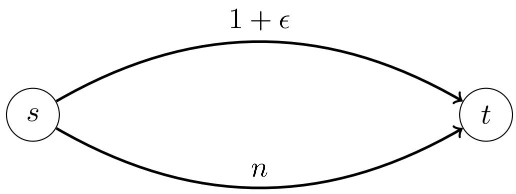

It is a pure Nash equilibrium if all players use the bottom edge, whereas the social optimum would be that all users use the top edge.

Although this is a very stylized example, there are indeed examples of such bad equilibria occurring in reality. A prime example are mediocre technologies, which win against better ones just because they are in the market early and get their share. As a concrete example consider social networks or messaging services. Even if you were to design and code-up a revolutionary new social network or messaging service, would people actually switch over from, say, Facebook or Whatsapp, if all their friends are using it?

On the other hand, there are several reasons why one could deem the above equilibrium as unrealistic. We will explore two of them.

2 Price of Stability

Let us first turn to the *price of stability*, which compares the best performance at equilibrium to the social optimum. Comparing the price of anarchy to the price of stability, allows us to gain insights on whether our assumption that players play some equilibrium causes the inefficiency or whether it is primarily a question of equilibrium selection.

Recall that for an equilibrium concept Eq, we defined the price of stability as follows

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

Observation 5.2. *In a symmetric cost sharing game, every social optimum is a pure Nash equilibrium. Therefore, the price of stability for pure Nash equilibria is* 1*.*

For general, asymmetric games, the social optimum is not necessarily a pure Nash equilibrium.

Example 5.3 (Lower bound on price of stability)**.** *Consider the following game with n players. Each player i has source node s_i and destination node t.*

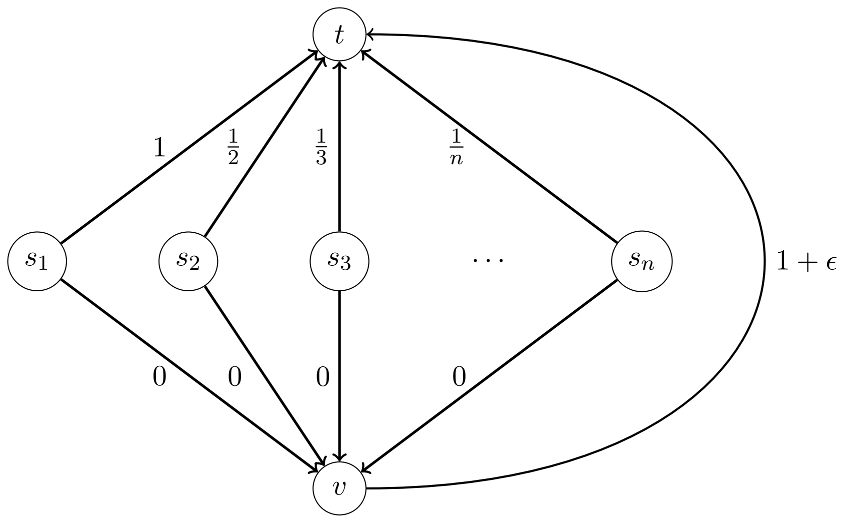

A player has two possible strategies: Either take the direct edge or take the detour via v. The social optimum lets all players choose the indirect path, which leads to a social cost of $1+\epsilon$. This, however, is not a Nash equilibrium. Player n, who currently faces a cost of $(1+\epsilon)/n$, would opt out and take the direct edge, which would give him cost $1/n$.

Therefore, the only pure Nash equilibrium lets all players choose their direct edge, yielding a social cost of $\mathcal{H}_n$, where $\mathcal{H}_n = \sum_{i=1}^{n} \frac{1}{i} = 1 + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{n}$ is the n-th harmonic number. We have $\mathcal{H}_n = \Theta(\log n)$.

Theorem 5.4. *The price of stability for pure Nash equilibria in fair cost sharing games is at most $\mathcal{H}_n$.*

Proof. Let's first derive upper and lower bounds on Rosenthal's potential function that apply to any state S, whether at equilibrium or not.

To obtain an upper bound we can use the definition of the potential function and the fact that each resource is used by no more than n players:

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

To obtain a lower bound we only need to observe that in $\Phi(S)$ we account for each player allocating resource r a contribution of c_r/i for some $i = 1, \dots, n_r$, whereas in his cost $c_i(S)$ we account only c_r/n_r. So,

$$cost(S) \leq \Phi(S) \ .$$

Now suppose we start at the optimum state S^* and iteratively perform improvement steps for single players. This eventually leads to a pure Nash equilibrium. Every such move decreases the potential function. For the resulting Nash equilibrium S' we thus have $\Phi(S') \leq \Phi(S^*)$ and

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

This proves that there is a pure Nash equilibrium S' that is only a factor of $\mathcal{H}_n$ more costly than S^*. □

We conclude that in both symmetric and asymmetric cost sharing games our assumption that players play an equilibrium has a rather mild impact, while equilibrium selection is a major source of inefficiency.

3 Strong Nash Equilibria

A second feature or our analysis up to this point that becomes salient in the introductory example is that so far we only considered equilibrium concepts which preclude beneficial deviations by individual and not, say, groups of players.

Specifically, the "bad" equilibrium that we have identified seems to be brittle in the sense that we only need to convince one additional player to switch from the expensive edge to the cheap edge to make this a beneficial deviation.[1]

We will call states in which no group of players can jointly deviate and achieve a strictly better outcome a *strong Nash equilibrium.*

Definition 5.5. *Let s be a state of a cost-minimization game. Consider a subset of players $A \subseteq \mathcal{N}$ (coalition). The strategy vector s' is a* beneficial deviation *for A if*

$$\begin{aligned} & c_i(s'_A, s_{-A}) \leq c_i(s) && \text{for all } i \in A \\ \text{and} \quad & c_i(s'_A, s_{-A}) < c_i(s) && \text{for at least one } i \in A \ . \end{aligned}$$

The state s is called a strong Nash equilibrium *if there is no coalition with a beneficial deviation.*

The concept of a strong Nash equilibrium is a *refinement* of the concept of a Nash equilibrium, in the sense that every strong Nash equilibrium is also a pure Nash equilibrium. Unilateral deviations correspond to coalitions of size one.

We will now consider the price of anarchy for strong Nash equilibria. Remember that generally, for an equilibrium concept Eq, it is defined as

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

So, given the respective equilibria exist, we have

$$PoS_{\mathsf{PNE}} \leq PoS_{\mathsf{SNE}} \leq PoA_{\mathsf{SNE}} \leq PoA_{\mathsf{PNE}} \ .$$

In other words, what we seek to do next is to limit the impact of equilibrium (mis-)selection by refining the equilibrium concept.

One caveat that one needs to be aware of when refining an equilibrium concept, is that one may jeopardize equilibrium existence.

[1] This is assuming that $\epsilon < 1$, so that the cost after the deviation is $(1+\epsilon)/2 < 1$. The assumption $\epsilon < 1$ is not crucial; for $\epsilon > 1$ we only need to convince more players.

Observation 5.6. *A state s of a symmetric cost sharing game is a strong Nash equilibrium if and only if it is socially optimal. Therefore, in symmetric cost sharing games the price of anarchy for strong Nash equilibria is* 1*.*

Asymmetric cost-sharing games, in contrast, need not possess a strong Nash equilibrium as the following example shows:

Example 5.7 (Non-existence of strong Nash equilibria)**.** *Consider the following two-player game, in which player 1 wants to allocate a path from s_1 to t_1 and player 2 wants to allocate a path from s_2 to t_2:*

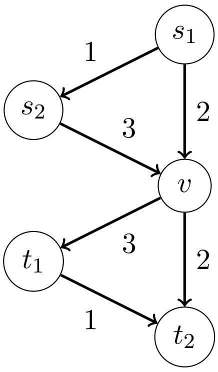

The two players have two strategies each, go left (L) or go right (R). The profile (L,L) is not a Nash equilibrium, as both players would want to deviate to R. Similarly, the profiles (L,R) and (R,L) are not a Nash equilibrium as the player playing L would rather play R. Finally, the profile (R,R) is not a strong Nash equilibrium, as the players would prefer forming a coalition and group-deviate to L.

On the other hand if strong Nash equilibria exist, then we can provide tighter bounds on the inefficiency of equilibria.

Theorem 5.8. *The price of anarchy for strong Nash equilibria in fair cost sharing games is at most $\mathcal{H}_n$.*

Proof. Let S be a strong Nash equilibrium, S^* be a socially optimal state.

First, we consider the coalition that consists of all players. Letting all players deviate to S^* is not beneficial. Therefore, there has to be one player i for which $c_i(S) \leq c_i(S^*)$. Without loss of generality, let this player be n.

Next, we consider the coalition of all players except n. Again, it is not beneficial if these players deviate to S^*. So, again, there has to be a player i for which $c_i(S) \leq c_i(S^*_{-n}, S_n)$. Let this player be $n-1$.

Following the argument, after renaming players, we get strategy profiles S^t for $t \in \{1, 2, \ldots, n\}$ such that

$$S_i^t = \begin{cases} S_i^* & \text{for } i \leq t \\ S_i & \text{for } i > t \end{cases}$$

and $c_t(S) \leq c_t(S^t)$. Note that $S_t^t = S_t^*$.

For $r \in \mathcal{R}$, define $k_r^t = |\{i \leq t \mid r \in S_i^*\}|$. We now have

$$c_t(S^t) = \sum_{r \in S_t^t} \frac{c_r}{n_r(S^t)} \leq \sum_{r \in S_t^t} \frac{c_r}{k_r^t} \quad .$$

This gives us

$$\sum_{i \in \mathcal{N}} c_i(S) \leq \sum_{t=1}^{n} c_t(S^t) \leq \sum_{t=1}^{n} \sum_{r \in S_t^t} \frac{c_r}{k_r^t} = \sum_{r \in \mathcal{R}} c_r \sum_{t : r \in S_t^*} \frac{1}{k_r^t} \quad .$$

Now observe that

$$\sum_{t:r\in S_t^*} \frac{1}{k_r^t} = 1 + \frac{1}{2} + \frac{1}{3} + \ldots + \frac{1}{n_r(S^*)} \enspace .$$

Therefore,

$$\sum_{r\in\mathcal{R}} c_r \sum_{t:r\in S_t^*} \frac{1}{k_r^t} = \Phi(S^*) \enspace .$$

Consequently,

$$cost(S) = \sum_{i\in\mathcal{N}} c_i(S) \leq \Phi(S^*) \leq \mathcal{H}_n \cdot cost(S^*) \enspace . \qquad \square$$

To see that this bound is tight consider the example from Section 2, which showed that the price of stability with respect to pure Nash equilibria is at least $\mathcal{H}_n$. The only pure Nash equilibrium in this example is also a strong Nash equilibrium.

Recommended Literature

- Chapter 19.3 in the AGT book. (Price of stability bound)
- Tim Roughgarden's lecture notes `http://theory.stanford.edu/~tim/f13/l/l15.pdf` and lecture video `https://youtu.be/VjCKNl-GENI`
- A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. Games and Economic Behavior, 67(1):5168, 2009. (Bound for strong Nash equilibria)