CS 6840: Algorithmic Game Theory — Spring 2017

Lecture 10: February 15

Lecturer: *Éva Tardos* — Scribe: *Jake Chen*

10.1 Best-case Nash Equilibria and Price of Stability

In the first part of the course we have primarily focused on upper-bounding the cost of Nash equilibria by analyzing the price of anarchy. For this purpose we have been assuming worst-case Nash equilibria. We will now take a different perspective and focus on best-case Nash equilibria. The price of stability is related to best-case Nash equilibria in the same way that the price anarchy and worst-case Nash equilibria are related. The motivation for analyzing the best-case is that we would like to urge players to cooperate with each other to get as close to optimal as possible, but the resulting configuration must be a stable Nash equilibrium or else players will refuse to cooperate.

Definition 10.1 ***Price of Stability.***

$$\textit{Price of Stability} = \frac{\textit{cost of best Nash}}{\textit{cost of OPT}}$$

10.1.1 A Preliminary Example

There is a variation of a congestion game where the cost of an element decreases the more players use it. This is called a **cost-sharing game**. Consider a routing game of this type, defined as follows:

- We are given a directed graph $G = (V, E)$ where each edge $e \in E$ has cost c_e.
- Each player i would like to get from a starting vertex s_i to a terminal vertex t_i.
- A strategy for player i consists of a path p_i from s_i to t_i.
- The flow on an edge $f(e)$ is the number of players using edge e in their chosen paths.
- Given a flow on this graph, the cost incurred by player i is given by the following expression:

  $$c_i(f) = \sum_{e \in p_i} \frac{c_e}{f(e)}$$

  Notice that since the summation is over edges in the player's path, the denominator is never zero.
- The total cost of a flow f is just the sum of $c_i(f)$ over all players and is equivalent to

  $$\sum_{e: f(e) > 0} c_e$$

Consider the following instance of a cost-sharing routing game with k players and some very small $\epsilon > 0$.

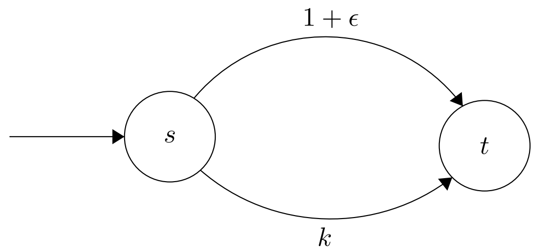

One Nash equilibrium (denoted by f) is where all k players take the bottom edge. Then each player i incurs cost

$$c_i(f) = 1$$

However, there is another significantly better Nash equilibrium (denoted by f^*) where all k players take the top edge. Then each player i incurs cost

$$c_i(f^*) = \frac{1+\epsilon}{k}$$

In fact, the equilibrium where all players take the top edge is socially optimal. This example shows that the best-case Nash equilibrium can potentially be significantly better than the worst-case Nash equilibrium. In general, we would hope that the best-case Nash equilibrium is pretty close to the social optimum. While we hope that players will converge to best-case Nash equilibria, it remains an open question whether under certain conditions, we can give an argument that players will naturally converge to them.

10.1.2 Analysis of the Cost-Sharing Routing Game

Since this cost-sharing routing game is a congestion game, there is also a corresponding potential game. We know that Nash equilibria correspond to local minima of the potential function Φ, but unfortunately, finding the global minimum of the Φ in general is NP-complete. We didn't prove this fact, but notice that finding the true minimum cost solution is the generalized Steiner tree problem, which is known to be NP-complete, it is maybe not surprising the the version minimizing the potential function is also NP-complete.

We now put some bounds on the cost of flows in this routing game. Recall that the cost of a flow f was defined to be

$$c(f) = \sum_{e: f(e)>0} c_e$$

and that the potential function of this game is given by

$$\Phi(f) = \sum_{e: f(e)>0} \sum_{x=1}^{f(e)} \frac{c_e}{x}$$

Rearranging the potential function by pulling the c_e out of the inner summation, we obtain

$$\Phi(f) = \sum_{e: f(e)>0} c_e \left(\sum_{x=1}^{f(e)} \frac{1}{x} \right)$$

This is clearly at least as big as $c(f)$, because the inner summation is at least 1. Furthermore, since there are k players, we have the following inequality:

$$\Phi(f) = \sum_{e: f(e)>0} c_e \left(\sum_{x=1}^{f(e)} \frac{1}{x} \right) \leq \sum_{e: f(e)>0} c_e \left(\sum_{x=1}^{k} \frac{1}{x} \right) = H_k c(f)$$

where we identify the inner summation as being the kth harmonic number, which asymptotically approaches $\ln k$. We arrive at the following:

Lemma 10.2 *For any flow f:*

$$c(f) \le \Phi(f) \le H_k c(f)$$

Proof: See above. ■

Theorem 10.3 *Let f be a Nash flow minimizing $\Phi(f)$ in this cost-sharing routing game and let f^* be a minimum-cost solution. Then the following expression holds:*

$$c(f) \le H_k c(f^*)$$

Proof: By the lemma, $c(f) \le \Phi(f)$. Since we chose f to minimize $\Phi(f)$, $\Phi(f) \le \Phi(f^*)$. Apply the lemma one more time so that we have $\Phi(f^*) \le H_k c(f^*)$. ■

As it turns out, this upper bound is tight. We now present an example demonstrating the case where the cost of a best-case Nash flow arbitrarily approaches H_k times the cost of the optimal.

Consider the following instance of the cost-sharing routing game.

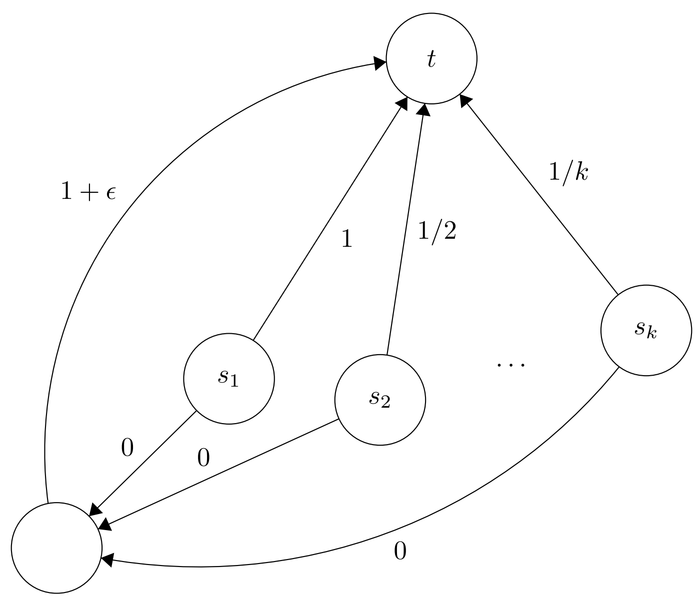

The only Nash equilibrium occurs when each player i uses their direct link to their destination t, incurring a total cost of H_k. The social optimum if every player takes the edge of 0 cost to the special node before taking the edge of cost $1+\epsilon$. Unfortunately, this is not Nash because if all k players are taking this sort of path, the kth player would prefer to deviate to her own personal path because a cost of $\frac{1}{k}$ is preferable to a cost of $\frac{1+\epsilon}{k}$. After this deviation, the $k-1$th player would prefer to deviate too because a cost of $\frac{1}{k-1}$ is preferable to a cost of $\frac{1+\epsilon}{k-1}$, and so on.

10.2 Strong Nash Equilibria

We have just seen an example with a Nash equilibrium where even if we allow players to collaborate, there is no group of players that can make a decision to change strategies that will make all players in the group better off. This is called a **strong Nash equilibrium**.

We will not go too much into the details of strong Nash equilibria, but there are two important takeaways (see book for more details):

1. If f is a Strong Nash equilibrium and f^* is a social optimum in the cost-sharing routing game then $c(f) \leq H_k c(f^*)$
2. Strong Nash equilibria are not guaranteed to exist. See the Braess paradox example for a unique Nash equilibrium that is not strong.