CS 6840 Algorithmic Game Theory February 3, 2020

Lecture 6: February 3

Instructor: Eva Tardos *Scribe: Kyra Wisniewski, Talia Turnham*

Last Time

Continuing with last Friday's discussion, we first recap with the discrete version of routing games. An instance of a discrete routing game has graph $G = (V, E)$, players i, source-sink s_i, t_i pairs, where each has a rate of exactly 1, i.e. $r_i = 1$. The delay on an edge e with x users is given by $d_e(x)$. Each player i aims to find a path P_i from s_i to t_i that minimizes their user cost. User cost is defined as $\sum_{e \in P_i} d_e(f(e))$, where the flow on an edge is given by $f(e) = \{\#i \mid e \in P_i\}$, i.e. the number of player who have edge e in their path. We can define social cost as $\sum_e f(e) d_e(f(e))$.

Last time, we proved the following Theorem:

Theorem 1. *This discrete version of a routing game is a potential game.*

Recall the following definition of a potential game:

Definition 1. A *potential game* is one for which there exists a *potential function* Φ with the property that, for every player who deviates from path P_i and switches to path Q_i, the change in the potential function value equals the change in the deviator's cost.

As a reminder, a potential function is defined as:

$$\Phi(P_1, \ldots, P_n) = \sum_e \sum_{k=1}^{f(e)} d_e(k)$$

where $P_1, \ldots, P_n$ are all paths that players take. If player i switches from path P_i to Q_i in a potential game, the change in the value of the potential function is given by:

$$\Phi(P_1, \ldots, P_i, \ldots, P_n) - \Phi(P_1, \ldots, Q_i, \ldots, P_n) = \sum_{e \in P_i} d_e(f(e)) - \sum_{e \in Q_i} d_e(\hat{f}(e))$$

where $\hat{f}(e)$ is the congestion if i switches to path Q_i.

Note that no assumptions on $d_e(x)$ are needed - it could be monotonically increasing, monotonically decreasing, or anything.

Decreasing Cost Functions

We will now shift our focus to decreasing cost functions, and examine the cost-sharing game. In a simple model of the cost-sharing game with decreasing cost, there is a cost c_e on every edge e. As before, player i is choosing a path and taking his share of each edge on his path's cost. So, we have that $d_e(x) = \frac{c_e}{x}$ for x users. Note that $d_e(x)$ is a decreasing cost function (the bigger x is, the smaller $d_e(x)$ is). This is also an example of Positive Externality (i.e. taking a certain edge benefits the other players taking that same edge).

Now, we will examine what we can prove about Nash equilibriums of games with decreasing cost functions. As an example, consider the following network, where $n > 0$ and the cost of each edge is given beside it:

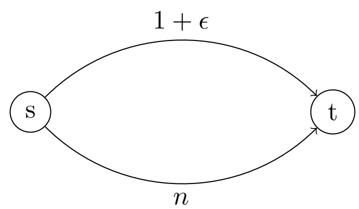

In this network, there are two pure strategy Nash solutions:

1. All n users take the top edge, so the total cost is $1+\epsilon$ and the user cost is $\frac{1+\epsilon}{n}$ (this is clearly the optimum solution)
2. All n users take the bottom edge, so the total cost is n and the user cost is 1

These results tell us that there is such a thing as a bad Nash solution and pushes us to consider how to evaluate the quality of the best possible Nash. To do so, we consider a measure called the Price of Stability defined as:

$$\text{Price of Stability} = \max_{games} \frac{mincost(NASH)}{OPT}$$

This gives rise to the following theorem, which we use to argue that the Nash that minimizes the potential function will get us close to the minimal cost solution:

Theorem 2. *In a cost-sharing game with c_e for $e \in E$, (s_i, t_i) pairs, and $d_e(x) = \frac{c_e}{x}$, we have that*

$$\text{Price of Stability} \le H_n \sim ln(n)$$

where H_n is the n-th harmonic number.

Proof. Consider paths $P_1, \ldots, P_n$ minimizing the potential function $\Phi(P_1, \ldots, P_n) = \sum_e \sum_{k=1}^{f(e)} d_e(k)$. If one player changes paths, the change in the potential function value is exactly the change in that player's cost. Since we know that the potential function in minimized by the original paths, we know that no player can minimize his own cost anymore. Thus, we know that the set of paths $P_1, \ldots, P_n$ is a Nash equilibrium.

Note that while our goal is to minimize the social cost, our Nash was found by minimizing the potential function. Usually, minimizing the wrong function gives the wrong minimum. However, if we can show that the potential function and the social cost give values that are not too far apart from each other, then we will know that our solution is not too far off from the optimum. We will first argue that $\Phi(f) \ge SC(f)$. We can write $SC(f)$ as $\sum_e \sum_{k=1}^{f(e)} d_e(f(e))$ and perform a term-by-term comparison of $\Phi(f) = \sum_e \sum_{k=1}^{f(e)} d_e(k)$ to $SC(f) = \sum_e \sum_{k=1}^{f(e)} d_e(f(e))$. We see that $d_e(k) \ge d_e(f(e))$ for every term, since $d_e(x)$ is decreasing. Thus, the sum of the $d_e(k)$ must be $\ge$ to the sum of $d_e(f(e))$ terms, and so we get that $\Phi(f) \geq SC(f)$ as desired.

Now, we want to look at how much bigger $\Phi(f)$ is than $SC(f)$. We claim that $\Phi(f) \leq SC(f) \cdot H_n$, where H_n is the n-th harmonic number, namely $1 + \frac{1}{2} + \cdots + \frac{1}{n}$. To show why this is true, we will look at any particular edge e and perform an edge-by-edge comparison of $\Phi(f) = \sum_e \sum_{k=1}^{f(e)} d_e(k)$ to $SC(f) = \sum_e \sum_{k=1}^{f(e)} d_e(f(e))$. Looking at $\Phi(f)$, for edge e we have that:

$$\begin{aligned}\sum_{k=1}^{f(e)} d_e(k) &= c_e + \frac{c_e}{2} + \cdots + \frac{c_e}{f(e)} \\ &= c_e\left(1 + \frac{1}{2} + \cdots + \frac{1}{f(e)}\right) \\ &\leq c_e \cdot H_n\end{aligned}$$

Looking at $SC(f)$, for edge e we have that:

$$\begin{aligned}\sum_{k=1}^{f(e)} d_e(f(e)) &= \frac{c_e}{f(e)} + \frac{c_e}{f(e)} + \cdots + \frac{c_e}{f(e)} \\ &= f(e)\left(\frac{c_e}{f(e)}\right) \\ &= c_e\end{aligned}$$

So, for any edge e, we see that the cost of e in $\Phi(f)$ is within a factor of H_n to the cost of e in $SC(f)$. Thus, summing over all edges, we get that $\Phi(f) \leq SC(f) \cdot H_n$.

Now, we have shown that $\Phi(f) \geq SC(f)$ and $\Phi(f) \leq SC(f) \cdot H_n$, which gives rise to a final lemma to complete the proof:

Lemma 1. *$SC(f) \leq \Phi(f) \leq SC(f) \cdot H_n$ implies that the Nash minimizing Φ is within an factor of H_n to the optimal solution.*

Proof. Letting f = Nash solution, and f^* = optimal solution minimizing the social cost, we see that:

$$SC(f) \leq \Phi(f) \leq \Phi(f^*) \leq SC(f^*) \cdot H_n$$

Thus, we get that $\Phi(f) \leq (f^*) \cdot H_n$, proving the lemma. ∎

So, since we have shown that $SC(f) \leq \Phi(f) \leq SC(f) \cdot H_n$, we know that the Nash minimizing Φ is within an factor of H_n to the optimal solution by Lemma 1. Price of Stability is defined as $\max_{games} \frac{mincost(NASH)}{OPT}$, and so we can take the following steps to reach the result of Theorem 2:

$$\text{Price of Stability} = \max_{games} \frac{mincost(NASH)}{OPT} = \frac{\Phi(f)}{SC(f^*)} \leq \frac{SC(f^*) \cdot H_n}{SC(f^*)} = H_n$$

$$\implies \text{Price of Stability} \leq H_n$$

This concludes the proof of Theorem 2. ∎

The general message here is that if we optimize a slightly wrong equation, we get a slightly wrong result. The result of Theorem 2 gives a bound for just how "slightly wrong" our result is.

Now, we ask: Is it possible to prove a tighter bound? This bound is bad, especially if there are many users involved. Unfortunately, the answer seems to be no - this bound is tight. To illustrate why, we consider the following example:

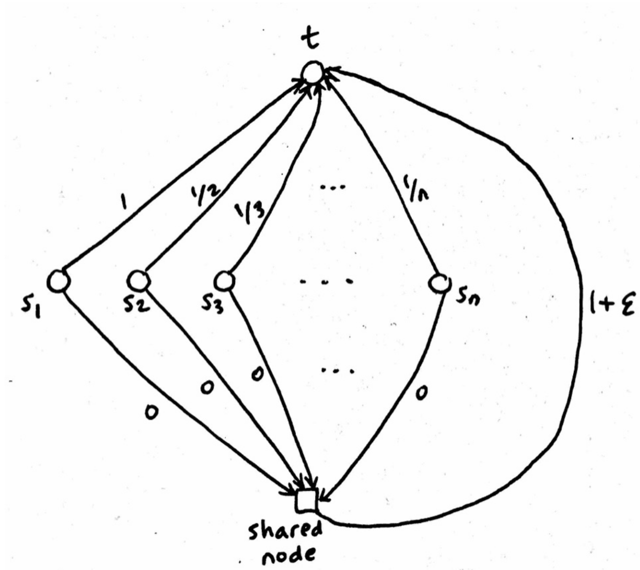

In this network, we have n sources $s_1, \ldots, s_n$ and a single sink t. Each player has two options: (1) use the edges with cost 0 to go down to the shared node and used the shared path with cost $1+\epsilon$ to get to t, or (2) use the above edges to go directly up to t. In this case, the optimal solution is for everyone to travel down to the shared node and through the shared path to t, giving a total cost of $1+\epsilon$. This network has a unique Nash of total cost H_n, where everyone goes directly up to t. Note that if one person doesn't take the path going up, his cost will be higher ($\frac{1}{k}$ if he goes up vs. $\frac{1+\epsilon}{k}$ if he goes down). Because of this example, we see that H_n is the tightest bound on the Price of Stability that we can achieve.