CS 6840: Algorithmic Game Theory **Spring 2017**

Lecture 4: February 1

Lecturer: Éva Tardos *Scribe: Cong Ding*

4 Over-Provisioning in routing

4.1 Logistics

- Problem Set 1 will be released through CMS by this evening, which dues in 2 weeks.
- It may help you do the homeworks if you find a partner (Piazza can be used for this purpose).
- There are 5 problems in total, and we have to solve 4 of them to get full credit.
- Extra credit will be available for solving all the problems (essentially) correctly, but multiple partial credits don't make up problems not fully solved. Extra credit is useful for all, and is needed for reaching a grade of A+.
- If we get stuck, we can submit partial solution to reflect our thought.
- Office hours and tutorials: Thursday morning, Monday afternoon, Tuesday.

4.2 Questions from previous class

4.2.1 Why should I care about price of anarchy?

In the example from previous class, we get the POA = 4/3 when $N = 1$; and POA$\to +\infty$ when $N \gg 1$ (very large).

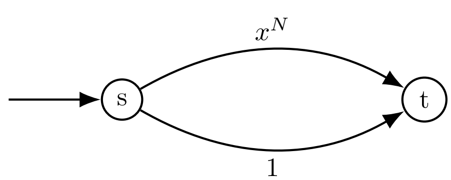

If we are talking about the downloading speed in the Internet, a factor of 4/3 may not be an issue (in part as our model of download speeds does not exactly match reality in any case), but PoA $\to \infty$ is always an issue.

If we are talking about money, for example 1 million dollars, a factor of 4/3 is also an issue.

If we are talking about database processing speed, a factor of 4/3 may or may not be an issue, but a factor of 6 is obviously an issue.

In summary, small POA may not be an issue, but large POA is an issue. The definition of small or large is somehow subjective. However, it really depends on applications.

4.3 Over-Provisioning in routing

In last lecture, we discussed the scenario that all the links can carry unlimited traffic. However, this is not true in real-world routing problems. In communication networks, all the links have capacities. We call these links capacitied links. In this class, we consider networks in which every cost function $c_e(x)$ has the form

$$c_e(x) = \begin{cases} \frac{1}{u_e - x} & \text{if } x < u_e; \\ +\infty & \text{if } x \geq u_e. \end{cases}$$

The parameter u_e represents the capacity of edge e. The following figure shows the intuition of the formula ($u_e = 1$).

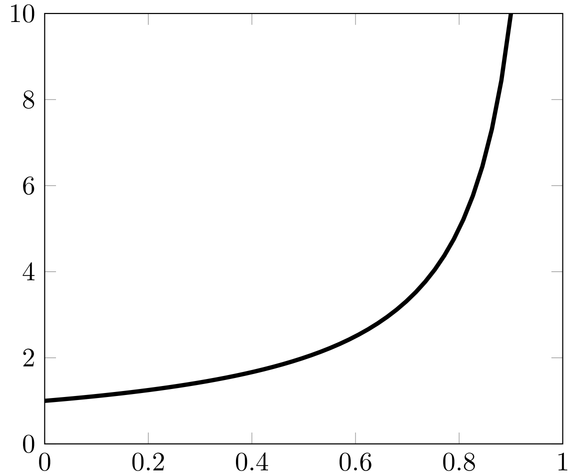

This function stays very flat until the amount of traffic nears the capacity, at which point the cost rapidly tends to $+\infty$. This is the simplest cost function used to model delays in communication networks.

We get the Nash of the beginning example, in the case that the traffic rate $r < u_e$.

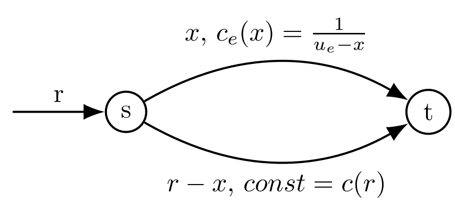

Theorem 4.1 *The POA of network with cost $c_e(x) = \frac{1}{u_e - x}$ is at most the worst POA in above example.*

Suppose $r = (1 - \beta)u_e$, we get

$$POA = \frac{1}{2}\left(1 + \frac{1}{\sqrt{\beta}}\right)$$

From the formula we see that the maximum POA is $+\infty$ when $r \leq u$ and $\beta \to 0$. However, it is not so bad, when $\beta \gg 0$. This observation extends to networks also, as shown by the following theorem.

Theorem 4.2 *If Nash flow f, $f(e) \leq (1-\beta)u_e$ for all edges, the network with cost $c_e(x) = \frac{1}{u_e - x}$, then POA of this flow $\leq \frac{1}{2}\left(1 + \frac{1}{\sqrt{\beta}}\right)$.*

For example, when $\beta = \frac{1}{2}$, POA $\leq \frac{1}{2}\left(1 + \frac{1}{\sqrt{2}}\right) = 1.21$; when $\beta = \frac{1}{4}$, POA $\leq \frac{1}{2}\left(1 + \frac{1}{\sqrt{4}}\right) = 1.50$; when $\beta = \frac{1}{10}$, POA $\leq \frac{1}{2}\left(1 + \frac{1}{\sqrt{10}}\right) = 2.08$. The POA is not large though a significant portion of the capacity is used.

Proof: One can prove this theorem along the lines we used last time, proving that the worst case price of anarchy always occurs on networks with two nodes and two links. Here we provide an alternate proof sketch that relies more on multivariate optimization, but makes the reason the theorem is true (maybe) more intuitive.

Suppose there is a Nash for $G = (V, E)$. We create a new network $G' = (V, E \cup E')$, where E' is the new copy of each edge. Therefore there are two links in graph G' for each link in the original graph G. To be detailed, for each edge e with cost $c_e(x)$ in G, we create two edges e and e' in graph G': e with cost $c_e(x)$ (same as the edge in G), and another edge e' with constant cost, s.t., $c_{e'}(x) = c_e(f(e))$.

In the new graph G', we have the following claims.

- Flow f is also a Nash in G', because the cost of f was not changed from graph G to G',a nd the new edges added do not offer cheaper alternate routes for any flow.
- The minimum cost of flow f' on $G' \leq$ the minimum cost of flow f on G, since new added edges can decrease the minimum cost or keep it unchanged.

The main claim is

Lemma 4.3 *The minimum cost flow f' on G' is an optimal way to rebalance flow $f(e)$ between e and e' for all edges $e \in E$.*

To prove this lemma, we need to use convex optimization (or at least multivariate calculus). We will not go into details here.

Using the lemma, it is not hard to see the theorem: Let f^* be this optimally rebalanced flow. Note that the rate of f^* on the pair of edges e and e' if $f(e) \leq (1-\beta)u_e$ by assumption. So for each edge pair e and e' the ratio of the cost of $f(e)$ and the sum of the costs of $f^*(e)$ and $f^*(e')$ is at most $\frac{1}{2}(1 + \frac{1}{\sqrt{\beta}})$. Now summing the costs and using the fact that $\frac{a+b}{c+d} \leq \max(\frac{a}{c}, \frac{b}{d})$ we get the theorem. ■

So far, we weer still evaluating networks and price of anarchy, rather than offering a suggestion of how to improve the price of anarchy. However, the theorem above suggest a natural act: increase the capacity of the edges, if the price of anarchy is bad. The next theorem makes this concrete:

Theorem 4.4 *For a flow on network capacity $c_e(x) = \frac{1}{u_e - x}$, the cost of optimal routing on $G \geq$ the cost of Nash flow with capacity $2u_e$ on every edge.*

We can prove this theorem using the following theorem for general capacities.

Theorem 4.5 *For a cost function, which is continuous, non-negative, monotonically non-decreasing, the minimum cost with traffic rate $2r_i \geq$ the cost of Nash at rate r_i for each edge.*

Proof: (Theorem 4.5)

See Section 12.3 in Roughgarden. ∎

Now we use this theorem to prove our claim about networks with increased capacity:

Proof: (Theorem 4.4)

Suppose f is the Nash of graph with capacity $2u_e$, then the cost of $f \leq$ the cost of f^* which carries $2r_i$ traffic rate with capacity $2u_e$ by Theorem 4.4.

We claim that $\frac{1}{2}f^*$ is the optimal flow with capacities u_e. To see why, notice that for any value x:

x flow with capacity u_e: cost is $\frac{1}{u_e - x}$, in total $x\frac{1}{u_e - x}$.

$2x$ flow with capacity $2u_e$: cost is $\frac{1}{2u_e - 2x}$, in total $2x\frac{1}{2u_e - 2x} = x\frac{1}{u_e - x}$.

The two total costs are the same, so any flow with values x and capacities u_e is in one-to-one correspondence with a flow of values $2x$ and capacities $2u_e$.

Therefore, $\frac{1}{2}f^*$ is optimal flow for capacity u_e, and it has the came cost is f^* for capacities $2u_e$. ∎