CS 6840 Algorithmic Game Theory October 2nd, 2024

Lecture 16: Price of Anarchy in Routing Games, Cont'd

Instructor: Eva Tardos *Scribe: Zhi Liu*

1 Problem setup: non-atomic network flow

Consider a network consisting of multiple origin-destination pairs (s_i, t_i), each carrying a flow of r_i.

On edge e, each unit of flow incurs a cost of $C_e(x)$, where $x \geq 0$ is the current total flow on this edge. We assume that $C_e(x) \geq 0$, is monotonically increasing, and is continuous.

Let $f_p \geq 0$ denote the amount of flow on path p from some origin s_i to destination t_i, then the total flow on all available paths should sum up to r_i:

$$\sum_{p: s_i \to t_i} f_p = r_i,$$

which results in total flows on each edge e:

$$f(e) = \sum_{p: e \in p} f_p.$$

Using these total flow amounts, we can calculate the cost incurred for a unit of flow on path p when the overall flow pattern is f:

$$C_p(f) = \sum_{e \in p} C_e(f(e)),$$

and subsequently, calculate the total cost of flow pattern f:

$$\begin{aligned}
\text{cost}(f) &= \sum_p f_p C_p(f) \\
&= \sum_p f_p \sum_{e \in p} C_e(f(e)) && \cdots \text{definition of } C_p(f) \\
&= \sum_e C_e(f(e)) \sum_{p: e \in p} f_p && \cdots \text{exchange summation} \\
&= \sum_e c_e(f(e)) f(e)
\end{aligned}$$

Remark: as the number of edges and vertices grows, this summation over p might not be realistic, but nevertheless it's useful to think of the cost this way.

2 Price of anarchy

2.1 Characteristics of equilibrium flow

We discussed last time that if f is an equilibrium flow, then for all path p from s_i to t_i with positive flow $f_p > 0$, the unit cost on this path must be equal to the minimum unit cost for any path from s_i to

t_i:

$$C_p(f) = \min_{q: s_i \to t_i} C_q(f)$$

Remark: this is a result of the edge costs being continuous; if there happens to be one path that does not satisfy this, an infinitesimal amount of flow can deviate and get lower cost on another path.

Consider f to be an equilibrium flow pattern, and consider a copy of this network with edge costs $\bar{C}_e = C_e(f(e))$, i.e., freeze these equilibrium edge costs. Then we claim that for this set of edge costs, the equilibrium flow f is the minimum cost flow. This can be seen from the previous argument, where only minimum cost paths have positive flow.

2.2 Introducing the α coefficient

To facilitate later arguments, consider the following two-edge network with the following costs:

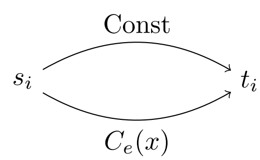

Definition 1. For a class of possible cost functions $\mathcal{C}$, let

$$\alpha(\mathcal{C}) = \max \frac{\text{Nash}}{\text{OPT}},$$

on the above network where $C_e \in \mathcal{C}$, where the maximum is taken over all possible C_e and all possible equilibria. Here Nash and OPT denote the social cost of any Nash equilibrium and the social optimum of this two-edge network, respectively.

2.3 Price of anarchy in routing games

We now introduce the following theorem.

Theorem 1. *If all edge costs $C_e \in \mathcal{C}$, then the price of anarchy in the network is bounded by $\alpha(\mathcal{C})$.*

Proof. Let f be the flow under any Nash, and let f^* be the minimum social cost flow. Consider the following two-edge network, for any edge e in the original network, with the flow from the left node to the right node being $f(e)$:

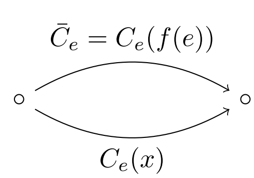

We claim that

$$\alpha \geq \frac{C_e(f(e))f(e)}{f^*(e)C_e(f^*(e)) + \bar{C}_e(f(e) - f^*(e))}. \tag{1}$$

To see this, first note that one Nash equilibrium of this two-edge network is if $f(e)$ unit flow through the lower edge, and 0 unit flow through the upper edge, since any flow in this scenario could not switch to the upper edge to gain a lower cost. In this Nash, the social cost is exactly $C_e(f(e))f(e)$.

On the other hand, for this two-edge network,

$$\text{OPT} \leq f^*(e)C_e(f^*(e)) + \bar{C}_e(f(e) - f^*(e)).$$

This is quite obviously true for $f^*(e) \leq f(e)$, as the way we calculate the OPT would be

$$\text{OPT} = \min_{x\in[0,f(e)]} xC_e(x) + \bar{C}_e(f(e) - x),$$

that is, minimizing the total cost of letting x unit flow through the lower edge and $f(e) - x$ through the upper edge. It is worth noting that this is also true for the case where $f^*(e) > f(e)$, if C_e is non-decreasing:

$$\text{OPT} \leq f(e)C_e(f(e)) \leq f^*(e)C_e(f^*(e)) - \bar{C}_e(f^*(e) - f(e)).$$

This is perhaps easier shown through a figure: the blue area is $f(e)C_e(f(e))$ and the yellow area plus the blue area is $f^*(e)C_e(f^*(e)) - C_e(f(e))(f^*(e) - f(e))$. The yellow area is non negative for increasing C_e and $f^*(e) \geq f(e)$.

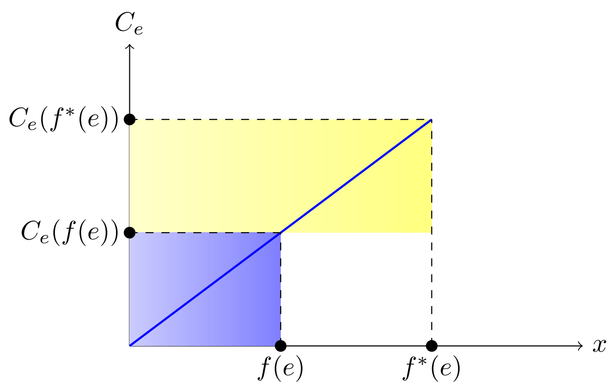

Now, combining the above arguments, we have confirmed Equation 1. Based on Equation 1, summing over all possible edges in the original network, we get

$$\begin{aligned}
\text{cost}(f) &= \sum_e C_e(f(e))f(e) \\
&\leq \alpha \left[\sum_e f^*(e)C_e(f^*(e)) + \sum_e C_e(f(e))(f(e) - f^*(e)) \right] \\
&= \alpha \left[\text{OPT} + \text{cost}(f) - \sum_e C_e(f(e))f^*(e) \right] \\
&\leq \alpha \left[\text{OPT} + \text{cost}(f) - \text{cost}(f) \right] \\
&= \alpha \text{OPT},
\end{aligned}$$

where the last inequality follows from our previous claim that the equilibrium flow f is the minimum cost flow on a network with fixed costs $C_e(f(e))$:

$$\sum_e C_e(f(e))f(e) \leq \sum_e C_e(f(e))f'(e), \quad \forall f' \text{ feasible.}$$

This completes the proof.

■

2.4 Intuition

The total cost of flow on a path p is $\sum_{e\in p} C_e(f(e))f(e)$. The derivative of this function captures the marginal gain in social cost from moving a bit of flow from or to this path. Particularly, if the derivative on path p is larger than that on path q, then we get global improvement from moving a bit of flow from path p to q, provided that they are both paths from s_i to t_i.

A closer look into the derivative of edge costs $C_e(f(e))f(e)$ tells us the composition of social cost:

$$[C_e(x)x]' = C_e(x) + xC'_e(x).$$

- The first part $C_e(x)$ is a selfish part, capturing the individual cost of this unit of flow;
- The second part $xC'_e(x)$ captures the 'social pain': how much other flow on this edge is impacted by the behavior of this unit of flow that we are about to move.

3 Next time

Consider adding an edge with fixed cost $C_e(f(e))$ alongside each edge e in the original network. Next time we will talk about three claims:

- Nash is unchanged in the new network;
- OPT can only improve in the new network;
- OPT on this new network can be found by optimizing flows edge by edge.