Stat 155 Lecture 16 Notes

Daniel Raban

March 20, 2018

1 The Price of Anarchy

1.1 Flows and latency in networks

Last time we saw Braess's paradox, in which a Nash equilibrium resulted in an inefficient flow in a network. How can we quantify the inefficiency of Nash equilibria in a network?

Definition 1.1. For a routing problem we define the *price of anarchy* as

$$\text{price of anarchy} = \frac{\text{average travel time in worst Nash equilibrium}}{\text{minimal average travel time}}.$$

Note that the minimum is over all flows. The flow minimizing average travel time is th socially optimal flow. The price of anarchy reflects how much average travel time can decrease in going from a Nash equilibium flow (where all individuals choose a path to minimize their travel time) to a prescribed flow.[1]

Example 1.1. Consider the following network.

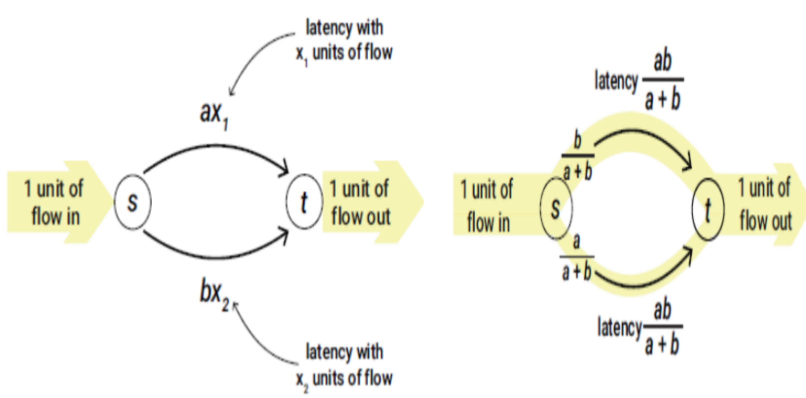

[1]This was first defined by Elias Koutsoupias and Christos Papadimitriou. They were awarded the 2012 Gödel Prize (with four others).

The price of anarchy of this network is 1. Finding the socially optimal strategy is equivalent to minimizing the function

$$f(x_1) = ax_1^2 + b(1 - x_1)^2.$$

Setting $f'(x_1) = 0$ is a equivalent to $ax_1 = b(1 - x_1)$, which is the Nash equilibrium condition.

Example 1.2. Consider the following network.

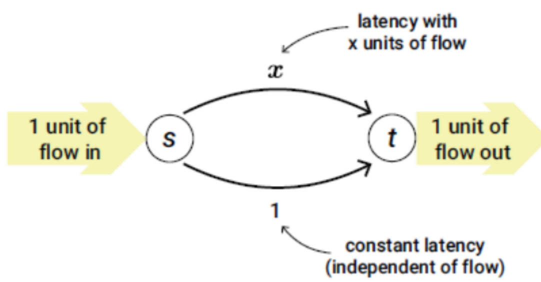

A Nash equilibrium flow occurs when $x = 1$. We can find an optimal flow by minimizing the function

$$f(x) = x^2 + (1 - x).$$

This is minimized at $x = 1/2$, so the socially optimal strategy gives an average time of $3/4$. So the price of anarchy is $4/3$.

Definition 1.2. A *flow* f from source s to destination t in a directed graph is a mixture of paths from s to t, with mixture weight f_P for path P. We write the flow on an edge e as

$$f_e = \sum_{P:e \in P} f_P.$$

Definition 1.3. *Latency* on an edge e is a non-decreasing function of F_e, written $\ell_e(F_e)$. The latency on a path P is the total latency

$$L_P(f) = \sum_{e \in P} \ell_e(F_e).$$

The average latency is

$$L(f) = \sum_P f_P L_P(f) = \sum_e F_e \ell_e(F_e).$$

Definition 1.4. A flow f is a *Nash equilibrium flow* if, for all P and P', if $f_P > 0$, then $L_P(f) \leq L_{P'}(f)$.

In equilibrium, each driver will choose some lowest latency path with respect to the current choices of other drivers.

1.2 The price of anarchy for linear and affine latencies

Theorem 1.1. *For a directed, acyclic graph (DAG) with latency functions ℓ_e that are continuous, non-decreasing, and non-negative, if there is a path from source to destination, there is a Nash equilibrium unit flow.*

Proof. Here is the idea of the proof. This is the non-atomic version of a congestion game. For the atomic version (finite number of players), we showed that there is a pure Nash equilibrium that can be found by descending a potential function. The same approach works here. The potential function is

$$\phi(f) = \sum_e \int_0^{F_e} \ell_e(x)\, dx.$$

If f is not a Nash equilibrium flow, then $\phi(f)$ is not minimal. ϕ is convex, on a convex, compact set, so it has a minimum. $\square$

Theorem 1.2. *For linear latencies, that is $\ell_e(x) = a_e x$ with $a_e \geq 0$, if f is a Nash equilibrium flow and f^* is a socially optimal flow (that is $L(f^*)$ is minimal, then*

$$L(f) = L(f^*).$$

Proof. Since f is a Nash equilibrium, there is no advantage to shifting any flow from f to any other flow. In particular, there is no advantage to shifting from f to f^*.

$$\begin{aligned}
L(f) &= \sum_{P: f_P > 0} f_P L_P(f) \\
&\leq \sum_P f_P^* L_P(f) \\
&= \sum_P f_P^* \sum_e \ell_e(F_e) \\
&= \sum_e \left(\sum_{P: e \in P} f_P^* \right) \ell_e(F_e) \\
&= \sum_e F_e^* \ell_e(F_e) \\
&= \sum_e a_e F_e^* F_e \\
&= \sum_e a_e \left(-(F_e - F_e^*)^2/2 + (F_e^{*2} + F_e^2)/2 \right) \qquad \text{(magic)} \\
&\leq \sum_e a_e (F_e^{*2} + F_e^2)/2
\end{aligned}$$

$$= \sum_e (F_e^* \ell_e(F_e^*) + F_e \ell_e(F_e))/2$$
$$= (L(f^*) + L(f))/2,$$

so $L(f) \le L(f^*)$. $\square$

Corollary 1.1. *For linear latency functions, the price of anarchy is 1.*

Remark 1.1. In the proof above, we used a quadratic inequality to bound $F_e^* F_e$; one could also use the Cauchy-Schwarz inequality to do the same. Quadratic inequalities are useful because for any α, we have

$$xy = -\left(\alpha x - \frac{y}{2\alpha}\right)^2 + \alpha^2 x^2 + \frac{y^2}{4\alpha^2} \le \alpha^2 x^2 + \frac{y^2}{4\alpha^2}.$$

This shows that

$$xy = \min_\alpha \left(\alpha^2 x^2 + \frac{1}{4\alpha} y^2\right).$$

If x and y have the same sign, then we could choose $\alpha^2 = y/(2x)$ to give $xy = \alpha^2 x^2 + y^2/(4\alpha^2)$, so in this case, these inequalities are tight. In bounding the price of anarchy, we could use any of these inequalities to gie a linear bound relating to $L(f)$ to $L(f^*)$. The choice of $\alpha^2 = 1/2$ givse the best linear bound.

Theorem 1.3. *For affine latencies, that is, $\ell_e(x) = a_e x + b_e$, with $a_e, b_e \ge 0$, if f is a Nash equilibrium flow and f^* is a socially optimal slow (that is $L(f^*)$ is minimal), then*

$$L(f) \le \frac{4}{3} L(f^*).$$

Proof. Recall, because there is no advantage to shifting from f to f^*,

$$L(f) = \sum_e F_e \ell_e(F_e) \le \sum_e F_e^* \ell_e(F_e).$$

$$\begin{aligned} L(f) - L(f^*) &= \sum_e (F_e \ell_e(F_e) - F_e^* \ell_e(F_e^*)) \\ &\le \sum_e F_e^* (\ell_e(F_e) - \ell_e(F_e^*)) \\ &= \sum_e F_e^* a_e (F_e - F_e^*) \\ &= \sum_e a_e ((F_e/2)^2 - (F_e^* - F_e/2)^2) \qquad \text{(more magic)} \\ &\le \frac{1}{4} \sum_e F_e (a_e F_e + b_e) \end{aligned}$$

$$= \frac{L(f)}{4}.$$

So $L(f) \leq (4/3)L(f^*)$. $\square$

Corollary 1.2. *For affine latency functions, the price of anarchy is* $\leq 4/3$.

1.3 The impact of adding edges

As we saw before, adding edges to a network can reduce efficiency. We can quantify this in relation to the price of anarchy.

Theorem 1.4. *Consider a network G with a Nash equilibrium from f_G and average latency $L_G(f_G)$ and a network H with additional roads added. Suppose that the price of anarchy in H is no more than α. Then any Nash equilibrium flow f_H has average latency*

$$L_H(f_H) \leq \alpha L_G(f_G).$$

Proof.

$$L_H(f_H) \leq \alpha L_H(f_H^*) \leq \alpha L_H(f_G^*) = \alpha L_G(f_G^*) \leq \alpha L_G(f_G).$$ $\square$

Removing edges might improve the Nash equilibrium flow's latency by up to the price of anarchy. Which edges should we remove? It turns out finding the best edges to remove is NP-hard. For affine latencies, even finding edges to remove that gives approximately the biggest reduction is NP-hard! It's east to efficiently compute a Nash equilibrium flow that approximates the minimal latency Nash equilibrium flow within a factor of 4/3.; just compute a Nash equilibrium flow for the full graph. Nothing better is possible; assuming P $\neq$ NP, there is no $(4/3 - \varepsilon)$-approximation algorithm.