Selfish routing & the price of anarchy

(continued)

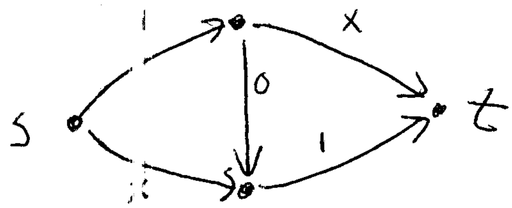

based on monograph by Tim Roughgarden

Previously, . . .

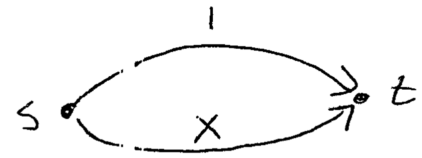

The Model: An instance (G, r, c)

G = a graph $G=(V,E)$ with directed edges

r = amount of traffic to be routed from source (s) to sink (t).

c = cost functions on edges (increasing, continuous, non-negative ...)

$\mathcal{P}$ = paths from s to t

f = flow vector

f_e = flow induced on edge e

- Cost of a path w.r.t. a flow f:

$$C_P(f) = \sum_{e \in P} c_e(f_e)$$

- Cost of a flow

$$C(f) = \sum_{P \in \mathcal{P}} C_P(f) f_P$$

or (reversing summation)

$$C(f) = \sum_{e \in E} c_e(f_e) f_e$$

Theorem 1: Instances have an (essentially unique) Nash Equilibrium.

Let f be a Nash flow & f^* a min-cost flow for an instance (G, r, c)

Define the price of anarchy

$$\rho(G, r, c) = \frac{C(f)}{C(f^*)}$$

Example (Pigou)

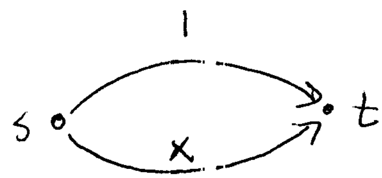

$r = 1$

Nash: all flow on bottom

$\text{Cost} = 0 \cdot 1 + 1 \cdot 1 = 1$

OPT: $\frac{1}{2}$ on each edge

$\text{cost} = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$

price of anarchy $= \frac{4}{3}$

<u>Theorem 2</u>: If edge costs are linear then $\rho \leq \frac{4}{3}$.

Example (Non-linear Pigou)

$$\begin{tikzcd} s \arrow[r, bend left, "1"] \arrow[r, bend right, "x^p"'] & t \end{tikzcd} \qquad r = 1$$

Price of anarchy $\xrightarrow{p \to \infty} \infty$

<u>Theorem 3:</u> If f is Nash for (G, r, c), and f^* is OPT for $(G, 2r, c)$, then

$$C(f) \leq C(f^*)$$

<u>Corollary 4:</u> Let f and f_δ be Nash flows for instances (G, r, c) and (G, r', c) respectively, with $r' \leq r/(1+\delta)$, $\delta > 0$. Then,

$$C(f_\delta) \leq \frac{C(f)}{\delta \cdot \rho(G, r, c)}$$

Question 1: What is the price of anarchy if we are allowed to choose the amount of traffic?

NL Pigou:

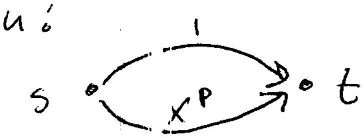

$r < 1$ Nash = OPT

$r > 1+\varepsilon$. $\rho \sim \frac{1+\varepsilon}{\varepsilon}$ ($\to 1$ as $r \to \infty$)

Def: Let f and f_1 be Nash flows for (G, r, c) and $(G, r/2, c)$ respectively. Then

$$\Pi(G, r, c) \stackrel{\text{def}}{=} \frac{C(f)}{C(f_1)}$$

Note $\rho(G, ; c) \le \Pi(G, r, c)$

(Proof: By thm. 3, $C(f_1) \le \frac{C(f)}{\rho}$)

Π gives us a way to bound ρ.

Theorem 5: Suppose (G, r, c) is an instance and $\rho(G, \lambda r, c) \ge \rho^*$ for all $\lambda \in [1-l, 1]$, $l \le 1/2$.

Then

$$l \cdot \rho^* = O(\ln \Pi(G, r, c))$$

Idea: Apply Cor. 4 repeatedly

Proof: Suppose $\rho(G, \lambda r, c) \geq \rho^*$ for all $\lambda \in [1-\ell, 1]$, $\ell \leq 1/2$

Let $\delta = 2/\rho^* < 1$

By induction & Cor. 4,

Nash cost for $(G, r, c) \geq 2^i$ (Nash cost for $(G, \lambda r, c)$)

if
- $\lambda \leq (1+\delta)^{-i}$
- $(1+\delta)^{-(i-1)} \geq 1-\ell$

$\ell \leq 1/2 \Rightarrow (1+\delta)^{-(i-1)} \geq 1-\ell$ if $i = O(\ell/\delta)$

Set $i = \Theta(\ell/\delta) = \Theta(\ell \cdot \rho^*)$

$\lambda = 1/2$

to get $T(G, r, c) \geq 2^{\Theta(\ell \cdot \rho^*)}$

$\Rightarrow \ell \rho^* = O(\ln \pi(G, r, c))$

<u>Question 2</u>: How much can the price of anarchy be reduced by a little central control?

Set-up: We'll consider only networks of parallel edges

Suppose a central authority is allowed to control a β-fraction of the traffic, with the rest routed selfishly?
(Roughgarden calls this "Stackelberg Routing")

Example: (can do better) $\beta = 1/2$

Pigou

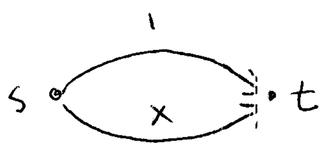

$r = 1$

If we route $\frac{1}{2}$ on the top edge, we get an optimal flow.

Example: (can't always get optimal, or even improvement)

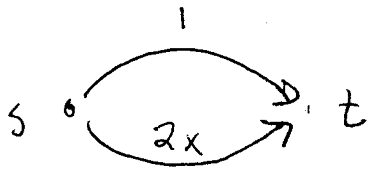

$r = 1$

$\beta = 1/2$

Nash flow is $1/2$ on each edge, cost $= \frac{1}{2}\cdot 1 + \frac{1}{2}\cdot 1 = 1$

No matter how we choose to route our flow,
we get the Nash flow.

OPT $= 3/4$ in top, $1/4$ on bottom

OPT cost $= \frac{3}{4}\cdot 1 + \frac{1}{4}\cdot\frac{1}{2} = \frac{7}{8}$

How should we route the traffic we control?

Strategy 1: "Aloof"

- Route traffic so it incurs min-cost

Bad example: Pigou

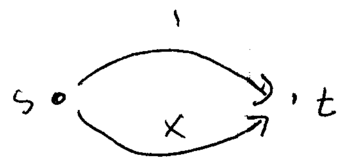

$r = 1$

$\beta = 1/2$

Aloof routes all traffic on bottom, inducing Nash flow.

No improvement.

Strategy 2: "Scale"

- Scale the optimal flow down ~~by~~ factor of β.

Bad example:

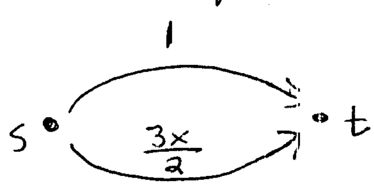

$r = 1$

$\beta = 1/2$

OPT = $2/3$ on top, $1/3$ on bottom

$$\text{cost} = \frac{2}{3} \cdot 1 + \frac{1}{3} \cdot \frac{1}{2} = \frac{5}{6}$$

Scale routes $\frac{1}{3}$ on top, $\frac{1}{6}$ on bottom

All selfish traffic goes on bottom $\Rightarrow$ Nash

$$\text{cost} = \frac{1}{3} \cdot 1 + \frac{2}{3} \cdot 1 = 1$$

Strategy 3: "Largest Cost First" (LCF)

- compute optimal flow
- saturate edges from most to least costly.

Example:

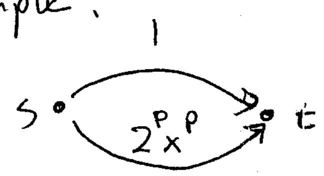

$r = 1$

$\beta = 1/2$

Nash: $\frac{1}{2}$ on each edge, cost $= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1$

OPT: $\frac{1}{2} + \varepsilon_p$ on top, $\frac{1}{2} - \varepsilon_p$ on bottom, cost $\to \frac{1}{2}$ as $p \to \infty$

Any strategy induces the Nash flow.

- So no strategy can do better than $\rho = 2$
- In general no strategy can do better than $1/\beta$

<u>Theorem 6:</u> LCF achieves $\rho = 1/\beta$ on networks of parallel edges.

Idea: Induction on # of edges

Proof:

<u>Lemma 1</u>: Saturation by LCF drives away selfish traffic.

Proof: OPT flow is an example of a flow where all other traffic can choose cheaper edges.

Thus, there is an equilibrium where no selfish traffic will choose the saturated edge. $\square$

WLOG assume $r=1$

Case 1: Suppose at least one edge is saturated.

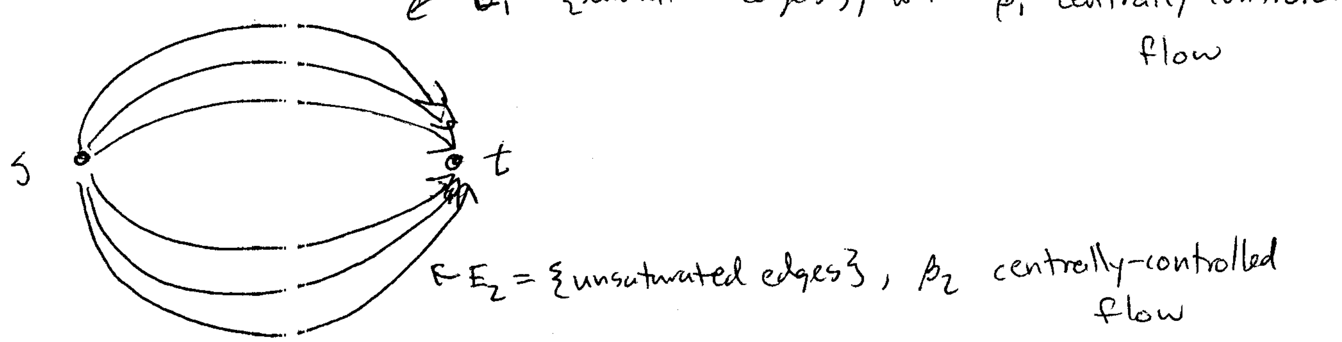

$\beta_1 + \beta_2 = \beta$

Let f^* be OPT flow, f be LCF flow, g selfish flow

Want to show $C(f+g) \le \frac{1}{\beta} C(f^*)$

Let L be cost of all edges with selfish flow >0.

C_1 = cost on E_1, C_2 = cost on E_2

(so $C(f+g) = C_1 + C_2$)

Restricting f to E_2 is LCF for instance $(G_2, 1-\beta_1, c)$ with $\beta = \beta'$

$$\beta' = \frac{\beta_2}{1-\beta_1}$$

Inductive hypothesis $\Rightarrow C(f^*) \geq C_1 + \beta' C_2$

Thus showing $C(f \cdot g) \leq \frac{1}{\beta} C(f^*)$

reduces to ~~showing~~ showing $\beta(C_1 + C_2) \leq C_1 + \beta' C_2$ (✡)

Note $C_2 = (1-\beta_1)L$, $C_1 \geq \beta_1 L$

(✡) is equivalent to

$$\beta(C_1 + (1-\beta_1)L) \leq C_1 + \beta'(1-\beta_1)L$$

$$\Leftrightarrow \left(\beta(1-\beta_1) - \beta_2\right) L \leq (1-\beta) C_1$$

if $(\beta(1-\beta) - \beta_2)L \leq (1-\beta)\beta_1 L$

$\Leftrightarrow \beta - \beta_2 \leq \beta_1$ ✓

Case 2: Suppose no edge is saturated by LCF.

Then $C(f+g) = L$

Note cost of most expensive edge in $f^* \geq L$

say this is the m^{th} edge.

We get

$$C(f^*) \geq f_m^* c_m(f_m^*) \geq \beta L = \beta C(f+g)$$

Theorem 6 applies only in networks of parallel edges.

It cannot be extended to more general networks:

Example:

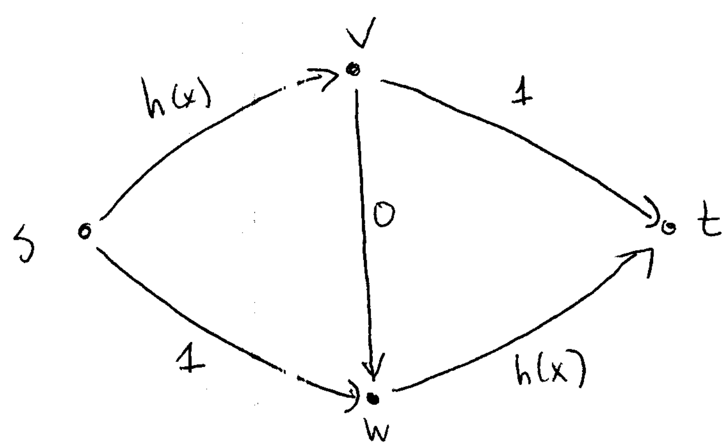

$$h(x) = \begin{cases} 0 & x \in [0, \frac{3}{4} - \varepsilon] \\ 1-\varepsilon & x \geq \frac{3}{4} \end{cases}$$

$r = 1$

OPT:

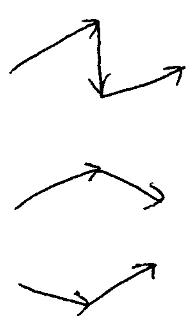

$\frac{1}{2} - 2\varepsilon$

$\frac{1}{4} + \varepsilon$

$\frac{1}{4} + \varepsilon$

cost $\approx \frac{1}{2}$

(Can show any strategy with $\beta = \frac{1}{2}$ gives cost > 1,

since: • must have $\frac{1}{4} + \varepsilon$ on $\{s,w\}$ or $\{v,t\}$ c.c. flow

• all selfish flow will take 3-hop path

$\Rightarrow$ cost $\geq \frac{5}{4}$

Open Problem: Does some version of Theorem 6 hold with $1/\beta$ replaced by some larger function of β for more general networks?