How Bad is Selfish Routing?

Tim Roughgarden
Cornell University

joint work with Éva Tardos

Traffic in Congested Networks

Mathematical model:

- A directed graph $G = (V,E)$
- source–sink pairs s_i, t_i for $i=1,..,k$
- rate $r_i \geq 0$ of traffic between s_i and t_i for each $i=1,..,k$
- For each edge e, a latency function $l_e(\bullet)$

$r_1 = 1$

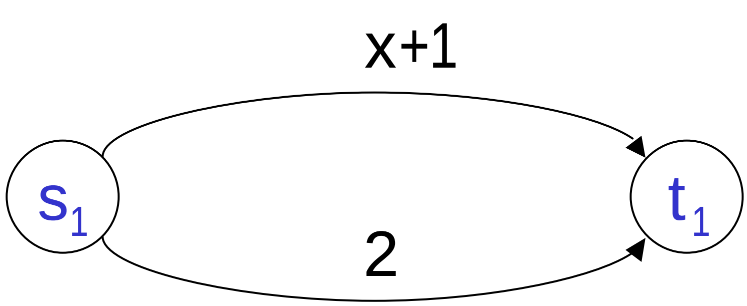

Example

Traffic rate: r = 1, one source-sink

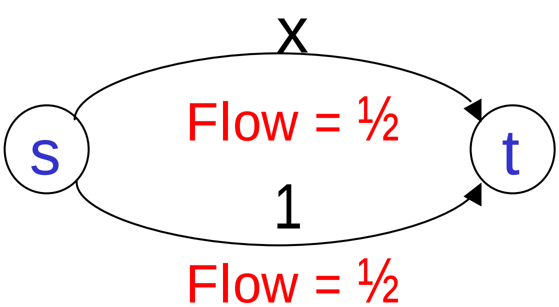

Total latency = $\frac{1}{2}\cdot\frac{1}{2} + \frac{1}{2}\cdot 1 = \frac{3}{4}$

But traffic on lower edge is envious.

An envy free flow:

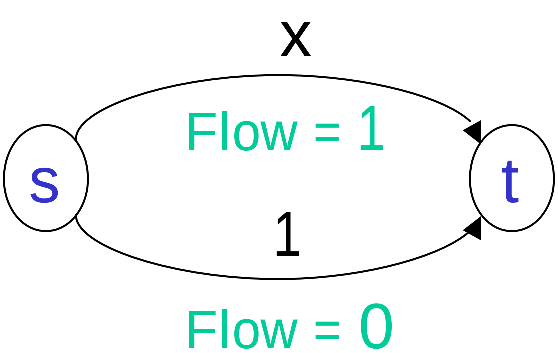

Total latency = 1

Flows

Traffic and Flows:

- f_P = amount routed on s_i-t_i path P

flow vector f ⇔ traffic pattern at steady-state

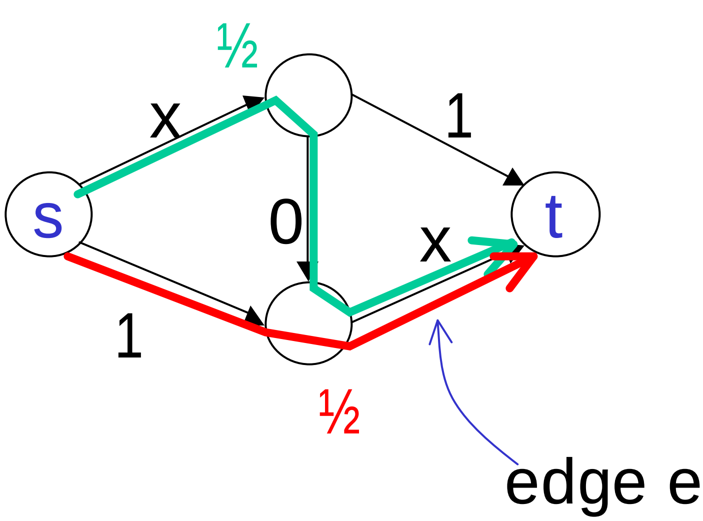

$f_e = ½ + ½ = 1$

$l_e(f) = 1$

Cost of a Flow

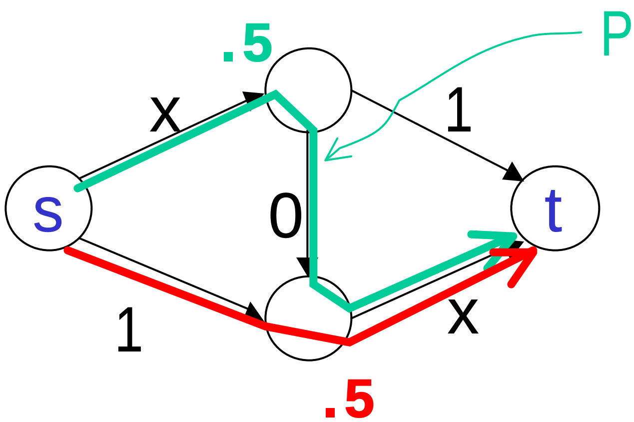

$l_P(f) = .5 + 0 + 1$

Latency along path P:

- $l_P(f)$ = sum of latencies of edges in P

The Cost of a Flow f:

= total latency

- $C(f) = S_P f_P \bullet l_P(f)$

Flows and Game Theory

- flow = routes of many noncooperative agents
- Examples:
  - cars in a highway system
  - packets in a network
    - [at steady-state]
- cost (total latency) of a flow as a measure of social welfare
- agents are selfish
  - do not care about social welfare
  - want to minimize personal latency

Flows at Nash Equilibrium

Defn: A flow is at Nash equilibrium (or is a Nash flow) if no agent can improve its latency by changing its path

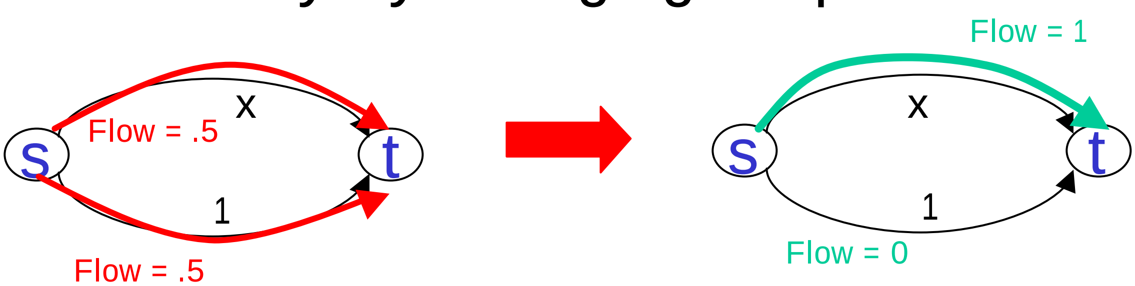

Assumption: edge latency functions are continuous, nondecreasing

Lemma: a flow f is a Nash flow if and only if all flow travels along minimum-latency paths (w.r.t. f).

Nash Flows and Social Welfare

Central Question:

- What is the cost of the lack of coordination in a Nash flow?

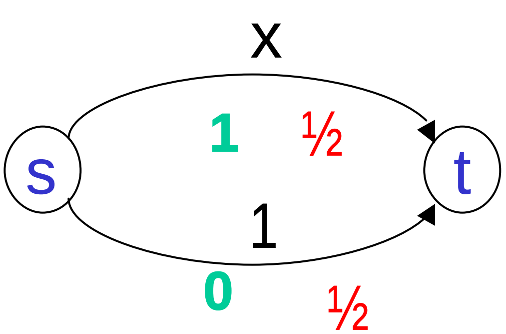

- Cost of Nash = 1
- min-cost

  = ½•½ + ½•1 = ¾

Analogous to IP versus ATM:

- ATM ≈ central control ≈ min cost
- IP ≈ no central control ≈ selfish

What Is Know About Nash?

Flow at Nash equilibrium exists and is essentially unique

[Beckmann et al. 56], ...

Nash and optimal flows can be computed efficiently [Dafermos/Sparrow 69], ...

Network design: what networks admit "good" Nash flows? [Braess 68], ...

The Braess Paradox

Better network, worse delays:

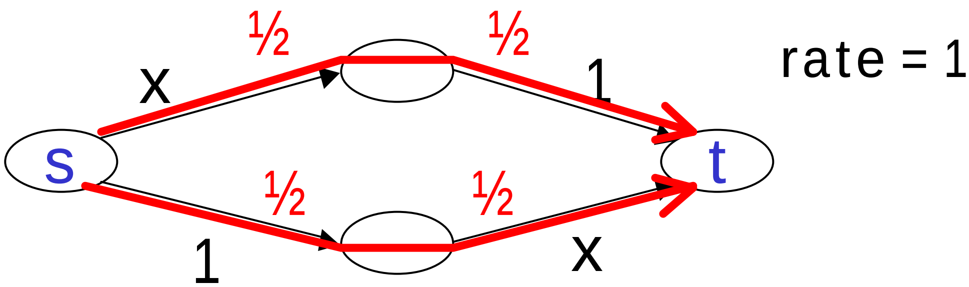

- Cost of Nash flow = 1.5

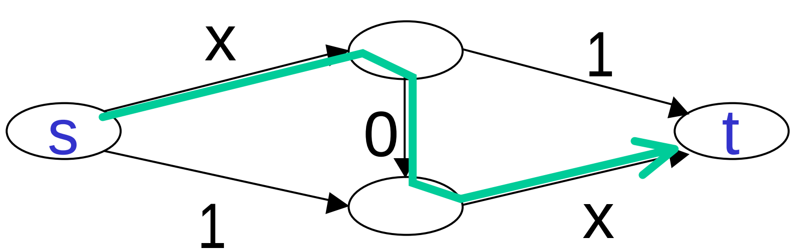

- Cost of Nash flow = 2

All the flow has increased delay!

Our Results for Linear Latency

latency functions of the form
$l_e(x)=a_e x+b_e$

the cost of a Nash flow is at most 4/3 times that of the minimum-latency flow

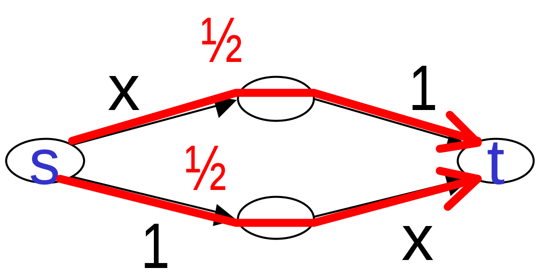

Delay = 1.5

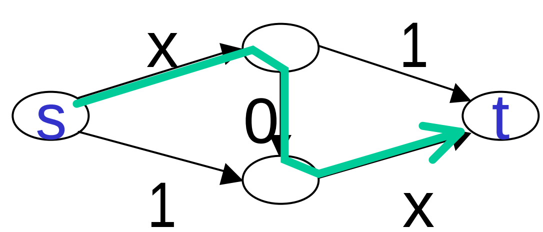

Delay = 2

General Latency Functions?

Bad Example: (r = 1, i large)

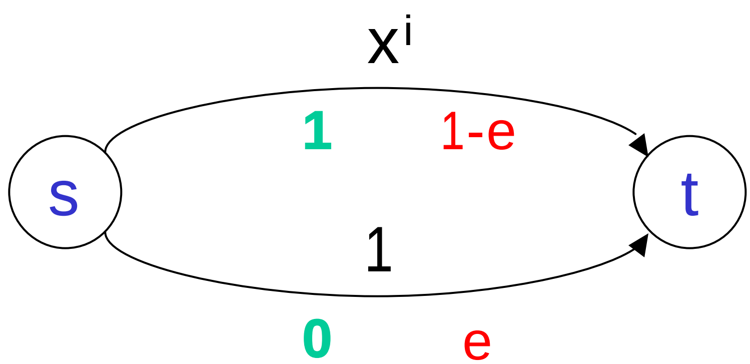

Nash flow cost =1, min cost ≈ 0

Þ Nash flow can cost arbitrarily more than the optimal flow

Our Results for General Latency

In any network with latency functions that are

- continuous,
- non-decreasing

the cost of a Nash flow with rates r_i for i=1,..,k

is at most the cost of a minimum cost flow with rates $2r_i$ for i=1,..,k

Morale for IP versus ATM?

IP today no worse than
ATM a year from now …

Instead of
- building central control
- build networks that support twice as much traffic

What Is the Minimum-cost Flow Like?

Minimize

$$C(f) = S_e\, f_e \bullet l_e(f_e)$$

- by summing over edges rather than paths
- f_e amount of flow on edge e

Cost $C(f)$ usually convex

- e.g., if $l_e(f_e)$ convex
- if $l_e(f_e) = a_e f_e + b_e$

$\Rightarrow$ $C(f) = S_e\, f_e \bullet (a_e f_e + b_e)$ convex quadratic

Why Is Convexity Good?

A solution is optimal for a convex cost if and only if

- tiny change in a locally feasible direction cannot decrease the cost

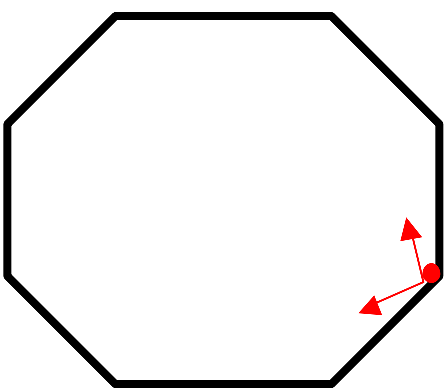

Characterizing the Optimal Flow

Direction of change: moving a tiny flow from one path to another

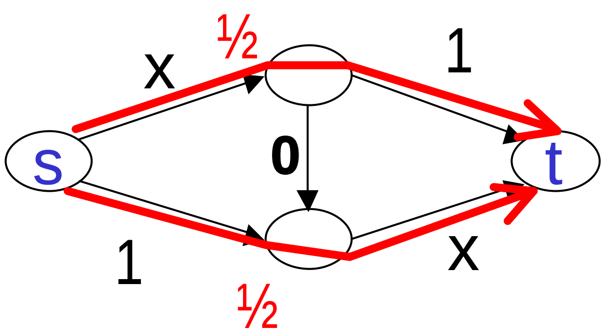

flow f is minimum cost if and only if cost cannot be improved by moving a tiny flow from one path to another

Characterizing the Optimal Flow

Cost $f_e \bullet l_e(f_e)$ Þ marginal cost of increasing flow on edge e is

$$l_e(f_e) + f_e \bullet l_e'(f_e)$$

latency of new flow

Added latency of flow already on edge

Key Lemma: a flow f is optimal if and only if all flow travels along paths with minimum marginal cost (w.r.t. f).

Min-cost Is a Socially Aware Nash

flow f is minimum cost if and only if all flow travels along paths with minimum marginal cost

Marginal cost: $l_e(f_e) + f_e \cdot l_e'(f_e)$

flow f is at Nash equilibrium if and only if all flow travels along minimum latency paths

Latency: $l_e(f_e)$

Consequences for Linear Latency Fns

Observation: if $l_e(f_e) = a_e f_e + b_e$

Þ marginal cost of P w.r.t. f is:

$$\sum_{e \in P} 2a_e f_e + b_e$$

Corollaries

- if $a_e = 0$ for all e, Nash and optimal flows coincide (obvious)
- if $b_e = 0$ for all e, Nash and optimal flows coincide (not as obvious)

Example

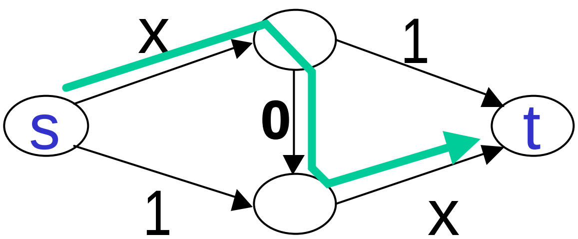

Edge cost = x^2 Þ
marginal cost = $2x$

- Nash flow of rate 1, latency L=2

- Note: Same flow for rate ½,
  - All paths have marginal cost = 2

⇒ it is min-cost for rate ½,

Key Observation

Nash flow f for rate r

- all flow paths have latency L

  $\Rightarrow C(f) = rL$

$\Rightarrow f/2$ is optimal with rate $r/2$ and

- all flow paths have marginal cost L

Bound for Nash: Linear Latency

Goal: prove that cost of opt flow is at least 3/4 times the cost of a Nash flow f

Cost of opt at rate r = Cost of opt at rate r/2 + Cost of increasing rate from rate r/2 to rate r

opt is f/2

C(f/2) ≥ ¼C(f)

At least (r/2)•L ≥ ½C(f)

Nonlinear Latency

Goal: cost of a Nash flow with rate r is at most the cost of the optimal flow with rate 2r

Analogous proof sketch??

Cost of opt at rate 2r = Cost of opt at rate r + Cost of augmenting opt flow at rate r to opt at rate 2r

Troubles:

Can be close to zero

What is opt at rate r? and what is its marginal cost?

Other Models?

- An approximate version of Theorem for non-linear latency with imprecise evaluation of path latency
- Analogue for the case of finitely many agents (splittable flow)
- Impossibility results for finitely many agents, unsplittable flow, i.e.,
  - if each agent i controls a positive amount of flow $r_i \geq 0$
  - flow of a single agent has to be routed on a single path

Other Games?

Koutsoupias & Papadimitriou

STACS'99

- scheduling with two parallel machines
- Negative results for more machines

First paper to propose quantifying the cost of a lack of coordination

- What other games have good Nash equilibrium?

More Open Questions

- Is there any model in which positive results are possible for unsplittable flow?
- Consider models in which agents may control the amount of traffic (in addition to the routes)
  - Problem: how to avoid the "tragedy of the commons"?