How Unfair is Optimal Routing?

Tim Roughgarden
Cornell University

Traffic in Congested Networks

The Model:

- A directed graph $G = (V,E)$
- A source s and a sink t
- A rate r of traffic from s to t
- For each edge e, a latency function $l_e(\cdot)$

Example: $(r=1)$

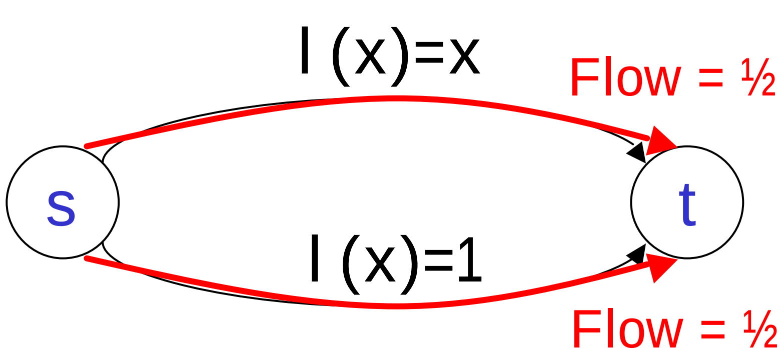

Flows and their Cost

Traffic and Flows:

- f_P = amount of traffic routed on s-t path P
- flow vector $f \Leftrightarrow$ traffic pattern at steady-state

The Cost of a Flow:

- $l_P(f)$ = sum of latencies of edges on P (w.r.t. the flow f)
- $C(f)$ = cost or total latency of flow f:
  $S_P f_P \bullet l_P(f)$

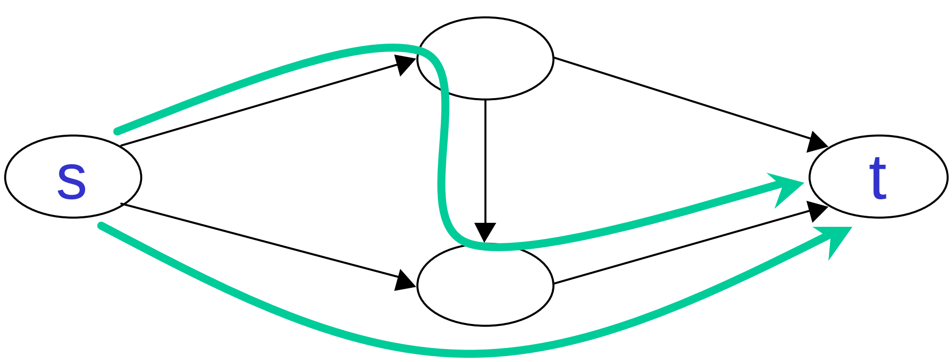

Flows and Game Theory

- flow = routes of many noncooperative agents
- Examples:
  - cars in a highway system [Wardrop 52]
  - packets in a network
- 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

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

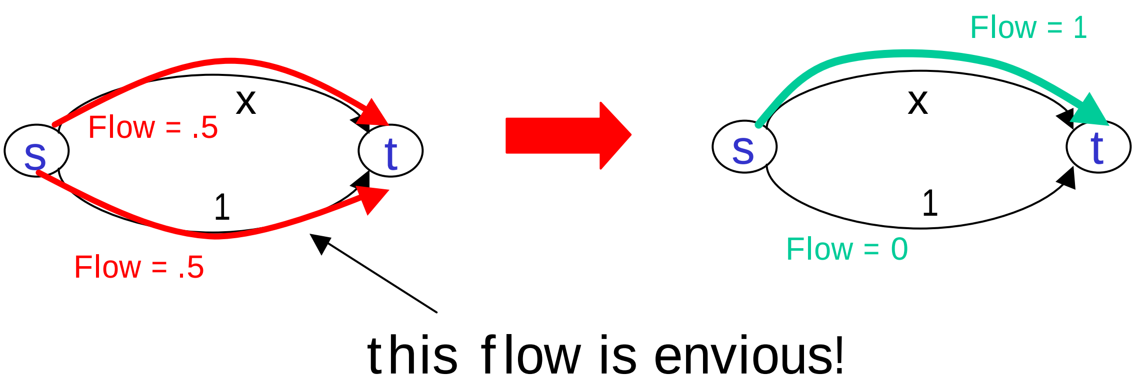

Assumption: edge latency functions are continuous, nondecreasing

Lemma: f is a Nash flow ⟺ all flow on minimum-latency paths (w.r.t. f)

Fact: have existence, uniqueness

Nash Flows and Social Welfare

Fact: Nash flows do not optimize total latency

- Þ lack of coordination leads to inefficiency

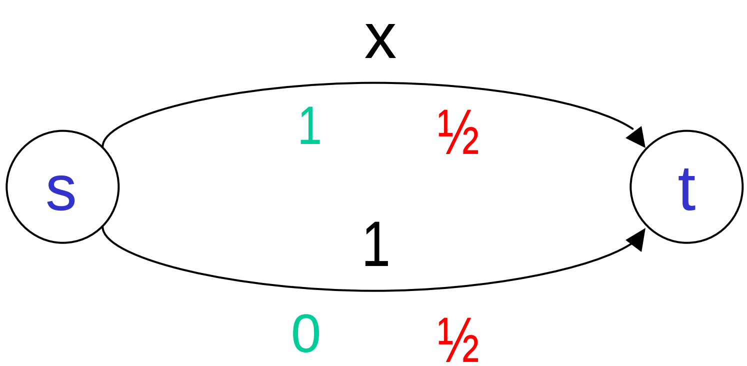

Cost of Nash flow = 1•1 + 0•1 = 1

Cost of optimal (min-cost) flow
= ½•½ +½•1 = ¾

How Bad is Selfish Routing?

- [Roughgarden/Tardos 00]
  - linear latency functions Þ
    cost of Nash = 4/3 × cost of OPT
  - bicriteria result for arbitrary fns
- [Roughgarden 01,02]: other latency fns
- [Friedman 01]: includes flow control
- Different model, objective fn:
  - [Koutsoupias/Papadimitriou 99], [Mavronicolas/Spirakas 01], [Czumaj/Vöcking 02]

Question: Is the optimal (min-cost) routing really what we want?

- what about fairness?

Bad example

r = 1, ? small

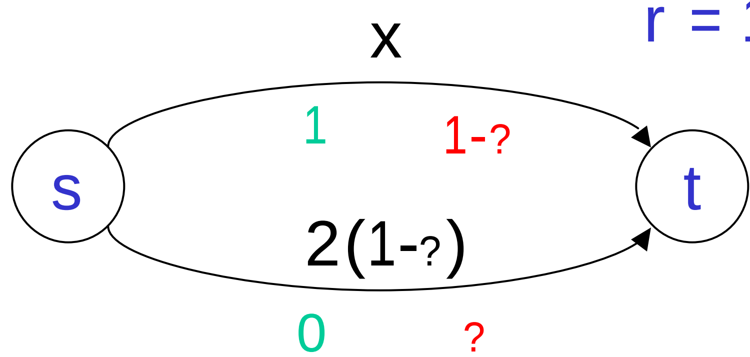

Þ some “martyrs” incur twice as much latency in OPT as in Nash!

Even worse:

r = 1, k large

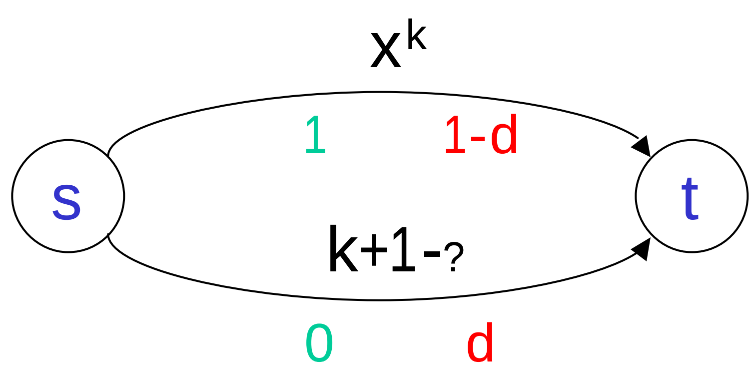

Þ some traffic can be arbitrarily worse off in OPT than in Nash

How Unfair is Optimal Routing?

Def: Given a network G, latency fns l, traffic rate r:

$$\text{unfairness of } (G, r, l) := \frac{\text{max latency in OPT}}{\text{common latency in Nash}}$$

Examples:

- Braess's Paradox (unfairness = ¾)
- bad example (unfairness ≈ k+1)

Central Question: What is the worst-possible unfairness?

- for a restricted class of latency fns

Informal Statement of Main Results

"Thm": In any network with latency fns that are "not too steep", unfairness is "small".

Special case: A network with with polynomial latency fns, max degree = k, has unfairness = k+1

Matching lower bound:

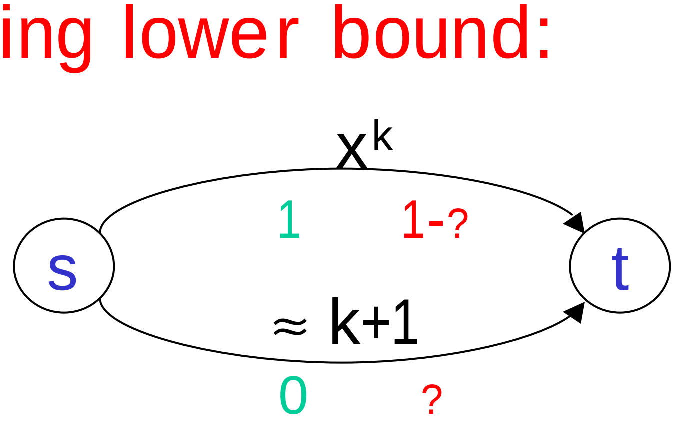

Characterizing the Optimal Flow

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).

The Optimal Flow as a Socially Aware Nash

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

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

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

Latency: $l_e(f_e)$

Main Theorem

Thm: For a network G w/latency fns l, suppose worst-case marginal cost vs. latency discrepancy is:

$$\max_{e,x} \frac{l_e(x) + x \cdot l_e'(x)}{l_e(x)} = ?.$$

Then, unfairness of G is $= ?$.

Example: if $l_e(x) = x^k$ get

$$\frac{x^k + k \cdot x^k}{x^k} = k+1$$

Proof Sketch

Lemma 1: Minimum-latency path in OPT = common latency in Nash.

- otherwise, Nash would have smaller total latency than OPT

Lemma 2: Latencies of OPT's flow paths differ by = a ? factor.

- OPT flow paths have equal marginal cost
- marginal costs, latencies differ only by a ? factor

Conclusions

Remark: proof actually shows that for any feasible flow f,

max-latency of an OPT path =? × max-latency of an f path

Fact: [Meyerson 01] False with multiple commodities!

- OPT can be arbitrarily less fair than other feasible flows
  - even with linear latency functions
  - under many definitions of fairness

Open: is OPT almost as fair as Nash w/many commodities?