Selfish Routing

Tim Roughgarden
Cornell University

joint work with Éva Tardos

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 $l_e(\bullet)$

Example: $(r=1)$

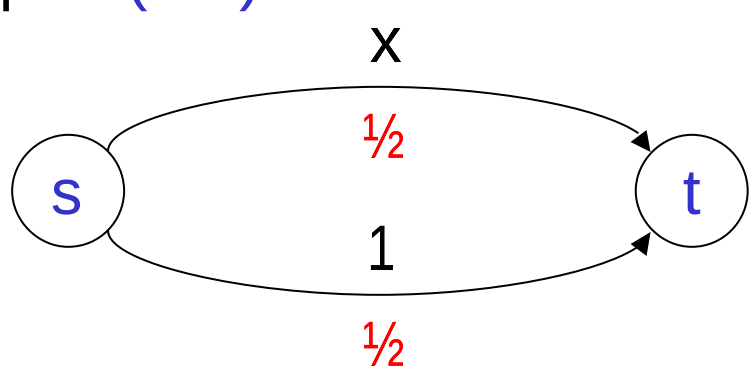

Flows and their Cost

Traffic and Flows:

- f_P = amount of traffic routed on s-t path P
- flow vector f ⇔ 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: $\Sigma_P f_P \bullet l_P(f)$

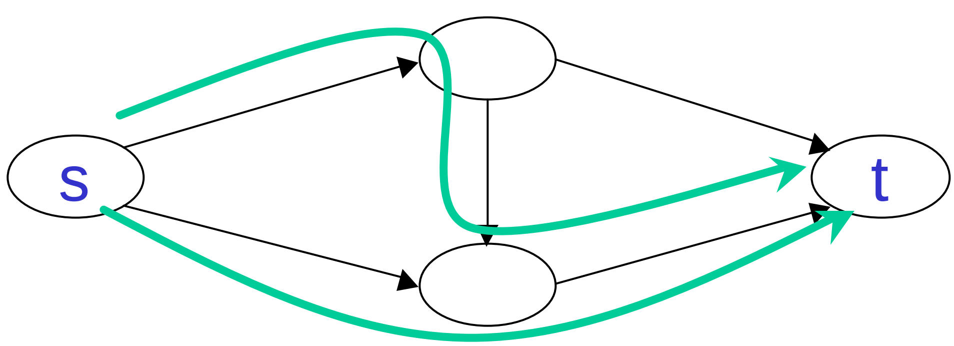

Flows and Game Theory

- 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

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

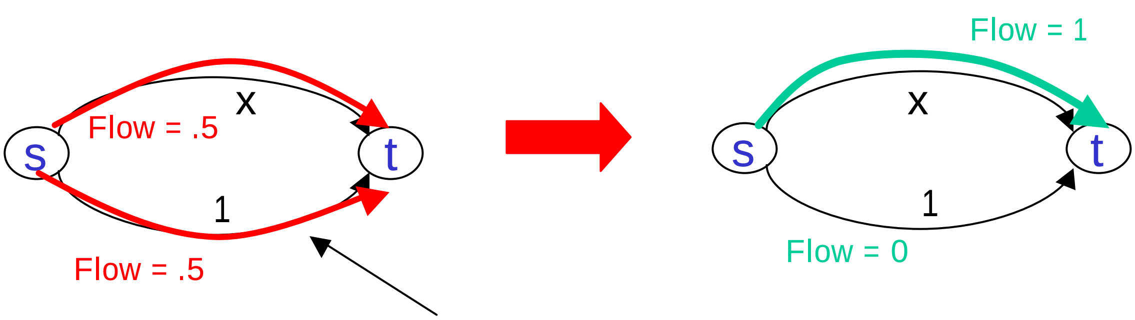

this flow is envious!

Assumption: edge latency functions are continuous, nondecreasing

Lemma: 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?

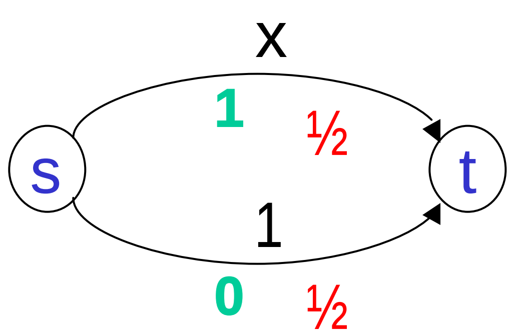

- Cost of Nash = 1
- min-cost
  = ½•½ + ½•1 = ¾

Analogous to IP versus ATM:

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

Previous Work

- [Beckmann et al. 56], ...
  - Existence, uniqueness of flows at Nash equilibrium
- [Dafermos/Sparrow 69], ...
  - Efficiently computing Nash and optimal flows
- [Braess 68], ...
  - Network design
- [Koutsoupias/Papadimitriou 99]
  - Quantifying the cost of a lack of coordination

Braess's Paradox

Better network, worse Nash flow:

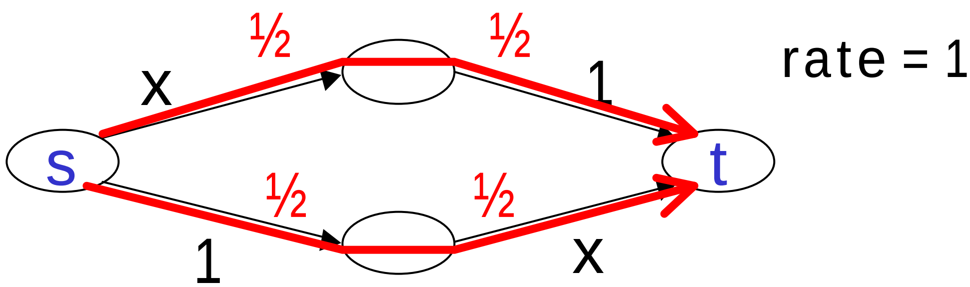

Cost of Nash flow = 1.5

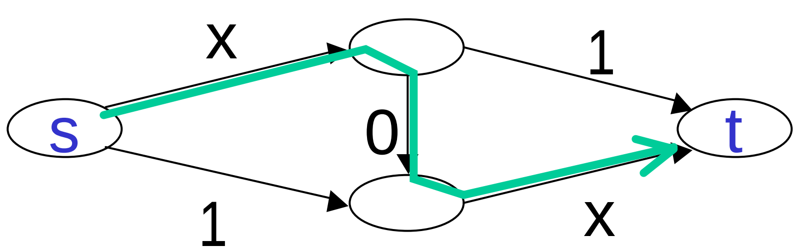

Cost of Nash flow = 2

All traffic experiences increased latency!

Our Results for
Linear Latency

Def: a linear latency function is of the form $l_e(x)=a_e x+b_e$

Theorem 1:

linear latency fns Þ

cost of Nash flow = 4/3 × cost of opt flow

General Latency Functions?

Bad Example: (r = 1, k large)

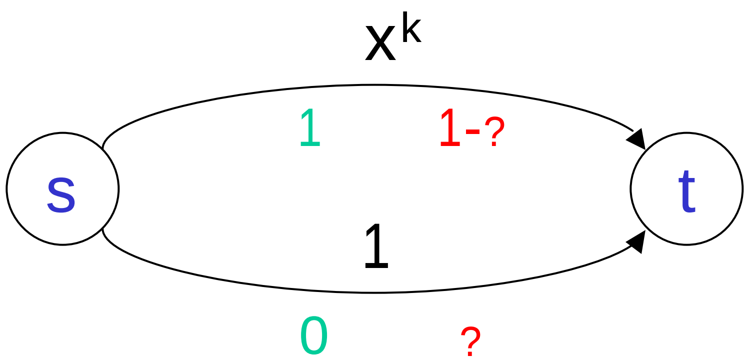

Nash flow has cost 1, min cost ≈ 0

Þ Nash flow can cost arbitrarily more than the optimal (min-cost) flow

- even if latency functions are polynomials

Our Results for General Latency

All is not lost: the previous example does not preclude interesting bicriteria results.

Theorem 2:
continuous, nondecreasing latency functions Þ

cost of Nash flow at rate r	=	cost of opt flow at rate 2r

Optimal Flows + Convexity

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 fn if and only if

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

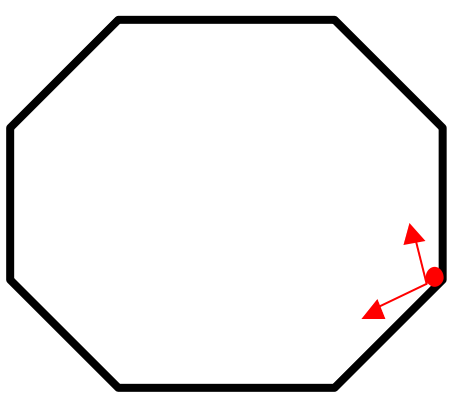

feasible directions

Characterizing the Optimal Flow

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

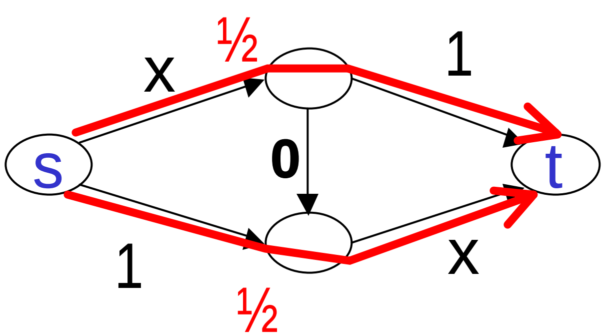

flow f is minimum cost iff its cost cannot be improved in this way

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: $l_e(f_e)$

Consequences for Linear Latency Fns

Observation: if $l_e(f_e) = a_e f_e + b_e$ (latency functions are linear) $\Rightarrow$ marginal cost of P w.r.t. f is:

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

Corollary: if $b_e = 0$ for all e, Nash and optimal flows coincide

Key Corollary

Corollary:

marginal costs of $f/2$ = latencies of f

$$2a_e(f_e/2) + b_e = a_e f_e + b_e$$

- f a Nash flow with rate r in a network with linear latency fns Þ $f/2$ is optimal with rate $r/2$

common marginal cost of $f/2$'s paths = common latency of f's paths } := L

Proof of Theorem

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

$C(f/2) \geq \frac{1}{4}C(f)$

At least $(r/2)\bullet L \geq \frac{1}{2}C(f)$

Extensions

- Multicommodity networks
  - many source-sink pairs
- More general games
  - proofs did not use structure provided by a network
- Other classes of latency fns?
  - proof for linear case is fragile

Sources of Inefficiency

Corollary of main Theorem:

- worst-case Nash/OPT ratio is realized on a two-link network!

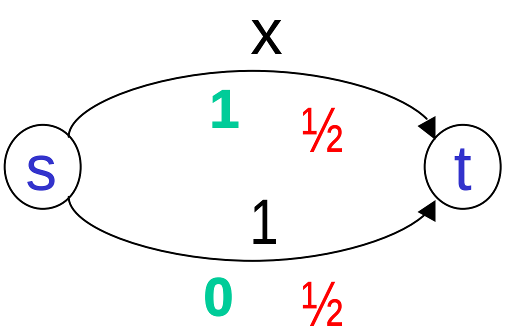

- Cost of Nash = 1
- Cost of OPT = ¾

Thus:

- one source of inefficiency:
  - confronted w/ two routes, selfish users choose incorrectly
- Corollary ⇒ that's all, folks!
  - network topology plays no role

No Dependence on Network Topology

Theorem: for (almost) any class of allowable latency fns, worst-case Nash/OPT ratio is realized on a two-link network.

Corollary: worst-case for bounded-degree polynomials is:

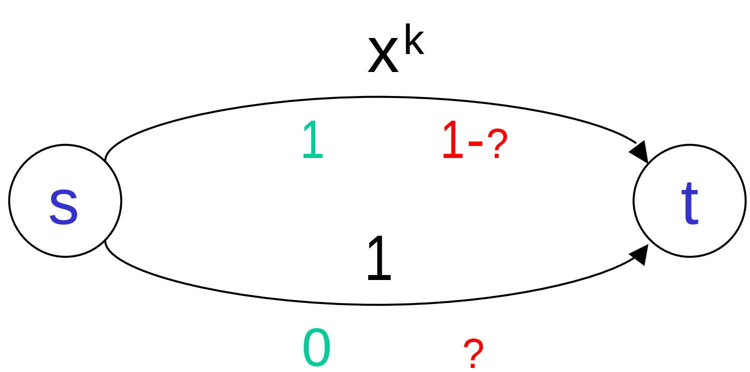

Part 2: Coping with Selfishness

Deleting Arcs to Improve a Nash Flow

Motivating Question: how can we "fix up" networks with a bad Nash flow?

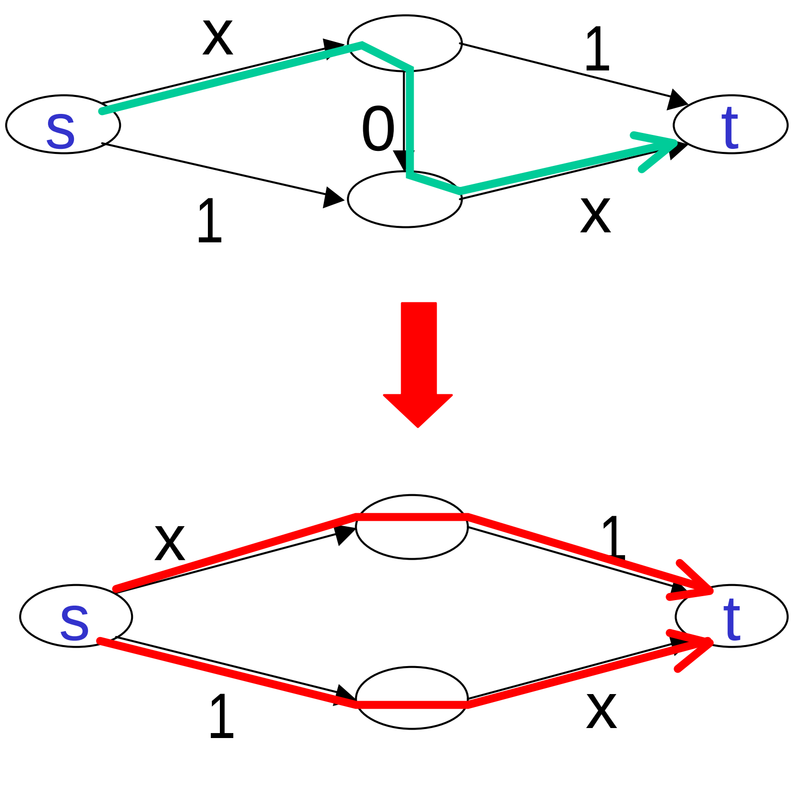

Designing Networks for Selfish Users

Formally:

- given network G = (V,E,l)
- find subnetwork minimizing latency experienced by all selfish users in a Nash flow

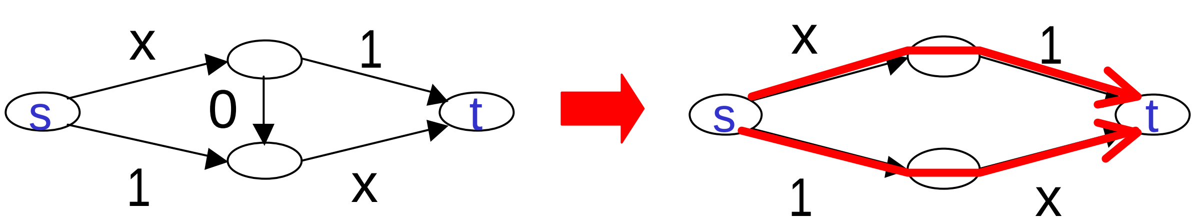

Def: The trivial algorithm is to build the entire network.

Designing Networks for Selfish Users is Hard

Our Results: the trivial algorithm is

- a 4/3-approx alg with linear latency fns (follows from [RT00])
- an n/2-approx alg with general latency fns (new analysis needed)

and

- nothing better exists! (unless P=NP)

Corollary: in general, "bad edges" cannot be detected efficiently.

Motivation

Goal: prove that Nash flows are near-optimal

- want a laissez faire approach to regulating users

Recall: false in general!

Bad example: (r=1, k large)

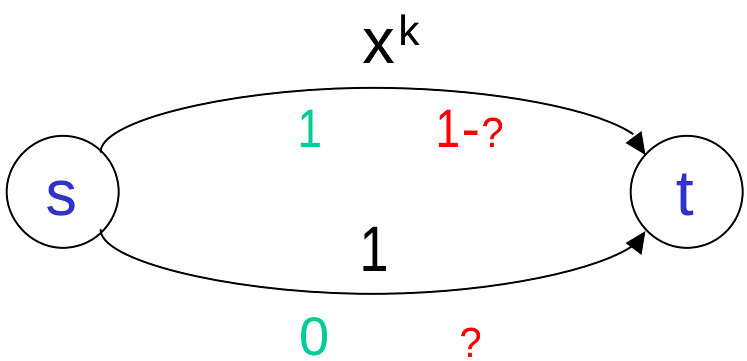

Recall our previous solutions:

- relax notion of OPT
- restrict class of allowable latency fns

Taming Selfishness through a Manager

New Approach:

- not all traffic need be controlled by selfish users
  - “centrally controlled” vs. “selfishly controlled” traffic
  - behavior of selfish users depends on assignment of managed traffic

Goal:

- assign centrally controlled traffic to induce “good” selfish behavior
  - see also [Korilis, Lazar, Orda 97]

Stackelberg Strategies

- Stackelberg strategy = assignment of centrally controlled traffic
  - Þ yields an induced equilibrium

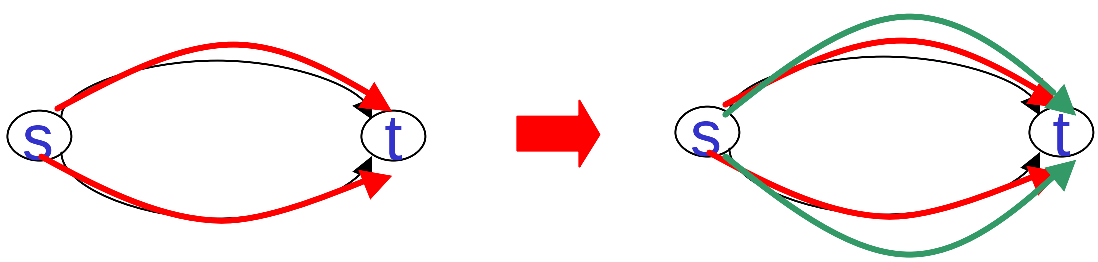

- Basic Questions:
  - what's the best strategy?
    - can we compute/characterize it?
  - how inefficient is the best induced equilibrium?
    - are we provably near-optimal?

A Constant-Factor Guarantee

Assumption: G = parallel links

Theorem: Can efficiently compute a strategy inducing an equilibrium with cost

$= (1/ß) \times$ cost of opt flow

(ß = fraction of managed traffic)

Fact: $(1/ß) \times$ OPT is best possible.

Also: can get $= [4/(3+ß)] \times$ OPT for linear latency functions.

General Graphs

Open:

- for general latency fns, fixed ß: a strategy inducing an equilibrium w/cost = f(ß) × opt
  - 1/ß not achievable in general graphs (!)
  - maybe 2/ß? (or O(n))
- for linear latency fns, a strategy w/cost < 4/3 × opt
  - e.g., is 8/7 achievable for ß=½?

Conclusions

Moral: you can be near-optimal in the presence of selfishness!

- Nash = 4/3 × OPT (linear latency fns)
  - similar results for other classes
  - multicommodity network no worse than a pair of parallel links
- Nash at rate r = OPT at rate 2r
- Stackelberg equilibrium = 1/ß × OPT
  - with a good Stackelberg strategy
  - open: general networks?

Conclusions

Bad news: selfishness leads to increased complexity

- intractability of network design

Big picture: other games?

- we've studied only one simple model
  - Nash equilibria well-characterized
- general themes:
  - are Nash equilibria near-optimal?
  - how to cope with selfishness?