The Price of Anarchy is Independent of the Network Topology

Tim Roughgarden
Cornell University

Traffic in Congested Networks

The Model:

- A directed graph $G = (V,E)$
- k source-destination pairs $(s_1, t_1), \ldots, (s_k, t_k)$
- A rate r_i of traffic from s_i to t_i
- For each edge e, a latency fn $l_e(\bullet)$ [cts, nondecreasing, convex]

Example: $(k, r=1)$

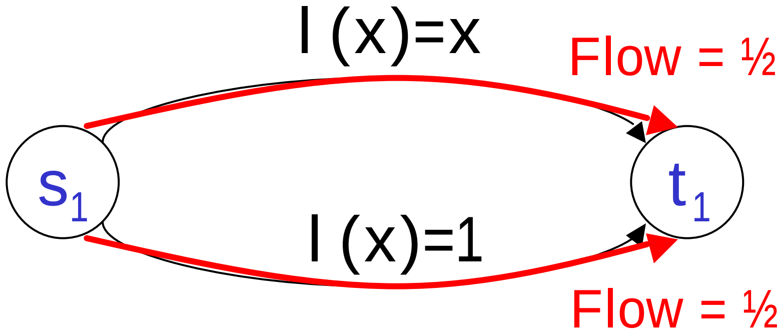

Selfish Routing

Traffic and Flows:

- f_P = amount of traffic routed on s_i-t_i path P
- flow vector f ⇔ routing of traffic

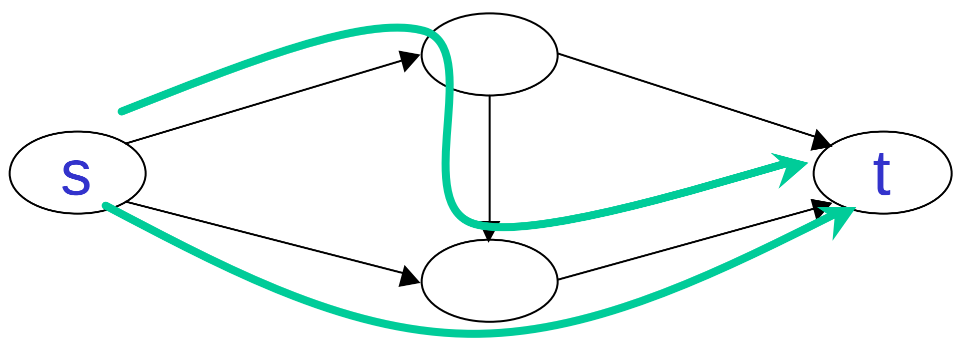

Selfish routing: what flows arise as the routes chosen by many noncooperative agents?

Nash Flows

Some assumptions:

- agents are small relative to network
- want to minimize personal latency

Def: A flow is at Nash equilibrium (or is a Nash flow) if all flow is routed on min-latency paths

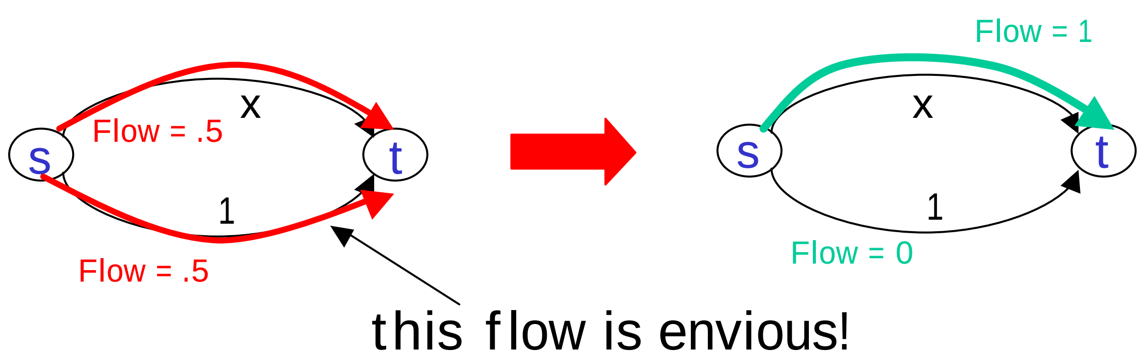

Fact: [Beckmann et al. 56] Nash flows always exist

The Cost of a Flow

Our objective function:

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

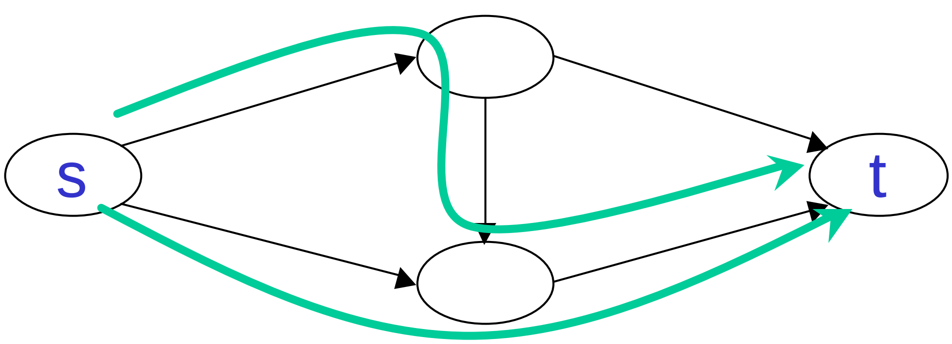

Key question: how good (or bad) are Nash flows?

The Inefficiency of Nash Flows

Fact: Nash flows do not optimize total latency [Pigou 1920]

Þ lack of coordination leads to inefficiency

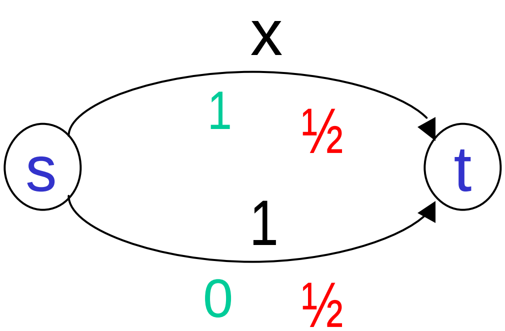

- Cost of Nash = 1
- Cost of OPT = ¾

Def: price of anarchy = worst-case Nash/OPT ratio

- also coordination ratio of [Koutsoupias/Papadimitriou 99]

Linear Latency Fns

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

Thm: [Roughgarden/Tardos 00]

network w/linear latency fns Þ

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

Cor: price of anarchy realized in a two-link network!

Point: worst-case Nash arises from overcongesting one of two available routes (and that's all)

No Dependence on Network Topology

Thm: for any class of latency fns including the constant fns, worst Nash/OPT ratio is in a two-link network.

- inefficiency of Nash flows always has simple explanation
- network topology plays no role

Note: worst ratio may be (much) larger than 4/3 with nonlinear latency fns (modify Pigou's ex)

Comparison to Previous Work

Remark: networks of parallel links are not worst-case examples for:

- Approximate Nash flows, integral Nash flows [Roughgarden/Tardos FOCS '00]
- Stackelberg equilibria [Roughgarden STOC '01]
- Braess's paradox, maximum travel time obj fn [Roughgarden FOCS '01]

Characterizing OPT

Def: f is at Nash equilibrium iff all flow travels along paths with minimum latency

Latency: $l_e(f_e)$

Lemma: [BMW 56] f is optimal iff all flow travels along paths with minimum marginal cost

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

- $l_e(f_e)$: latency of new flow
- $f_e \cdot l_e'(f_e)$: added latency for flow already on edge

Consequences for Linear Latency Fns

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

marginal cost of P w.r.t. f is:

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

Corollary: [RT00]

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

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

- f a Nash flow at rate r $\Rightarrow$ $f/2$ is optimal with rate $r/2$

Lower Bounding OPT (Linear Latency Fns)

Goal: prove that cost of OPT is $\geq 3/4$ times cost of Nash flow f

Idea: break cost of OPT into two pieces via previous Corollary

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

Cost of OPT at rate r/2: = cost of f/2, a "big chunk" of f [$\geq$ ¼C(f)]

Cost of increasing rate from r/2 to r: augmentation w.r.t. large marginal costs [$\geq$ ½ C(f)]

Proof Idea for Main Theorem

Problem: with nonlinear latency fns, f/2 (or f/c, any c) is not optimal!

Idea: scale flow by different factors on different edges

- can scale edge-by edge so that new marginal costs = old latencies

Þ equalizes marginal costs

Þ any "augmentation" should be costly

Proof Sketch (con'd)

Problem: scaling by different factors on different edges

Þ violates flow conservation!

- lower-bounding cost of the "augmentation" is tricky, must argue:
  - cut-by-cut (see proceedings)
  - edge-by-edge (simpler, see revision)

    [thanks to Amir Ronen]
- gives bound on price of anarchy; achieved in a Pigou-like example

Extensions

Thm: for any class of latency fns

- closed under scalar multiplication
- including a fn l s.t. $l(0) > 0$

the worst Nash/OPT ratio is in a network of parallel links.

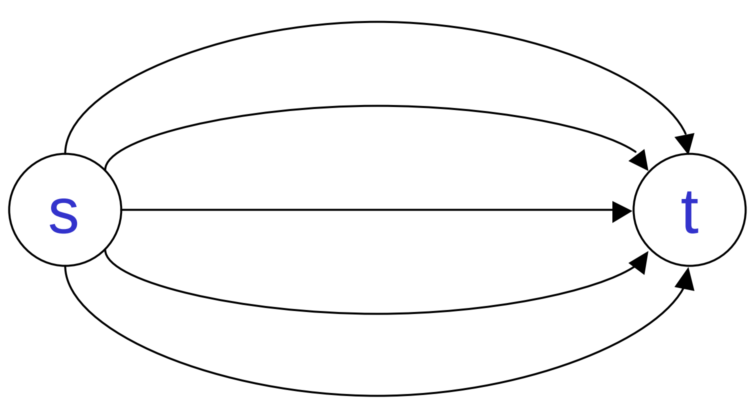

Computing the Price of Anarchy

Application: worst-case examples simple Þ price of anarchy easy to calculate

Example: polynomials with degree = d, nonnegative coeffs Þ price of anarchy T(d/log d)

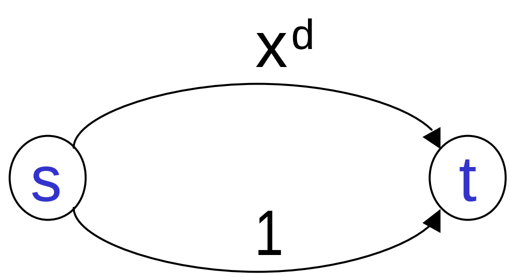

Also: M/M/1, M/G/1 queue delay fns, etc.

When does the price of selfishness have a succinct explanation?

Example: Braess's Paradox

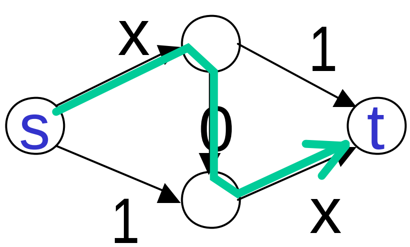

bad Nash flow

vs.

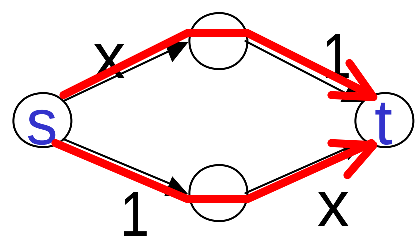

good Nash flow

Question: does this example explain Braess's Paradox?

- yes w/linear latency fns [RT 00]
- no otherwise (more complicated examples can be more severe) [R 01]