

Selfish Routing

Tim Roughgarden
Cornell University


Includes joint work with Éva Tardos

Traffic in Congested Networks

The Model:

- A directed graph $G = (V, E)$
- k source-destination pairs $(s_1, t_1), \dots, (s_k, t_k)$
- A rate r_i of traffic from s_i to t_i
- For each edge e , a latency function $l_e(\cdot)$

Example: $(k, r=1)$




Flows and their Cost

Traffic and Flows:

- f_p = amount of traffic routed on s_i - t_i path P
- flow vector $f \Leftrightarrow$ routing of traffic

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 :
$$\sum_p f_p \cdot l_p(f)$$




Selfish Routing

- 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



this flow is envious!

Assumption: edge latency functions are continuous, nondecreasing


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

Fact: have existence, uniqueness

The Inefficiency of Nash Flows

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

P lack of coordination leads to inefficiency




Cost of **Nash** flow = $1 \cdot 1 + 0 \cdot 1 = 1$

Cost of **optimal (min-cost)** flow
= $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = \frac{3}{4}$

How Bad is Selfish Routing?

Pigou's example is simple...



How inefficient are Nash flows:


- with more realistic **latency fns**?
- in more realistic **networks**?

Goal: prove that Nash flows are **near-optimal**

- want a *laissez-faire* approach to managing networks
 - also [Koutsoupias/Papadimitriou 99]

The Bad News

Bad Example: ($r = 1$, d large)



Nash flow has cost 1, min cost ≈ 0

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

- even if latency functions are polynomials

A Bicriteria Bound

Approach #1: settle for weaker type of guarantee

Theorem: [Roughgarden/Tardos 00]
network w/cts, nondecreasing latency functions \mathbb{P}

cost of **Nash** at rate r = cost of **opt** at rate $2r$

Corollary: M/M/1 delay fns
($l(x) = 1/(u-x)$, u = capacity) \mathbb{P}

Nash cost w/ capacities $2u$ = **opt** cost w/ capacities u

Linear Latency Functions

Approach #2: restrict class of allowable latency functions

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


Theorem: [Roughgarden/Tardos 00]
network w/linear latency fns \mathbb{P}

$$\text{cost of Nash flow} = \frac{4}{3} \times \text{cost of opt flow}$$

Sources of Inefficiency

Corollary of main Theorem:

- For linear latency fns, worst Nash/OPT ratio is realized in a two-link network!




- Cost of **Nash** = 1
- Cost of **OPT** = $\frac{3}{4}$

- one source of inefficiency:
 - confronted w/two routes, selfish users overcongest one of them
- Corollary \Rightarrow **that's all, folks!**
 - network topology plays no role

No Dependence on Network Topology

Theorem: [Roughgarden 02] for (almost) **any** class of latency fns including the constant fns, worst Nash/OPT ratio occurs in a two-link network.

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



Coping with Selfishness

Motivation:

- Nash flows inefficient
- centralized routing often infeasible


Goal: design/manage networks s.t. selfish routing “not too bad”
⇒ adds new algorithmic dimension

Two Approaches:


- network design (next)
- Stackelberg routing (see thesis)

Braess's Paradox

Better network, worse Nash flow:



Cost of **Nash flow** = 1.5




Cost of **Nash flow** = 2

All traffic experiences additional latency! [Braess 68]

Designing Networks for Selfish Users

The Problem:

- given network $G = (V, E, I)$
 - assume single-commodity
- find subnetwork minimizing latency experienced by all selfish users in a Nash flow




\Rightarrow want to avoid Braess's Paradox

Generalizing Braess's Paradox

Question: is Braess's Paradox more severe in bigger networks?

Fact: with linear latency fns, worst case is




Reason: with linear latency fns,


average latency of Nash flow = $4/3 \times$ average latency of any other flow

Braess's Paradox with General Latency Fns

A Bigger Braess Paradox:



Nash in whole graph
common latency = 4



Nash in opt subgraph
common latency = 1

\Rightarrow removing edges can improve
Nash by a $n/2$ factor ($n=|V|$)

Thm: [R 01] this is worst possible.

The Trivial Algorithm

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

We know: the trivial algorithm is

- a $4/3$ -approx alg with linear latency fns
- an $n/2$ -approx alg with general latency fns

Question: what about more sophisticated algorithms?

Designing Networks for Selfish Users is Hard

Thm: [R 01] For $\epsilon > 0$, no $(n/2 - \epsilon)$ -approximation algorithm exists (unless $P=NP$).

Thm: [R 01] For linear latency functions, no $(4/3 - \epsilon)$ -approx algorithm exists (unless $P=NP$).

Remark: similar results hold for other classes of latency fns.

Corollary: Braess's Paradox eludes efficient algorithms.

Directions for Further Research

Selfish Routing: many open questions, see thesis

Other Games: e.g., flow control, competitive facility location, auctions

Paradigm for studying selfishness:

- what is worst Nash/OPT objective fn value ratio?
- are other meaningful bounds (e.g., bicriteria) possible?
- sources of inefficiency?
- design/management strategies for coping with selfishness?