A New Look at Selfish Routing

Gregory Valiant

(joint work with Christos Papadimitriou)

What is the Internet?

- Composed of >10,000 smaller interconnected networks (Autonomous Systems)
  - Range in size from servicing single universities/businesses to multinational (operated by companies or governments)
- E.g Sprint, Qwest, MCI, Level 3, AT&T, Cogent, GBLX, Telianet….

ASes

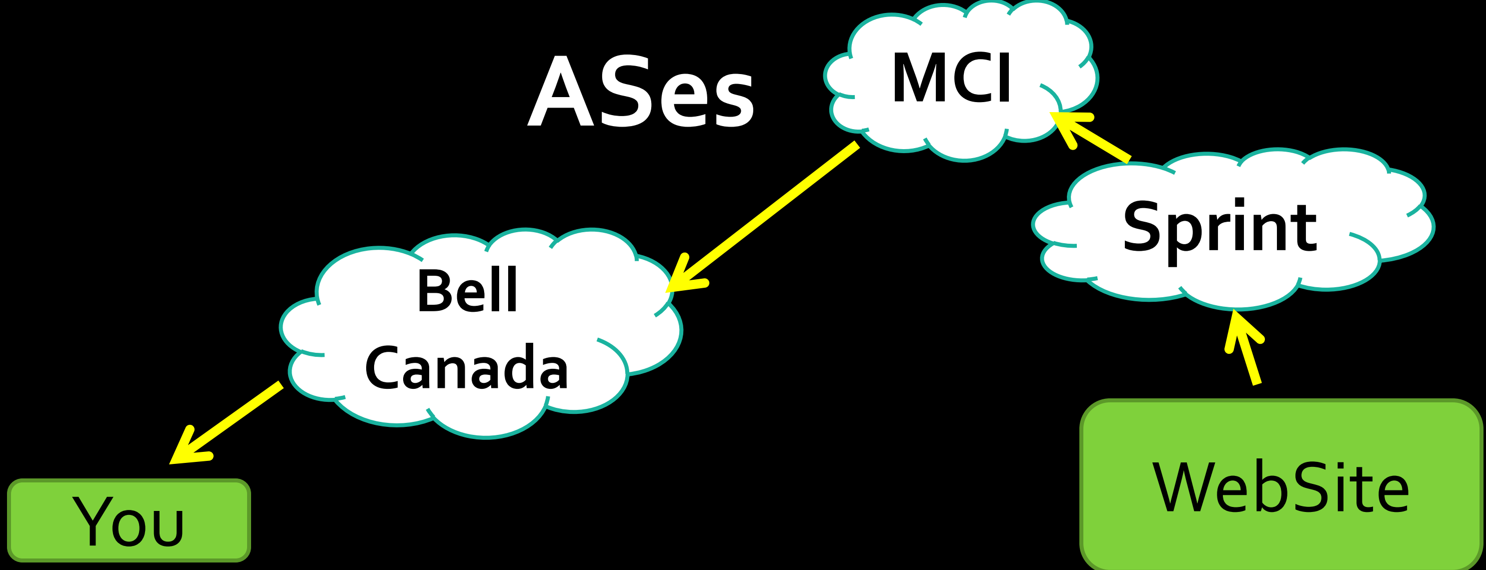

Complex economic relationships enable connectivity:

- Selfishly motivated
- Organically grown
- Decentralized and Unregulated

But: VERY SUCCESSFUL

low-latency and reliable

Question: Why does it work?

Is it surprising?

Economists
"no, competition begets efficiency"

- Theorems of Welfare Economics
- Bertrand and Cournot Competition
- Many others….

Computer Scientists
"yes, selfishness and decentralization lead to inefficiency"

- Prisoner's Dilemma/ inefficient Nash-eq.
- Distributed Systems

Economists
"no, competition begets efficiency"

Computer Scientists
"yes, selfishness and decentralization lead to inefficiency"

different notions of efficiency

- Market clears
- Pareto Efficient (Pareto Optimal)

very weak

Efficiency w.r.t global objective functions e.g.

- Computational eff.
- Social welfare

Formulating the Problem

Model the economic incentives of the agents.

Ask:

Is routing efficiently a plausible outcome?

Nash Equilibria

Price of Anarchy

$$\frac{\text{Cost of Nash}}{\text{Cost of Opt}}$$

Selfish Routing (Roughgarden/Tardos)

Similar spirit to our question

Some positive results: Price of Anarchy is bounded, for many congestion functions (no matter how ugly the network topology)

But model assumes that:

Flows make routing decisions

This is the basic assumption that we will change.

Our Models

As in Selfish Routing model:

Given a network, traffic travelling between s,t pairs, nondecreasing internal edge latencies.

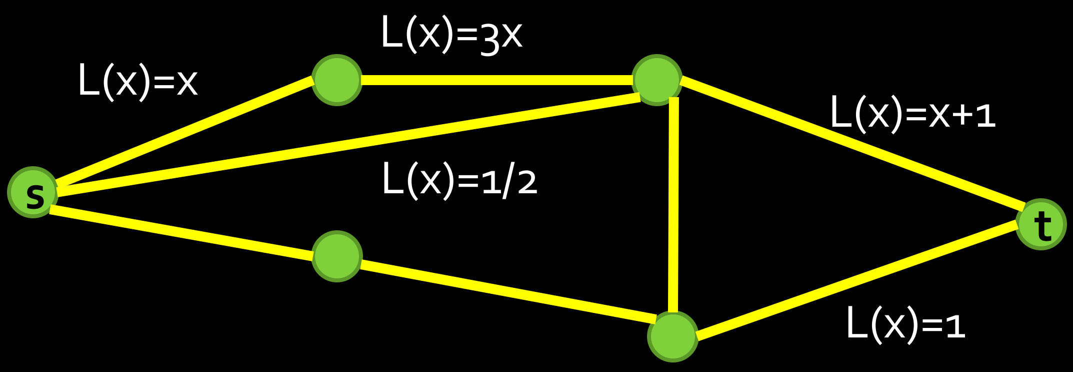

As in the Internet:

•**Network components make routing decisions**

Model economic incentives of ASes

What does an AS want?

Our Models

What does an AS want?

Answer 1: To route traffic really well (Latency Model)

•Each AS routes so as to minimize latency experienced by traffic passing through it.

Answer 2: $$$ (Pricing Model)

•Each AS advertises prices to neighbors, goal is to maximize profit.

Simplifying Assumptions: single source/sink, fixed topology, ‘network components’ = EDGES.

Latency Model

Edge *e*: Route so as to minimize latency experienced by traffic from *e* to *t*,
Assuming downstream edges continue to route as they have been routing.

Prop: Nash equilibria always exist

Main Theorem: Price of Anarchy/Stability unbounded!!! (even with internal latencies $L(x)=ax$)

Pf idea: Unlike in selfish routing,
`bad' networks are independent of traffic rate
=> a recursive construction amplifying badness.

Pricing Model

Each edge has its internal latency (per unit cost of processing, as a function of traffic)

Each edge makes the following decisions:

- price to advertise to upstream neighbors
- Routing of flow to downstream neighbors

Utility = money collected - money paid – internal cost

E.g. 1-unit total flow,

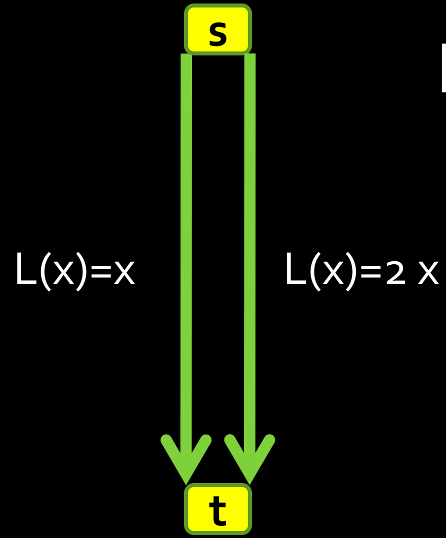

Both edges advertise price of 1,
many ways to split flow that are at eq.

too many equilibria!!!

Prices vs Pricing Schemes

Inefficiency stems from lack of expressiveness:

"I want to route a little bit of traffic at moderate price"

Modification: Allow players to advertise pricing schemes: cost(x)= ax + b

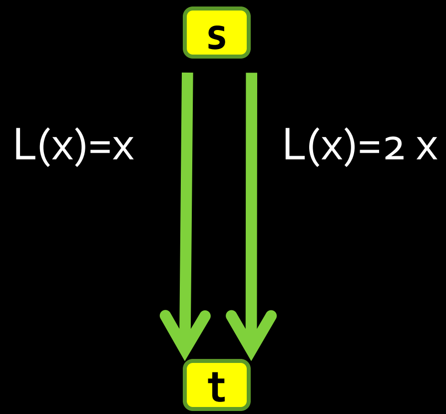

(or any nondec. scheme)

unique equilibrium will be 2/3, 1/3 split, each will get paid 4/3 per unit.

Monopolies [are Bad]

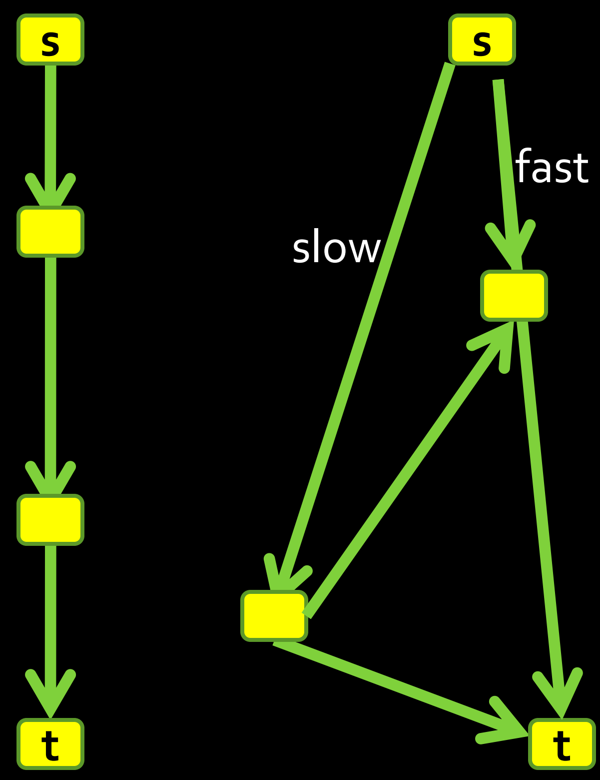

Monopolies => inefficiency (prevent info. propagation).

Def Monopoly Free:
In optimal routing, flow is always split.
(ie. no edge has an *effective* monopoly on upstream edge, at opt)

Main Theorem

Assume:

•'monopoly free' [at opt, players split flow]

•Internal edge latencies are of the form

$$L(x)=ax \quad (a>0)$$

Theorem:

Unique N. E. with optimal routing

(we believe conditions on latency can be weakened)

Proof Sketch

1. At equilibrium, all `competing players' advertise constant pricing schemes (equal to marginal cost)

2. By exaggerating true cost, cannot go from a competing player to a noncompeting player.

Proof Sketch: Part 1

At equilibrium, all `competing players' advertise constant pricing schemes

Observation 1: Get more flow at same price by 'flattening' pricing scheme

Observation 2: Raise price to get original flow at higher price

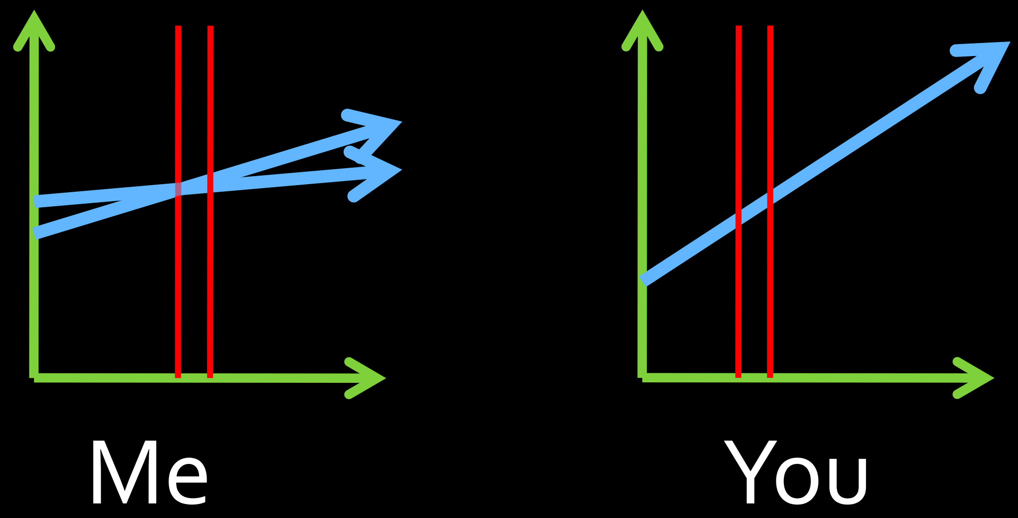

Proof Sketch: Part 2

By exaggerating true marginal cost, cannot go from ‘competing’ to ‘noncompeting’.

This implies that no ‘bad’ equilibria exist

Intuition: If I increase my price, that can only hurt me, ie more flow will go to competitors.

But what if a few of us all work together?

This is the heart of the proof, uses tools from circuit analysis

Lessons

- Prices bring efficiency in subtle way
- Preventing monopolistic situations (double sourcing) essential for efficiency
- Short-term competition between routing agents, informed by congestion, is crucial

Open Directions

Proof approach seems robust:

- Extend results to more general latencies
- Multiple source/sinks

Connection with Economic ideas:

- Trading networks
- Information propagation

Thanks