Network games

Brighten Godfrey
CS 538 October 11 2012

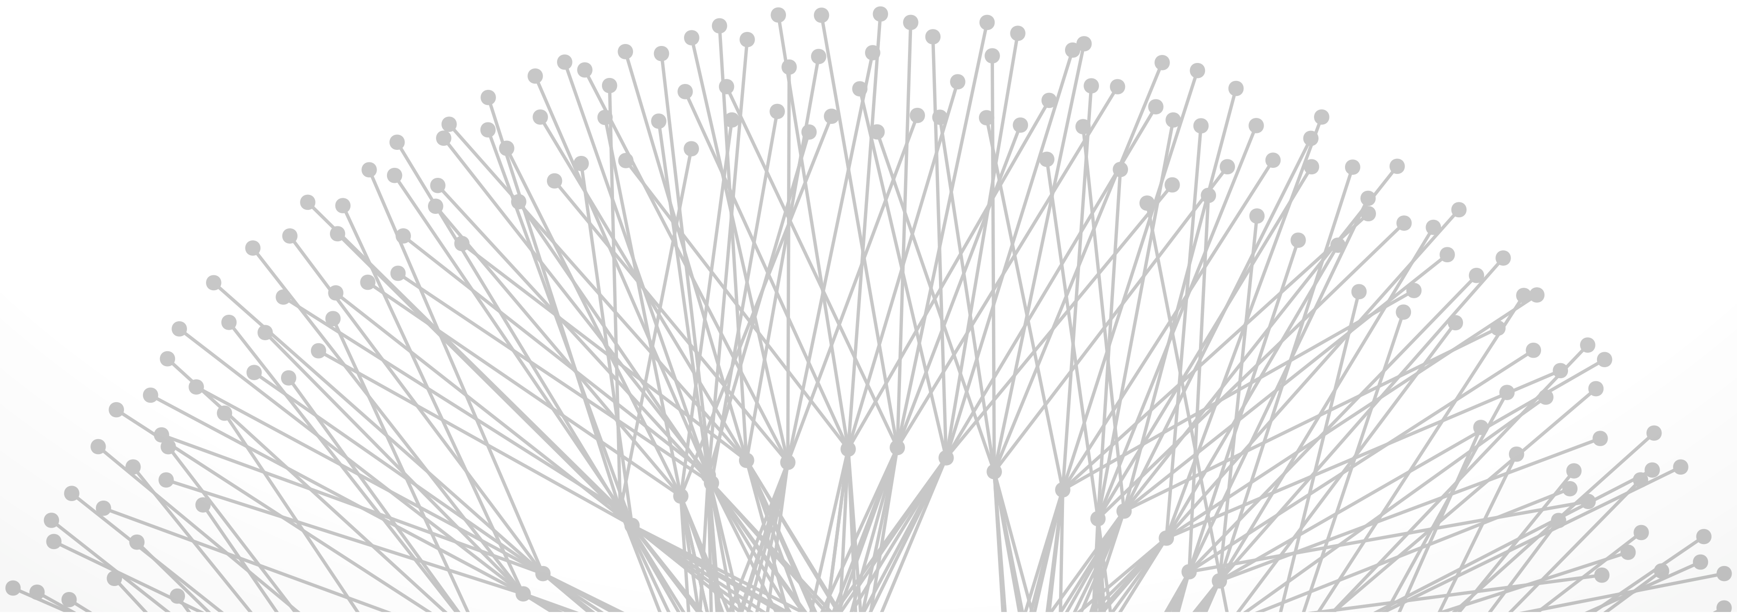

slides ©2010-2012 by Brighten Godfrey unless otherwise noted

Demo

Game theory basics

Games & networks: a natural fit

Game theory
Studies strategic interactions between selfish agents

Networking
Enables interaction between remote agents

Networks make games happen!

Game theory

Two or more players

For each player, a set of strategies

For each combination of played strategies, a payoff or utility for each player

		Blue player strategies		
		Rock	Paper	Scissors
Red player strategies	Rock	\$0, \$0	\$0, \$1	\$1, \$0
	Paper	\$1, \$0	\$0, \$0	\$0, \$1
	Scissors	\$0, \$1	\$1, \$0	\$0, \$0

(Pure) Nash equilibrium

A chosen strategy for each player such that no player can improve its utility by changing its strategy

- (In mixed Nash equilibrium: players randomize their strategies according to some distribution and no player can improve its *expected* utility)

Can you find a Nash equilibrium in R-P-S?

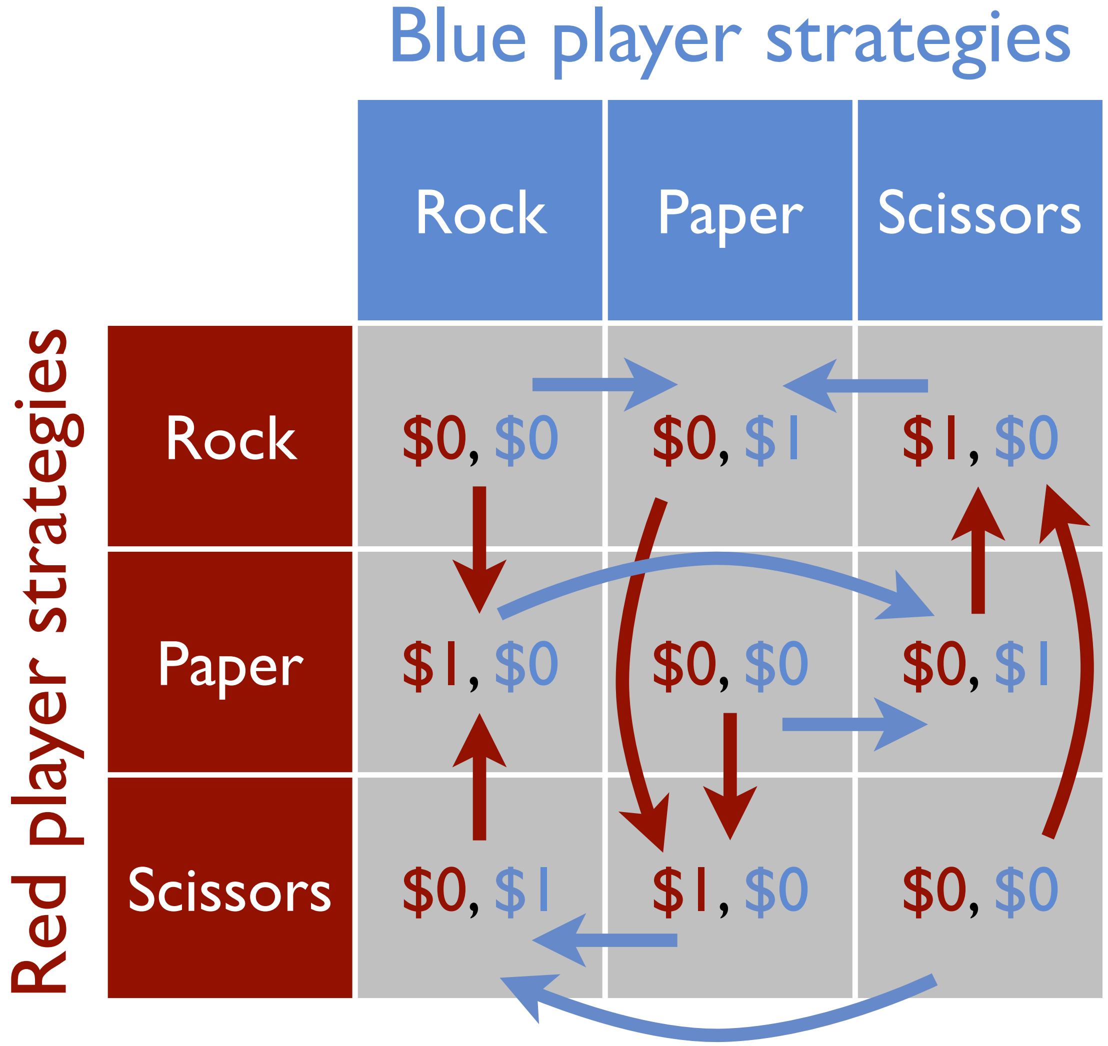

Blue player strategies

Red player strategies

	Rock	Paper	Scissors
Rock	\$0, \$0	\$0, \$1	\$1, \$0
Paper	\$1, \$0	\$0, \$0	\$0, \$1
Scissors	\$0, \$1	\$1, \$0	\$0, \$0

No pure Nash equilibrium!

Prisoner's dilemma

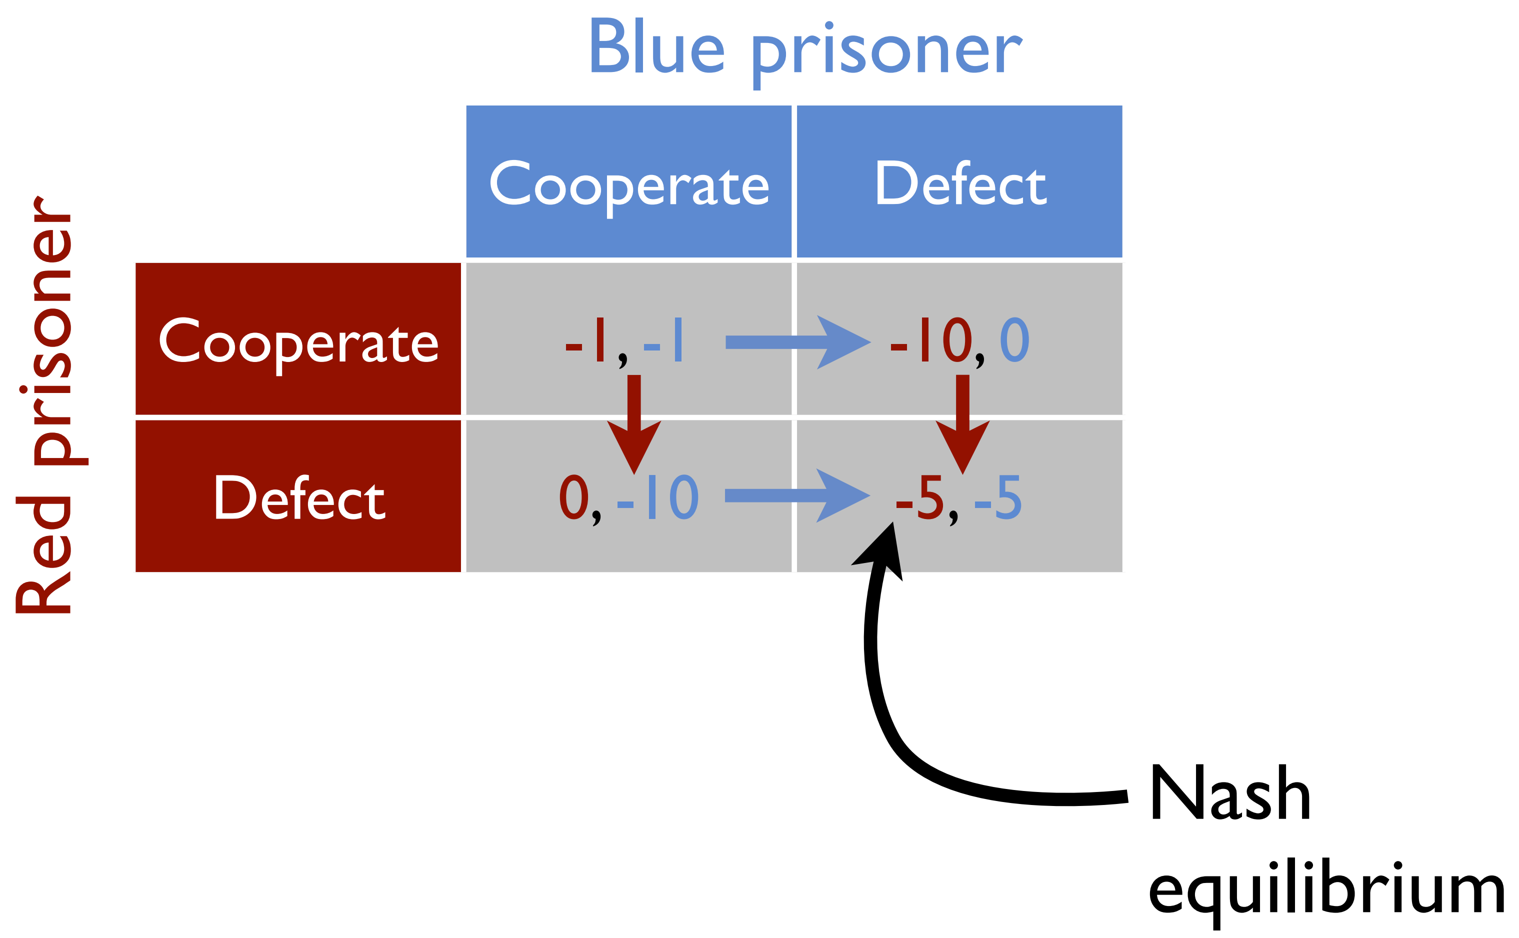

		Blue prisoner	
		Cooperate	Defect
Red prisoner	Cooperate	-1, -1	-10, 0
	Defect	0, -10	-5, -5

Nash equilibrium

Price of Anarchy

[C. Papadimitriou, "Algorithms, games and the Internet", STOC 2001]

Assumes some global "cost" objective, e.g., social utility (sum of players' payoffs).

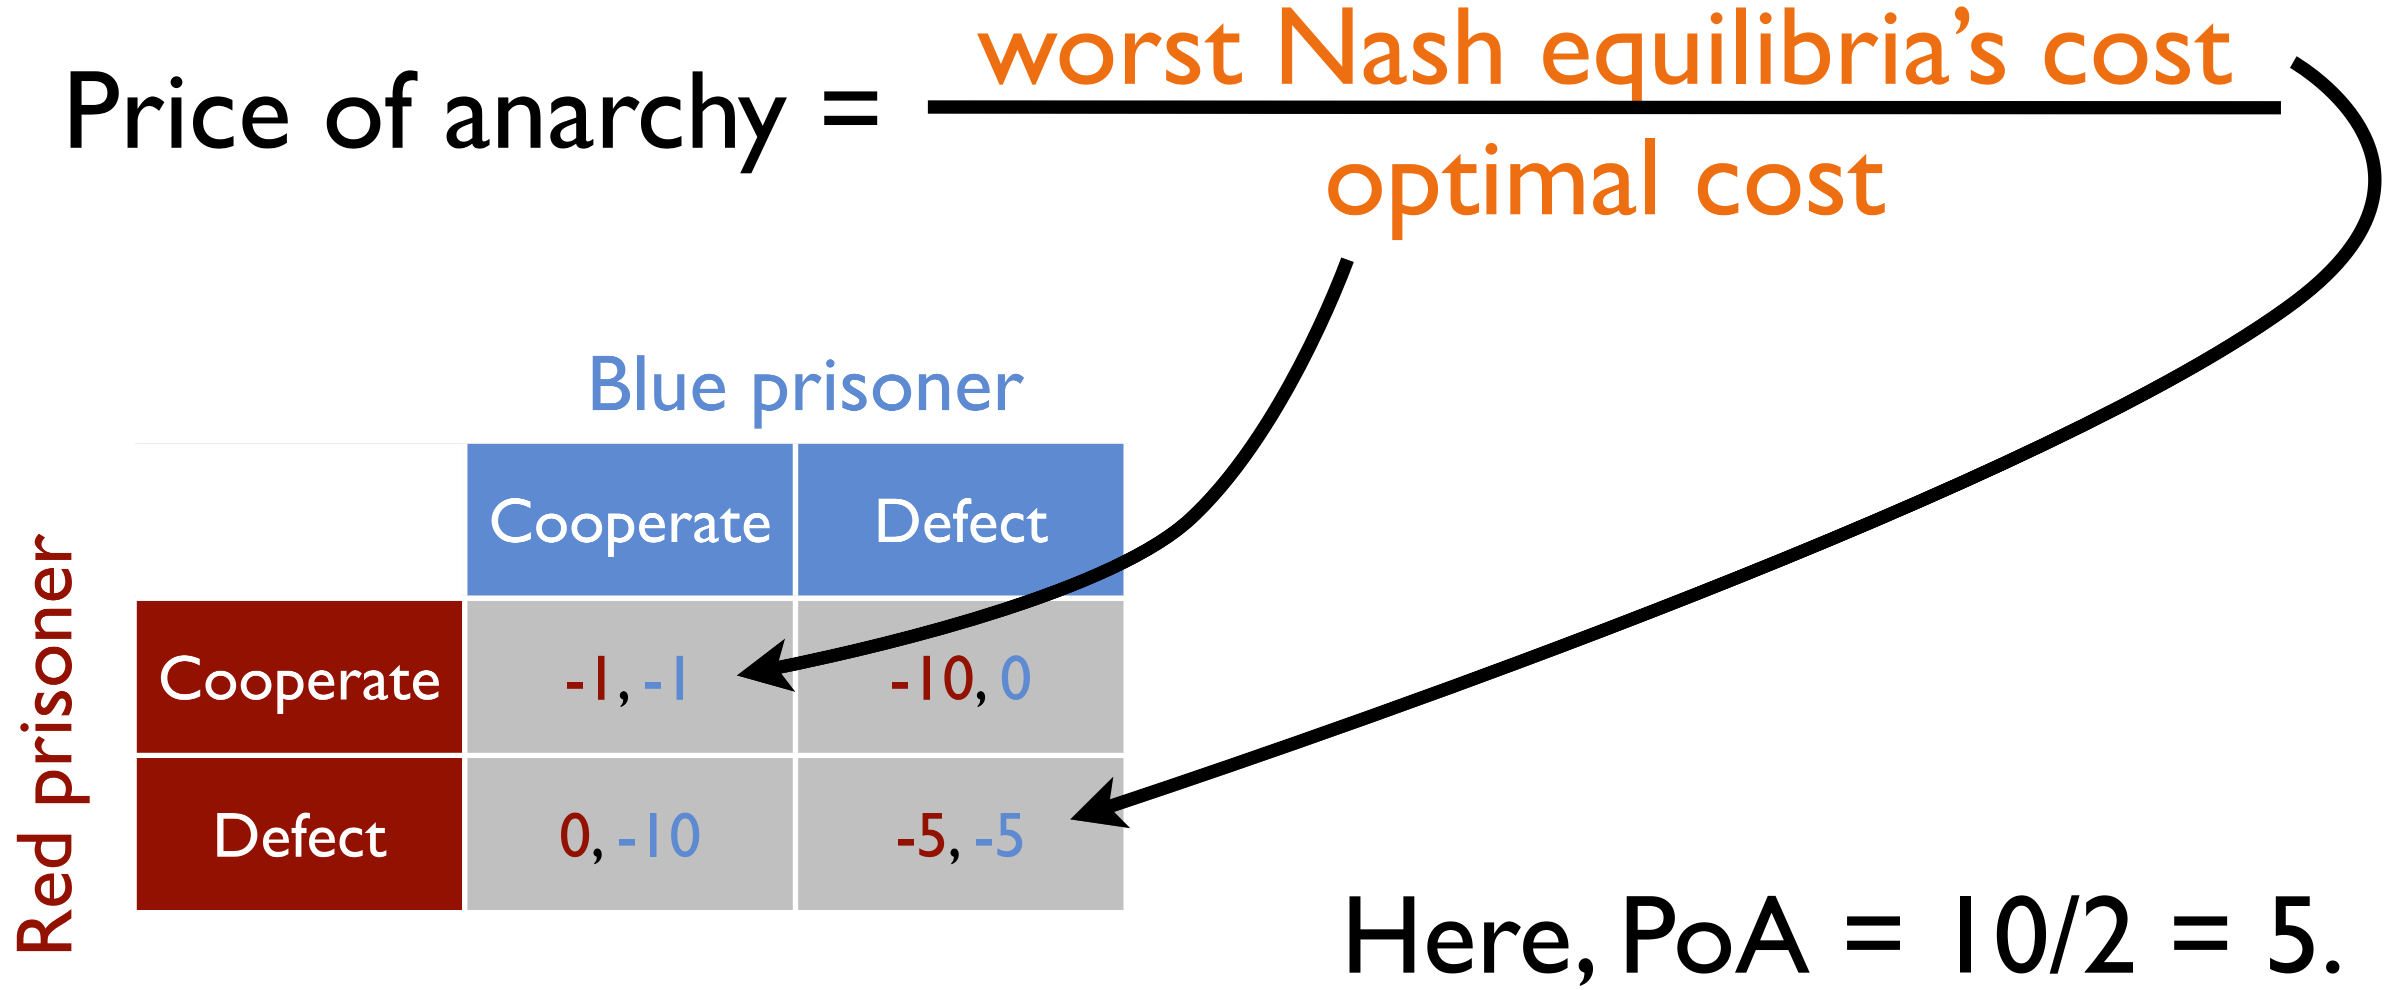

$$\text{Price of anarchy} = \frac{\text{worst Nash equilibria's cost}}{\text{optimal cost}}$$

		Blue prisoner	
		Cooperate	Defect
Red prisoner	Cooperate	-1, -1	-10, 0
	Defect	0, -10	-5, -5

Here, PoA = 10/2 = 5.

Today's papers

Stable paths problem

- [Tim Griffin, Bruce Shepherd, Gordon Wilfong, ToN'02]
- A game model of BGP

How bad is selfish routing?

- [Tim Roughgarden, Eva Tardos, JACM 2002]
- Analysis of price of anarchy of latency-optimized routing

Selfish routing in Internet-like environments

- [Lili Qiu, Richard Yang, Yin Zhang, Scott Shenker, SIGCOMM'03]
- What is the price of anarchy like in practice for latency-optimized routing?

Internet routing as a game

BGP routing as a game

players	autonomous systems
strategies	pick a route, any route... (to fixed dest.)
player's utility	arbitrary function of route (but $-\infty$ for 'illegal' route not offered by neighbor)

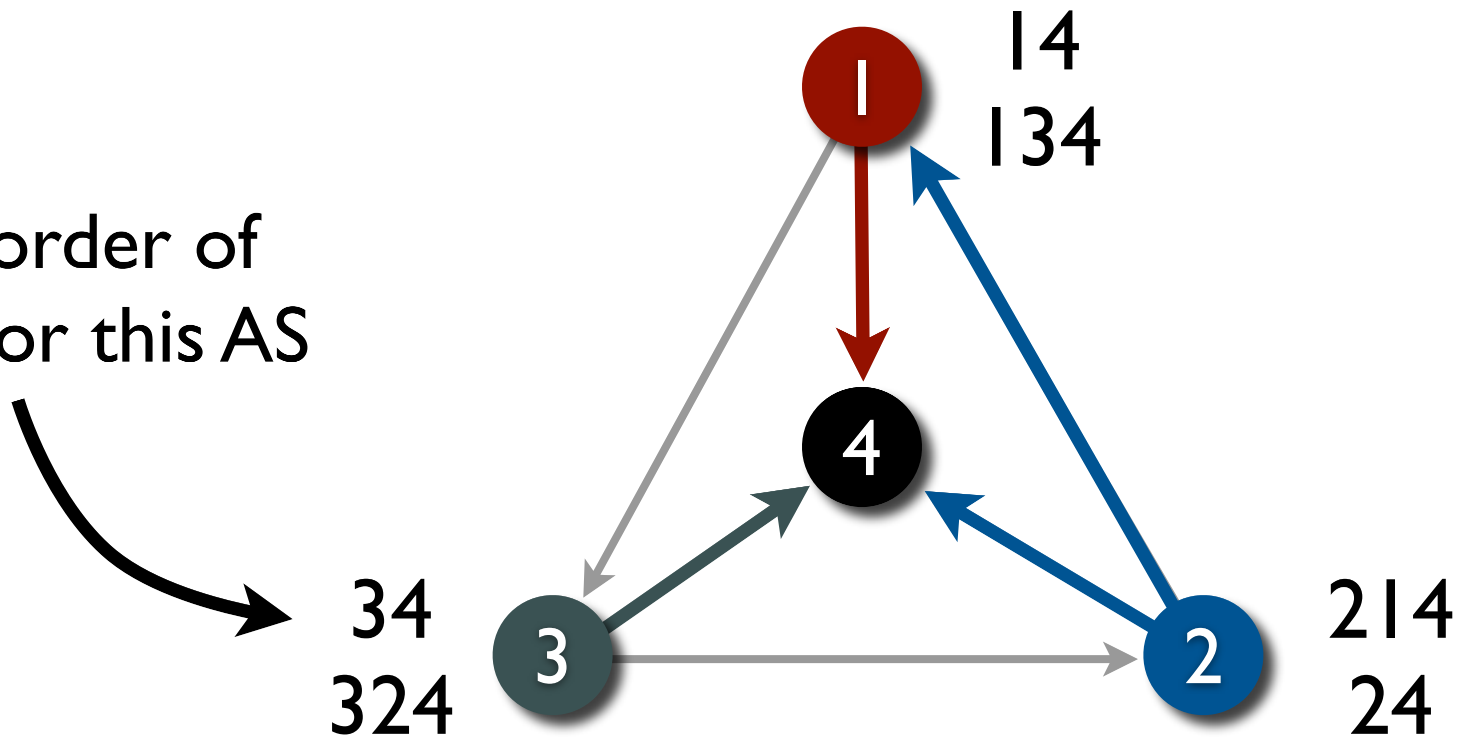
Routes in order of preference for this AS

BGP routing as a game

players	autonomous systems
strategies	pick a route, any route… (to fixed dest.)
player’s utility	arbitrary function of route (but –∞ for ‘illegal’ route not offered by neighbor)

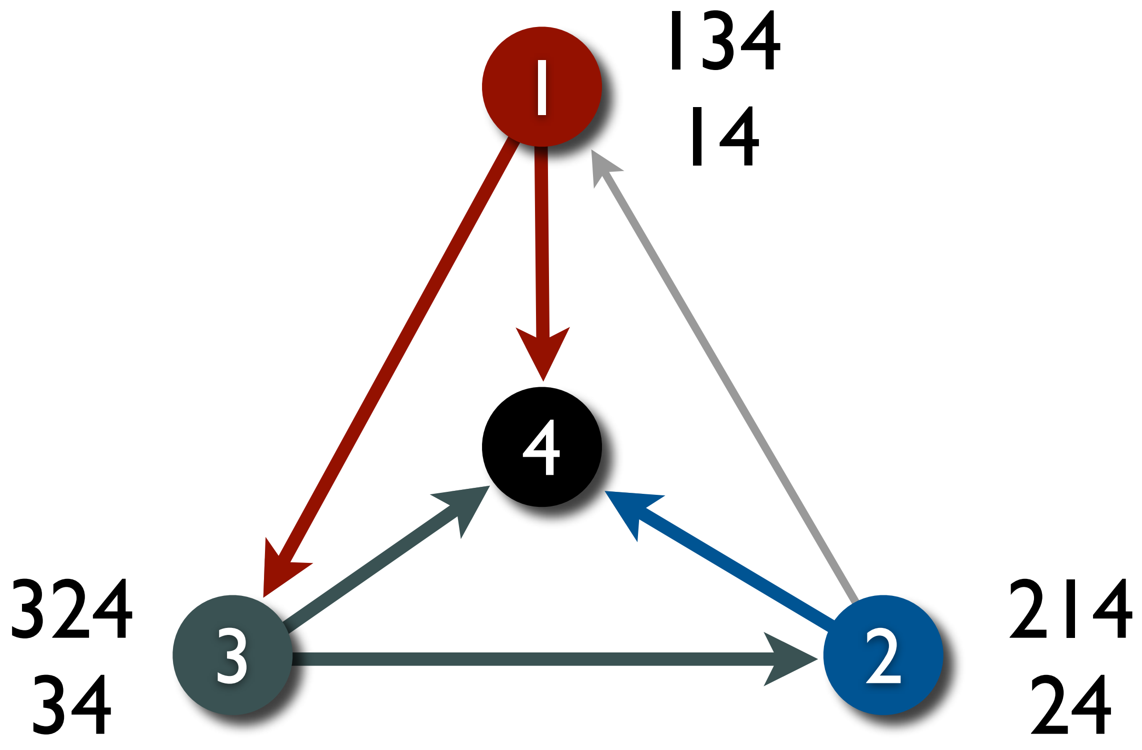

Routes in order of preference for this AS

No Nash equilibrium!

BGP routing as a game

In general, NP-complete to decide whether an equilibrium exists [Griffin, Shepherd, Wilfong, ToN'02]

Might have 0, 1, 2, 3, ... equilibria

Even if it has an equilibrium, might not converge to it

- Depends on starting state, message timing, ...
- PSPACE-complete to decide whether a given set of BGP preferences can oscillate [Fabrikant, Papadimitriou, SODA'08]

If we assume customer-provider-peer and valley-free routing, guaranteed to converge [Gao, Rexford]

How bad is selfish routing?

The selfish routing game

The game context:

- Directed graph
- Latency function on each edge specifying latency as function of total flow *x* on edge
- Path latency = sum of edge latencies

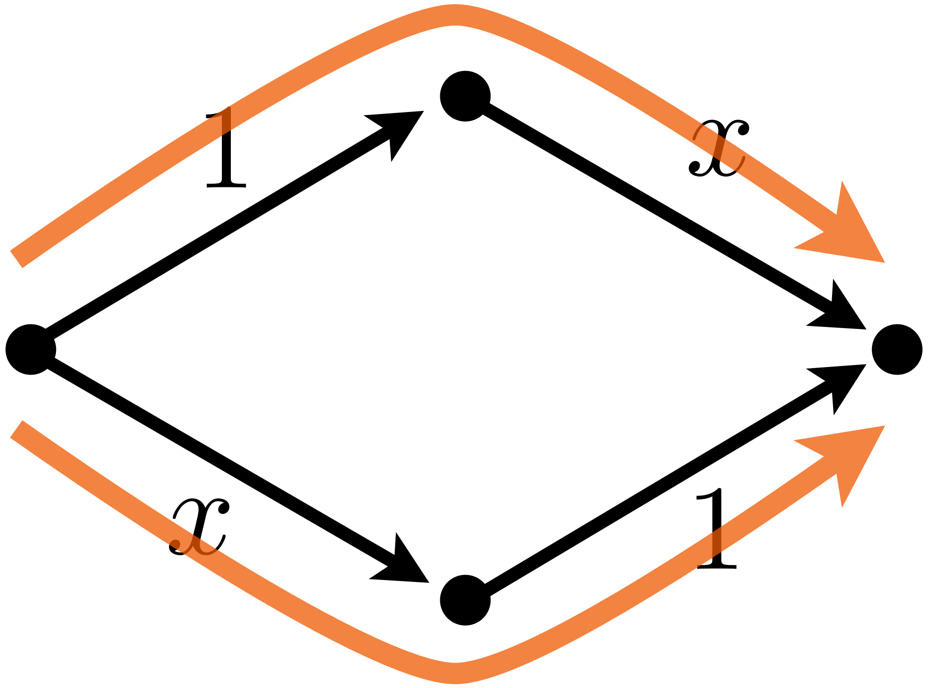

Flow *x* = 0.5 on each path;
Total latency = 1.5

The selfish routing game

Player strategy:

- Pick a path on which to route
- Players selfishly pick paths with lowest latency (source-controlled routing)

For now assume:

- many users
- each has negligible load
- total load = 1

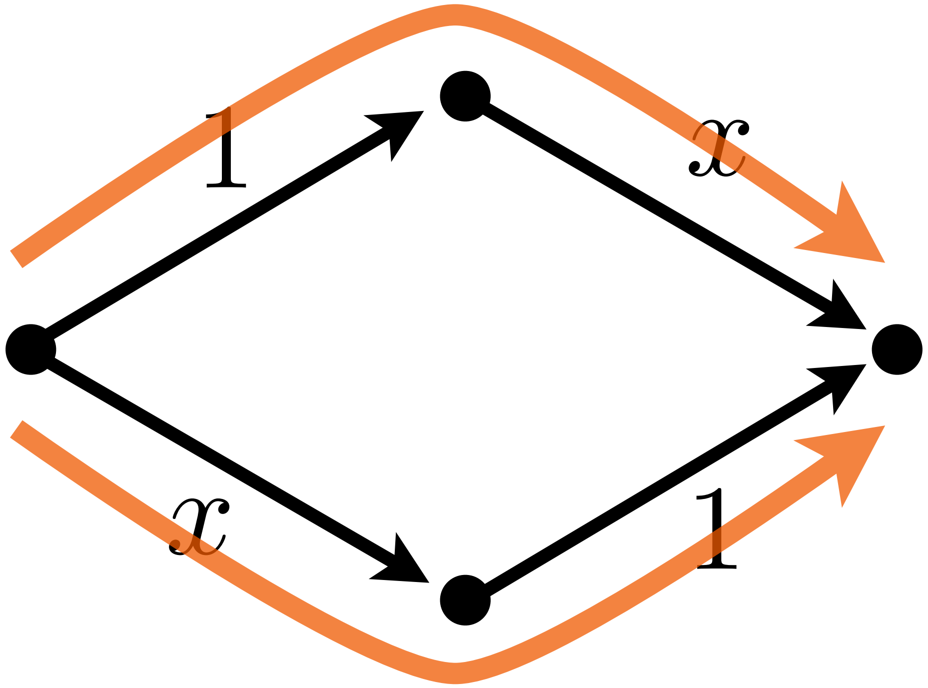

Flow *x* = 0.5 on each path;
Total latency = 1.5

Example: Braess's paradox

[Dietrich Braess, 1968]

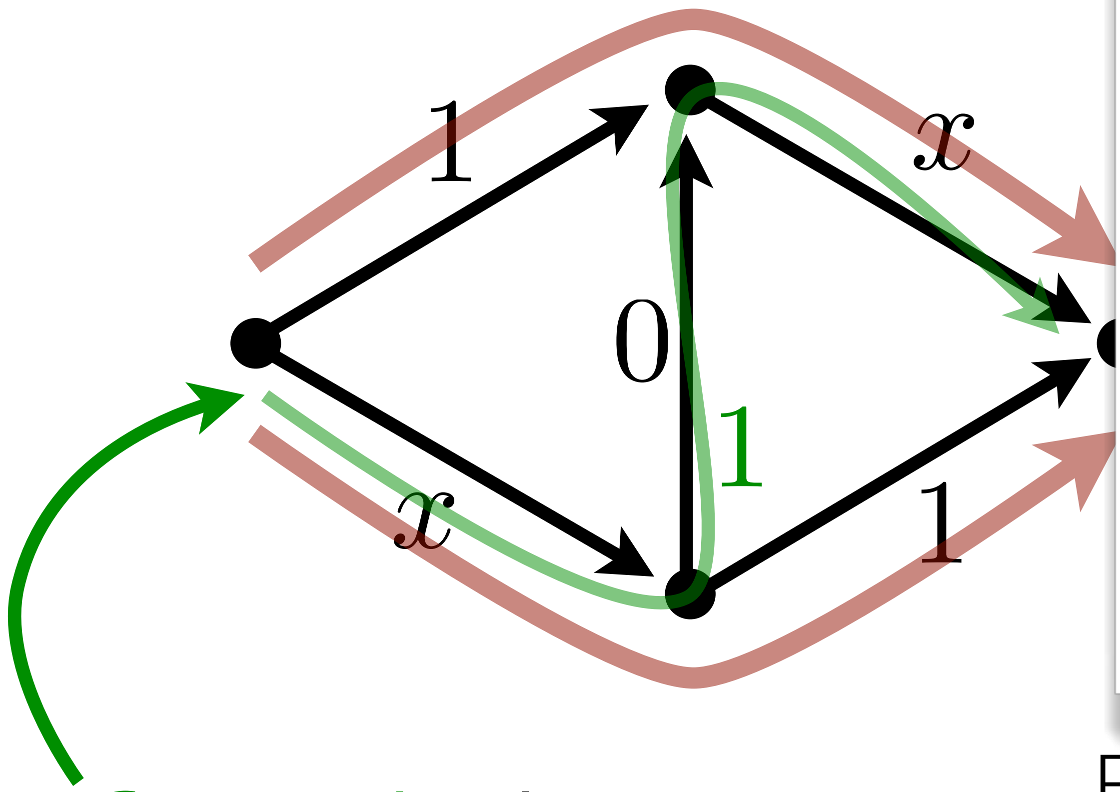

Green path is better.
Everyone switches to it!

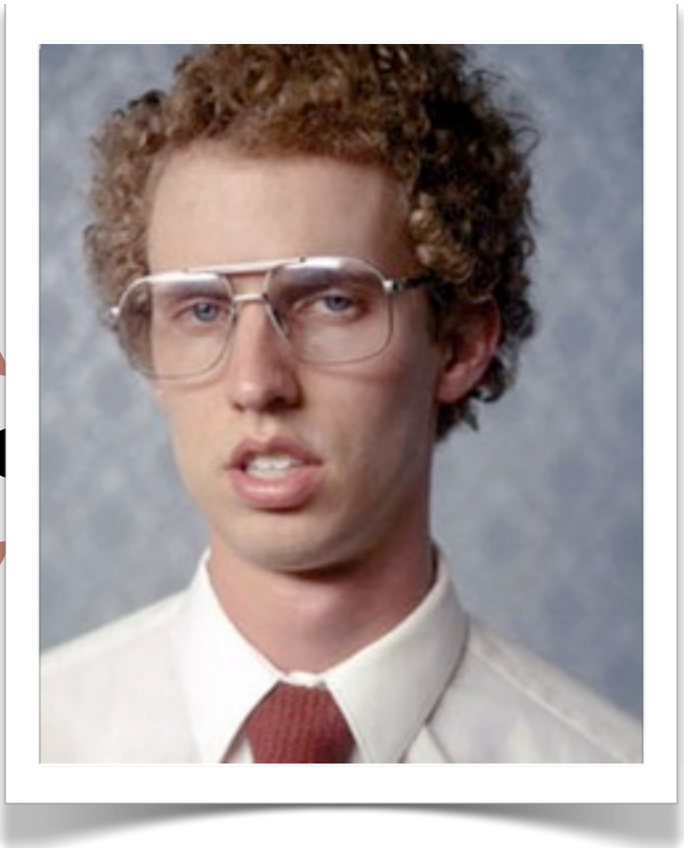

Fig 1b: N. Dynamite.

$\approx$

Fig 1a: D. Braess.

Initially: 0.5 flow along each path; latency 1+0.5 = 1.5

With new edge: all flow along greed path; latency = 2

Example: Braess’s paradox

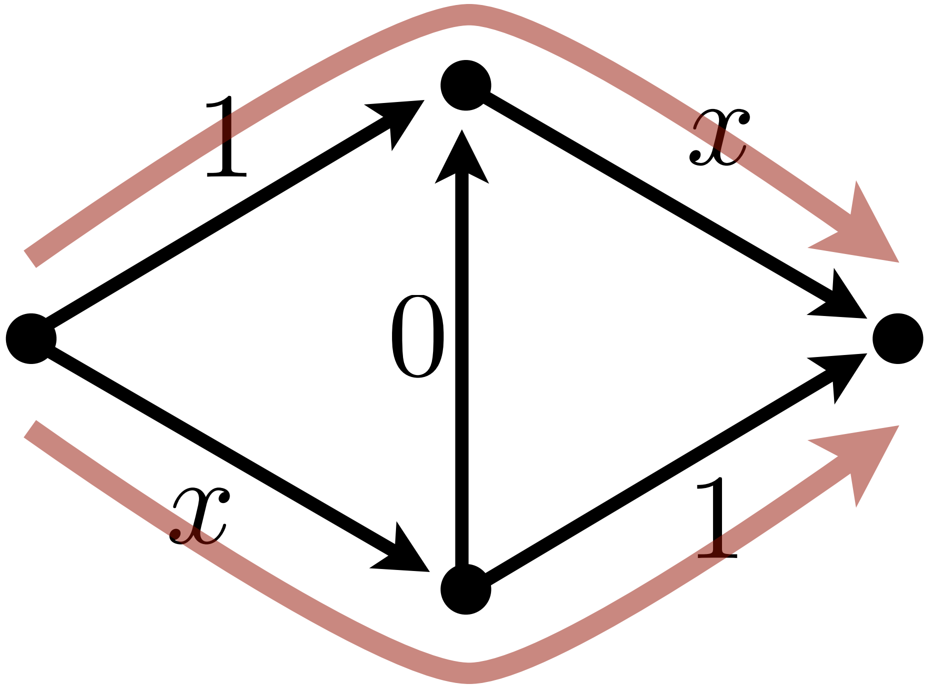

Optimal latency = 1.5

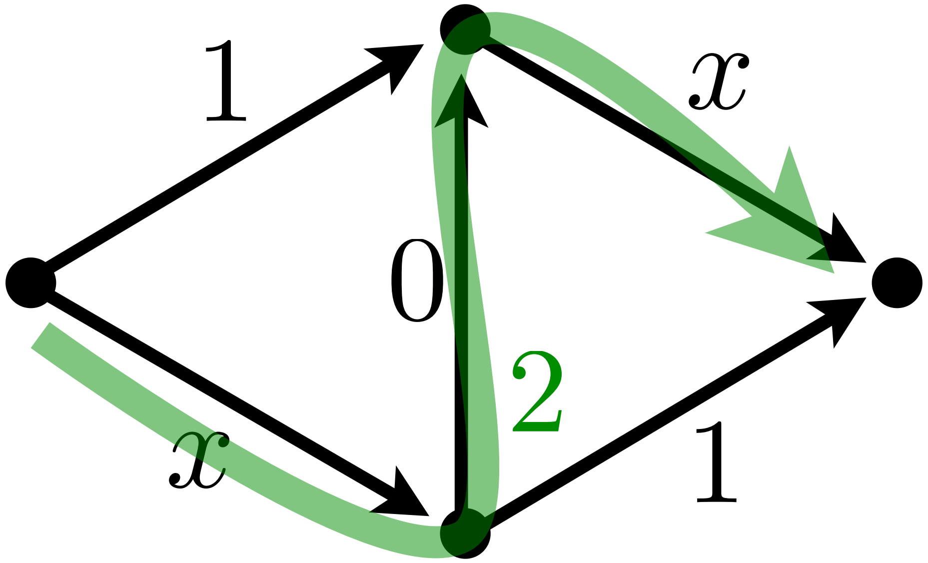

Nash equilibrium latency = 2

Thus, price of anarchy = 4/3

From links to springs

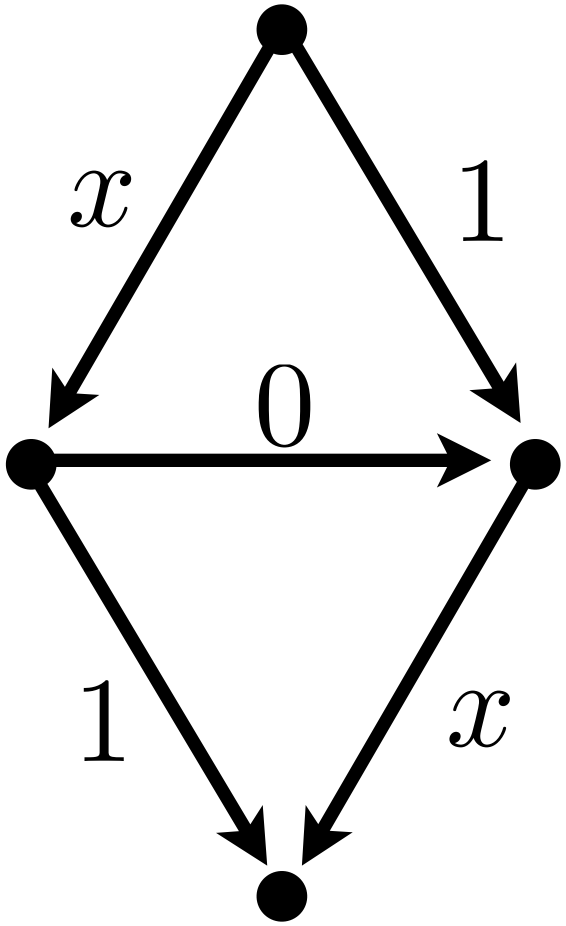

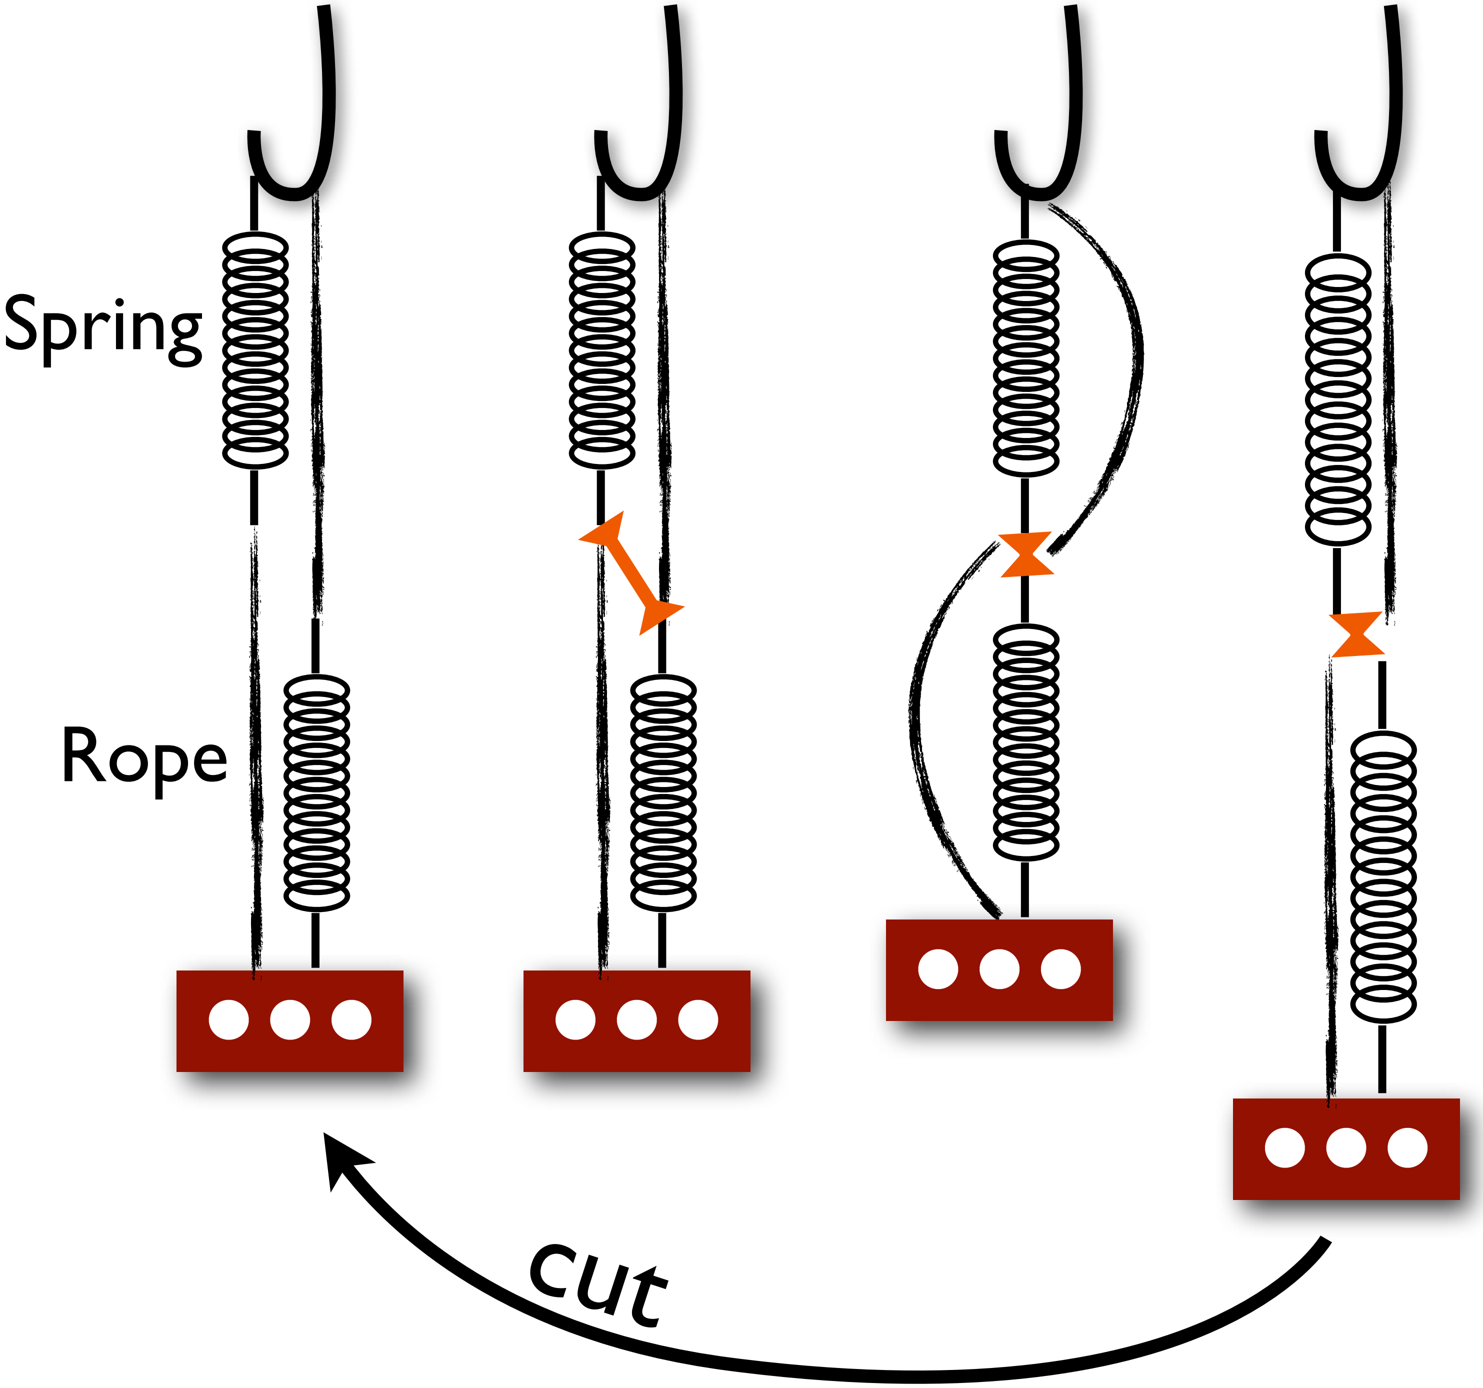

Example: arbitrarily bad

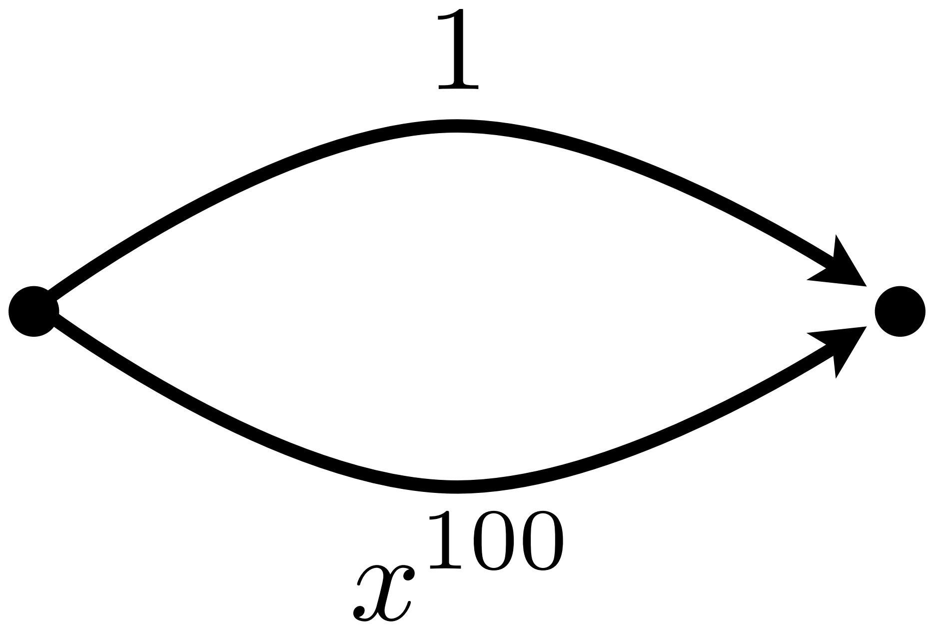

Optimal: almost all flow on bottom; total latency near zero

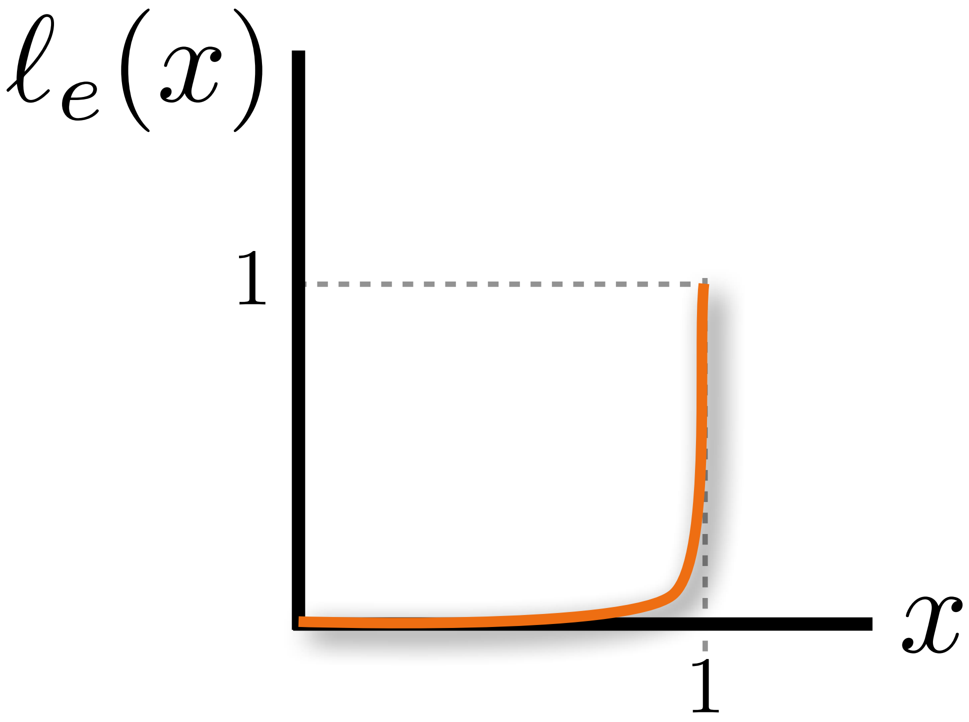

Nash: all flow on bottom; total latency = 1

Roughgarden et al. results

As we just saw, price of anarchy can be arbitrarily high

But for linear latency functions: PoA $\leq$ 4/3

For any latency function: Nash cost is at most optimal cost of 2x as much flow

Extension to finitely many agents

- i.e., a single agent might have a nontrivial fraction of the total bandwidth
- Splittable flow: similar "2x" result
- Unsplittable flow: can be very bad

Selfish routing in realistic networks

[Qiu et al., SIGCOMM 2003]

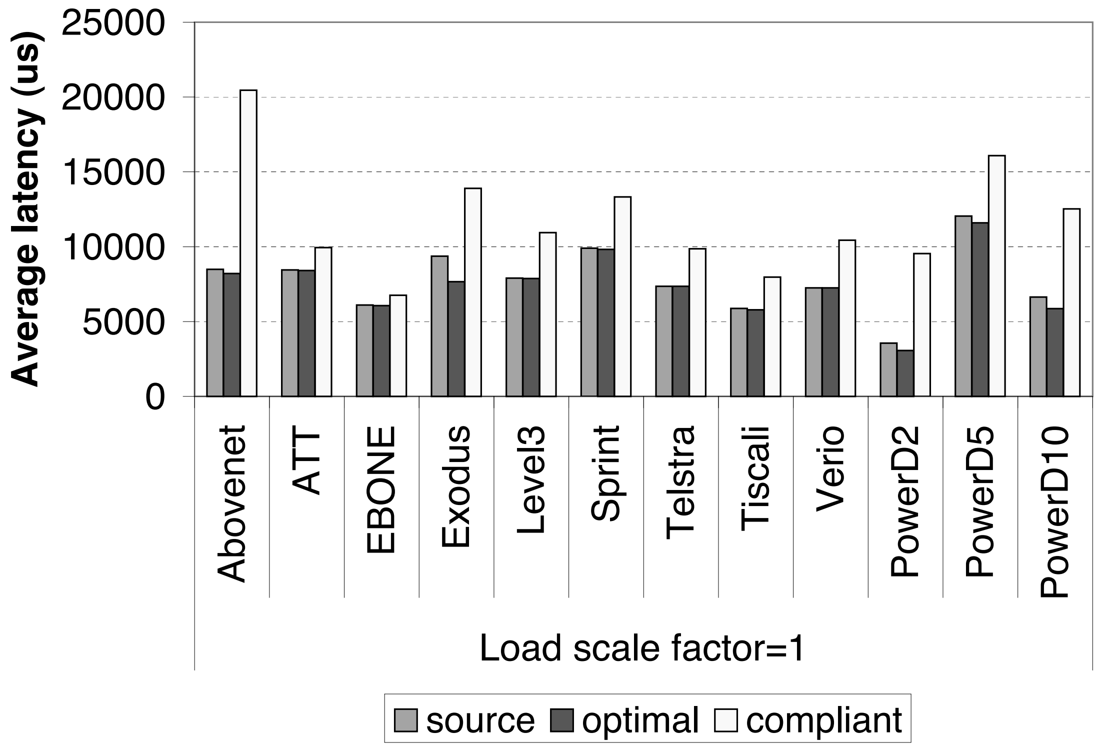

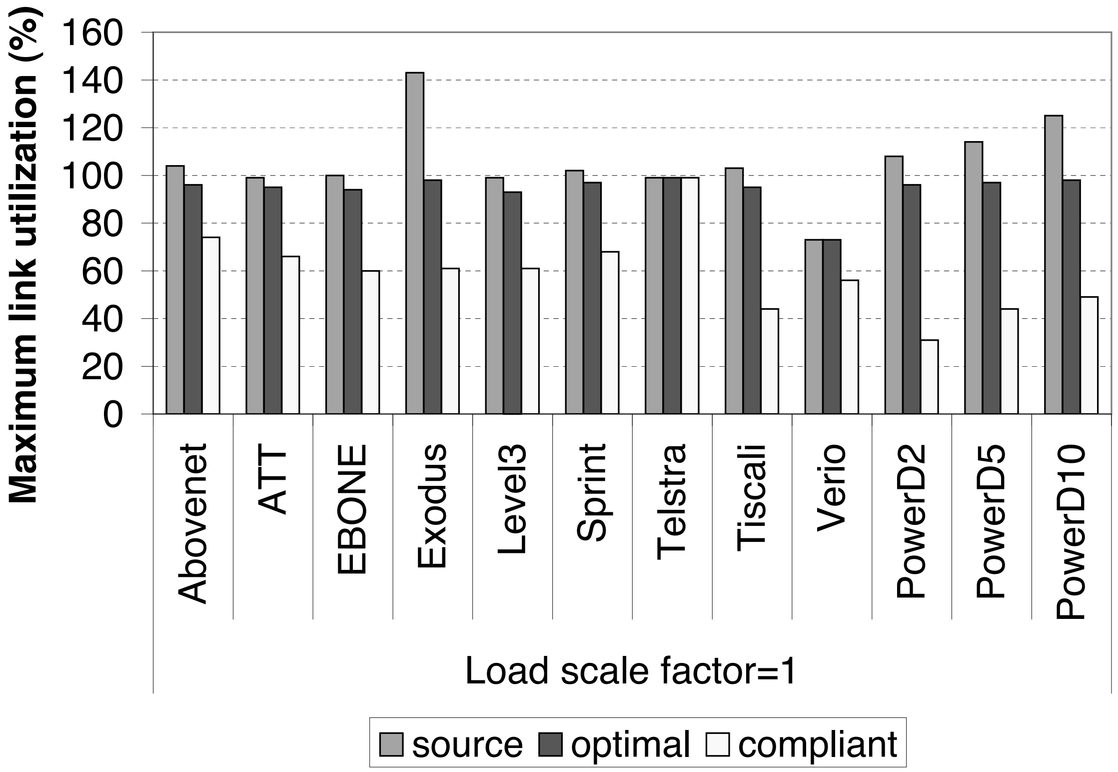

Close to optimal latency

...but higher maximum link utilization

Discussion

Max utilization is higher in selfish. Does it matter?

What happens before we reach equilibrium?

Discussion

Where else do ‘vertical interactions’ between the network and selfish agents appear?

Beyond routing...

Game theory used in networking to model

- Equilibria of distributed algorithms
- ISPs competing with each other
- Spread of new technology in social networks
- ...

Many more applications of game theory to CS

- ...and applications of CS to game theory!
- See Nisan, Roughgarden, Tardos, Vazirani's book Algorithmic Game Theory, available free online

Announcements

Next Tuesday

- Survey of BGP security Issues (Butler, 2010)
- Craig and Uttam present on secure routing

Get started now on your projects!

- Midterm presentations in ~ 3 weeks
- Benchmark: demonstrate concrete progress