The Price of Anarchy:
Some Old and New Results

Tim Roughgarden
Stanford University

Algorithms and Game Theory

Recent Trend: design and analysis of algorithms and systems with self-interested agents

Motivation: the Internet

- auctions (eBay, sponsored search, etc.)
- competition among end users, ISPs, etc.

Traditional approach:

- agents classified as obedient or adversarial
  - examples: distributed algorithms, cryptography

Inefficiency of Equilibria

Obvious fact: many modern applications in CS involve autonomous, self-interested agents

- motivates noncooperative games as modeling tool

Unsurprising fact: equilibria of noncooperative games typically inefficient

- i.e., don't optimize natural objective functions
- e.g., *Nash equilibrium*: an outcome such that no player better off by switching strategies

Price of anarchy: quantify inefficiency w.r.t some objective function.

Performance Guarantees

Good news: in theoretical CS, have lots of techniques for measuring inefficiency.

- motivated by NP-completeness, real-time algorithms, etc.

Definition: approximation ratio (w.r.t. some objective function):

$$\frac{\text{protagonists's obj fn value}}{\text{optimal obj fn value}}$$

the closer to 1 the better

Inefficiency of Nash Flows

Note: selfish routing does not minimize average delay (observed informally by [Pigou 1920])

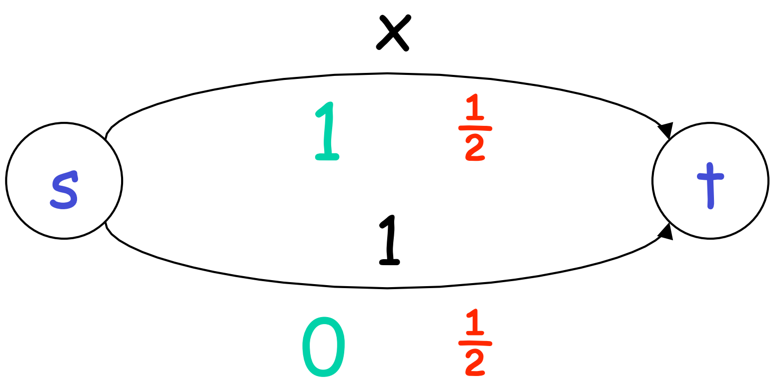

- Cost of equilibrium flow = $1 \cdot 1 + 0 \cdot 1 = 1$
- Cost of optimal (min-cost) flow = $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = \frac{3}{4}$
- *Price of anarchy := equilibrium/OPT ratio = 4/3*

Braess's Paradox

Initial Network:

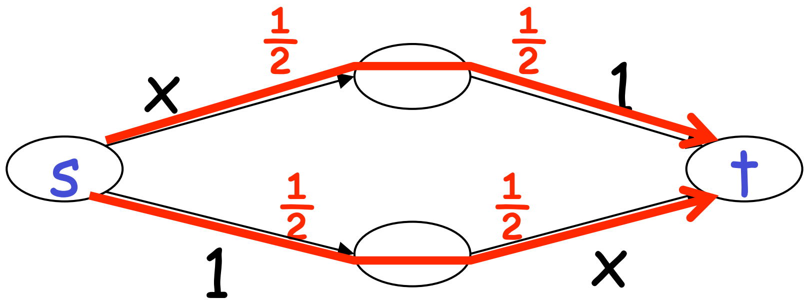

Cost = 1.5

Braess's Paradox

Initial Network:

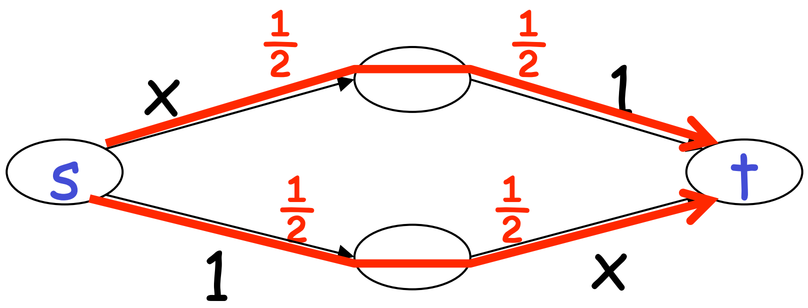

Cost = 1.5

Augmented Network:

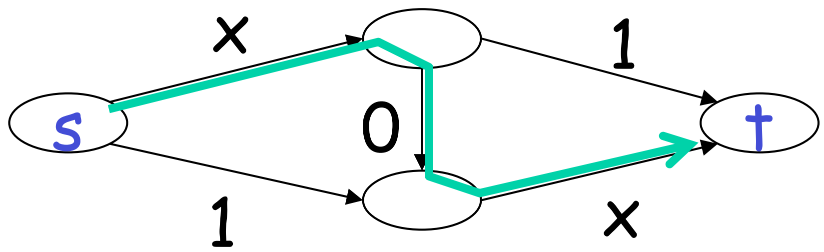

Cost = 2

Braess's Paradox

Initial Network:

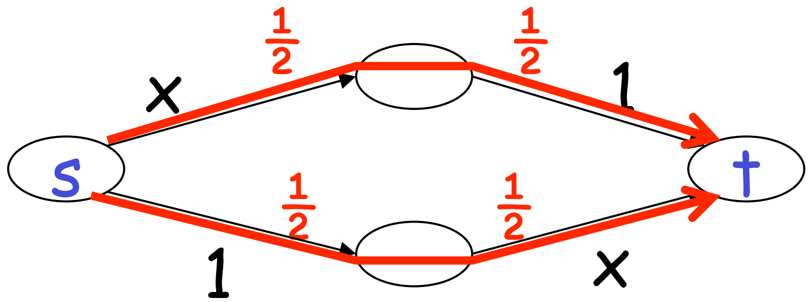

Cost = 1.5

Augmented Network:

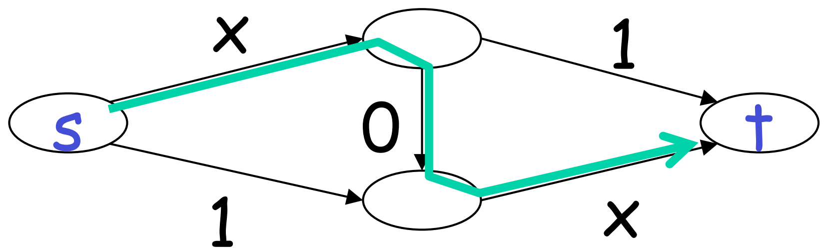

Cost = 2

All traffic incurs more cost! [Braess 68]

- also has physical analogs [Cohen/Horowitz 91]

Unbounded Inefficiency

Example: large prop delay + small queuing delay vs. small prop delay + large queuing delay

- one unit (comprising many flows) selfish traffic

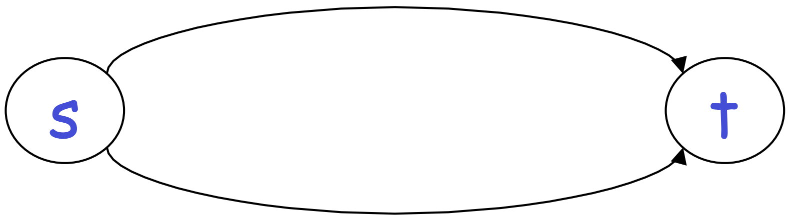

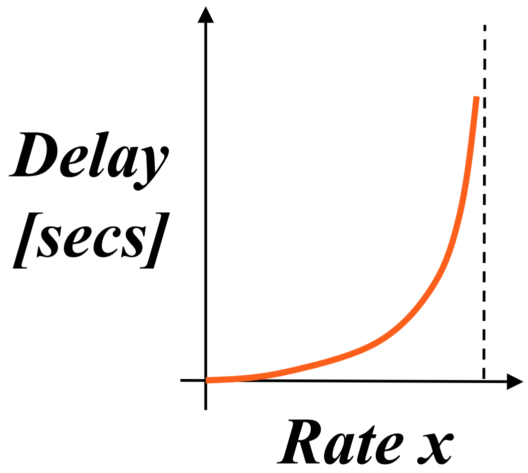

Unbounded Inefficiency

Example: large prop delay + small queuing delay vs. small prop delay + large queuing delay

- one unit (comprising many flows) selfish traffic

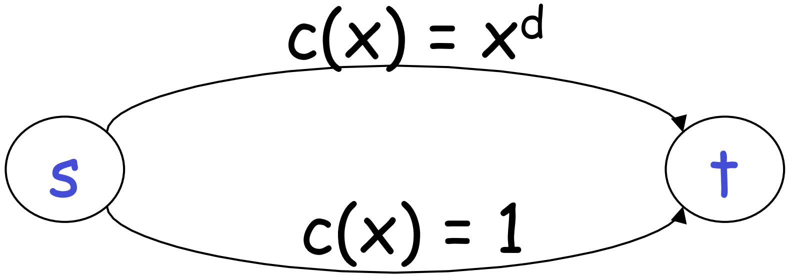

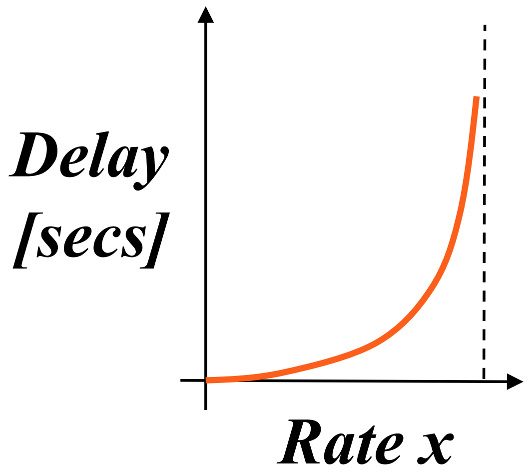

Unbounded Inefficiency

Example: large prop delay + small queuing delay vs. small prop delay + large queuing delay

- one unit (comprising many flows) selfish traffic

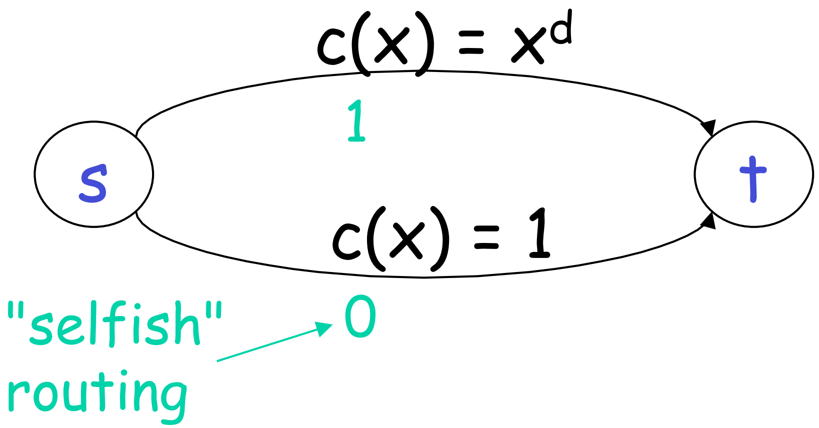

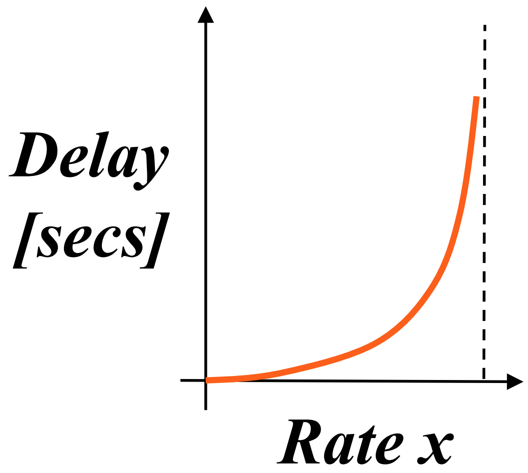

Unbounded Inefficiency

Example: large prop delay + small queuing delay vs. small prop delay + large queuing delay

- one unit (comprising many flows) selfish traffic

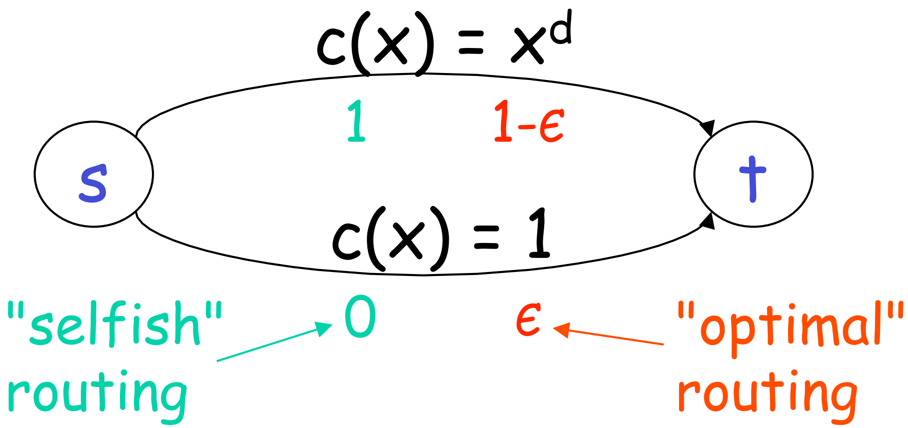

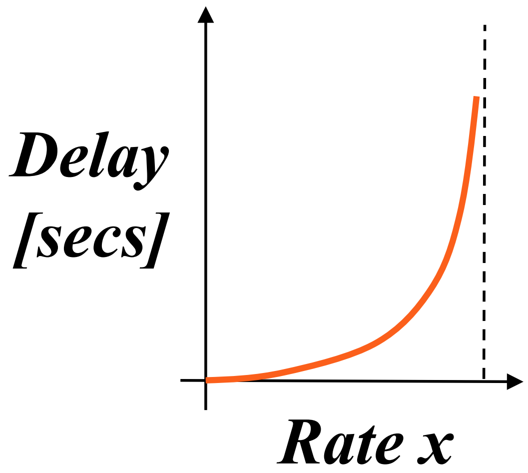

Unbounded Inefficiency

Example: large prop delay + small queuing delay vs. small prop delay + large queuing delay

- one unit (comprising many flows) selfish traffic

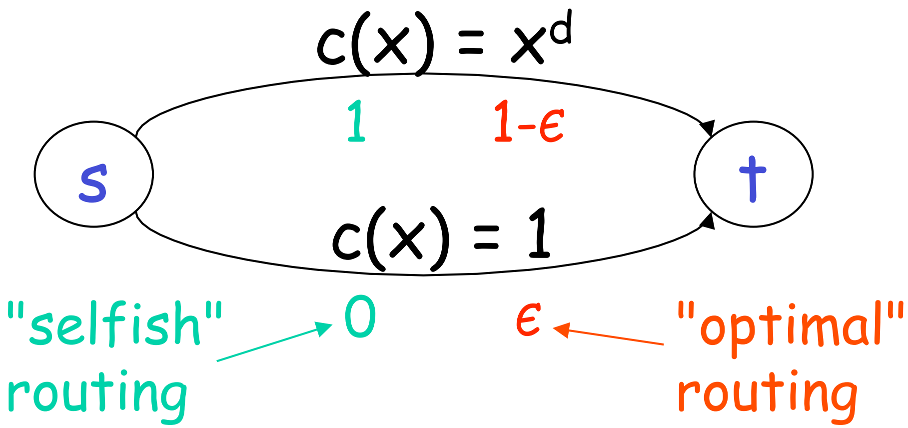

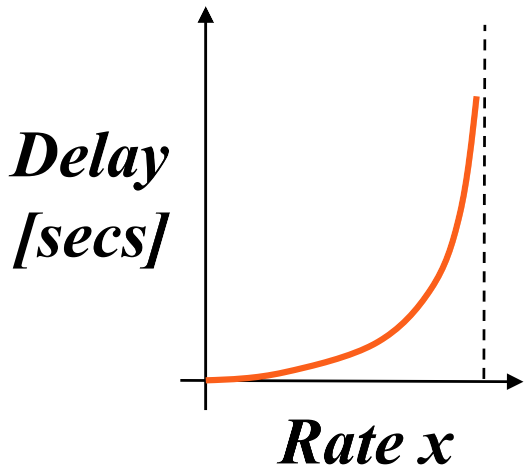

Hope: performance guarantees easier to achieve in overprovisioned network.

Benefit of Overprovisioning

Suppose: network is overprovisioned by $\beta > 0$ (β fraction of each edge unused).

Then: Delay of selfish routing at most $\frac{1}{2}(1+1/\sqrt{\beta})$ times that of optimal.

- arbitrary network size/topology, traffic matrix
- special case of [Roughgarden STOC 02]

Moral: Even modest (10%) over-provisioning sufficient for near-optimal routing.

But Are We at Equilibrium?

Since 2002: price of anarchy (i.e., worst eq/OPT ratio) analyzed in many models.

Possible critique: Interpretation of a POA bound presumes players reach equilibrium.

- assumes players are "rational" and *also* successfully coordinate on an equilibrium

Example Generalization

Definition: a sequence $s^1, s^2, \ldots, s^T$ of outcomes is *no-regret* if:

- for each player i, each fixed action q_i:
  - average cost player i incurs over sequence no worse than playing action q_i every time
  - simple hedging strategies can be used by players to enforce this (for suff large T)

Interpretation: players are at least "somewhat smart", but don't necessarily coordinate.

Intrinsic Robustness of the Price of Anarchy

Informal Theorem: [Roughgarden STOC 09] in many applications, every bound on the price of anarchy (for Nash equilibria) extends *automatically* to (e.g.) all no-regret sequences.

Example Application: selfish routing games ("nonatomic" or "atomic") with cost functions in an arbitrary fixed set.

Outline of Proof

- main definition: a "canonical way" to bound the price of anarchy (for pure equilibria)
- theorem 1: every POA bound proved "canonically" is *automatically far stronger*
  - e.g., even applies "out-of-equilibrium", assuming no-regret play
- theorem 2: canonical method provably yields optimal bounds in fundamental cases

Connections + Challenges

- dynamics in games + inefficiency bounds
  - e.g., how do details of dynamics affect which equilibrium is reached?
- *possible application in control theory*: worst-case performance guarantees for distributed approximations of a centralized optimum
- *possible application in control theory*: meaningful guarantees despite non-convergence of system