Networks, Potential Functions, and the Price of Anarchy

Tim Roughgarden
Stanford University

PRICE OF ANARCHY (1G)

Pigou's Example

Example: one unit of traffic wants to go from s to t

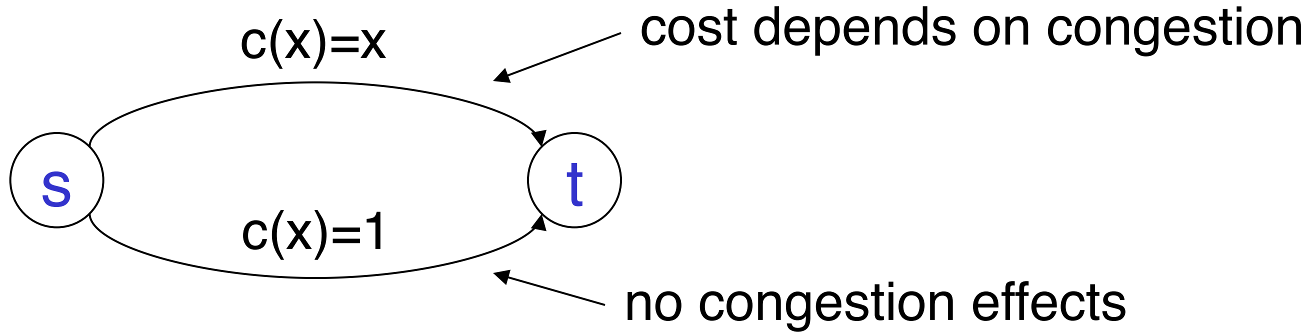

Question: what will selfish network users do?

- assume everyone wants smallest-possible cost
- [Pigou 1920]

Motivating Example

Claim: all traffic will take the top link.

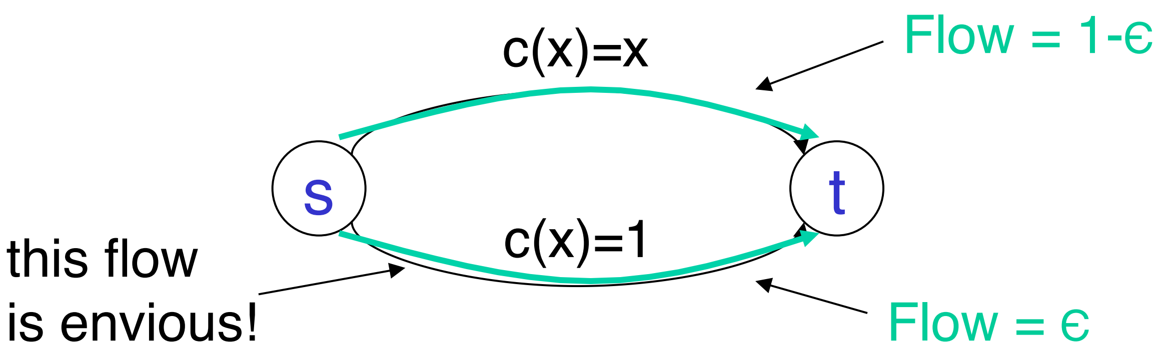

Reason:

- ϵ > 0 ⇒ traffic on bottom is envious
- ϵ = 0 ⇒ equilibrium
  - all traffic incurs one unit of cost

Can We Do Better?

Consider instead: traffic split equally

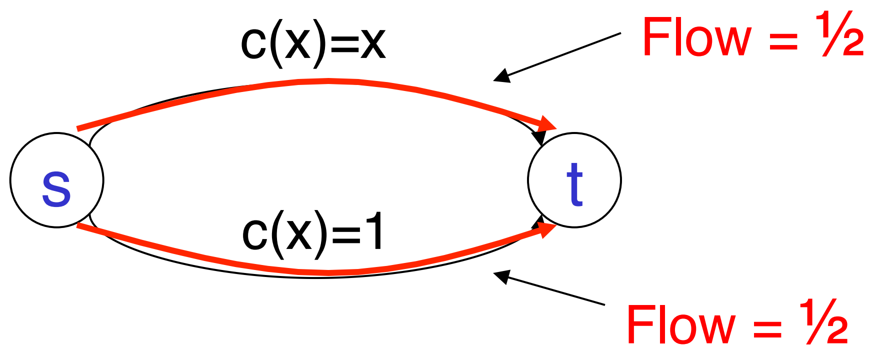

Improvement:

- half of traffic has cost 1 (same as before)
- half of traffic has cost ½ (much improved!)

Braess's Paradox

Initial Network:

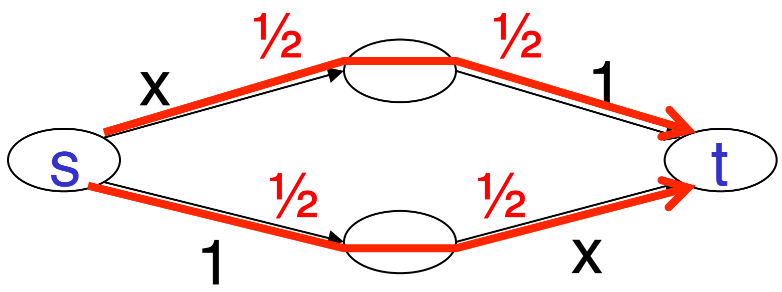

Cost = 1.5

Braess's Paradox

Initial Network:

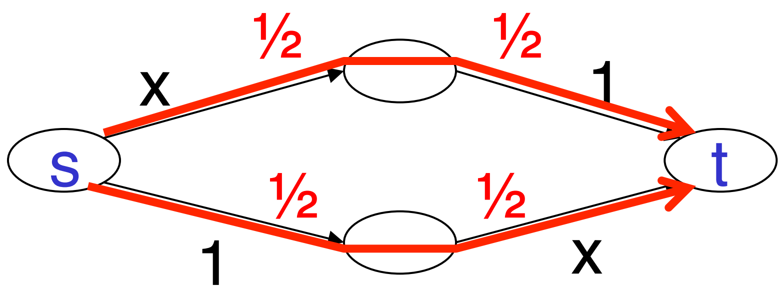

Cost = 1.5

Augmented Network:

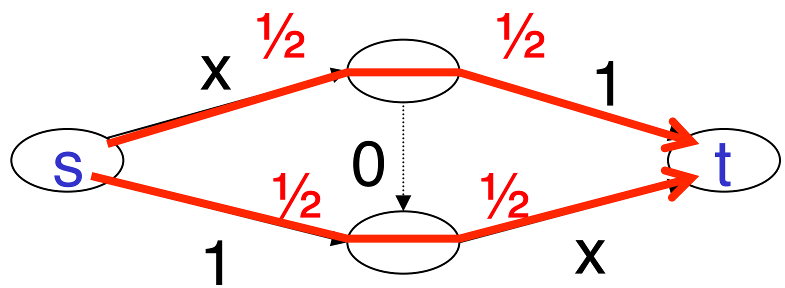

Now what?

Braess's Paradox

Initial Network:

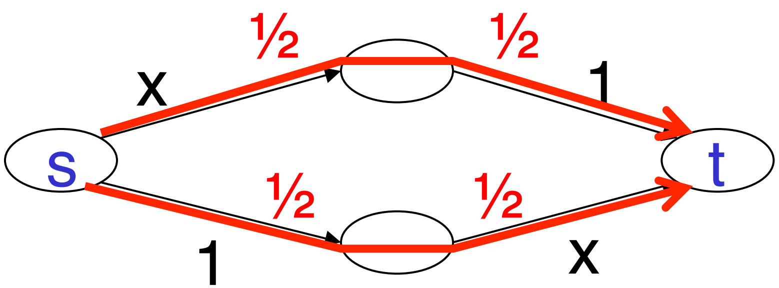

Cost = 1.5

Augmented Network:

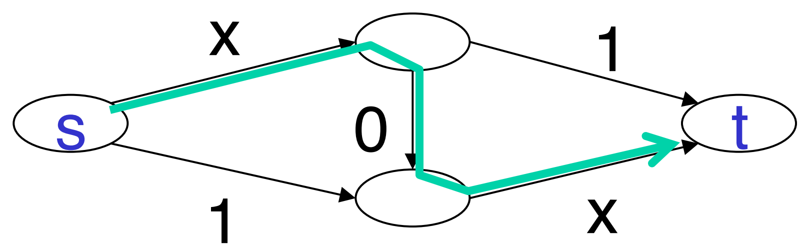

Cost = 2

Braess's Paradox

Initial Network:

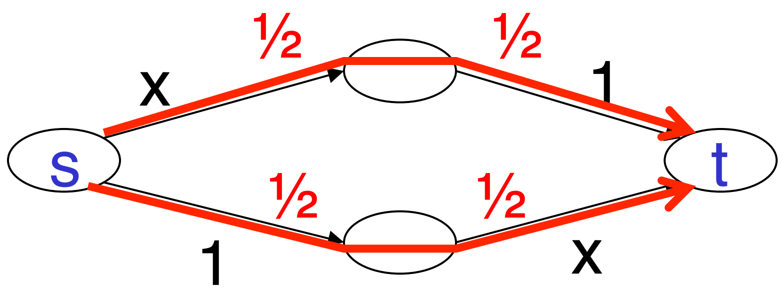

Cost = 1.5

Augmented Network:

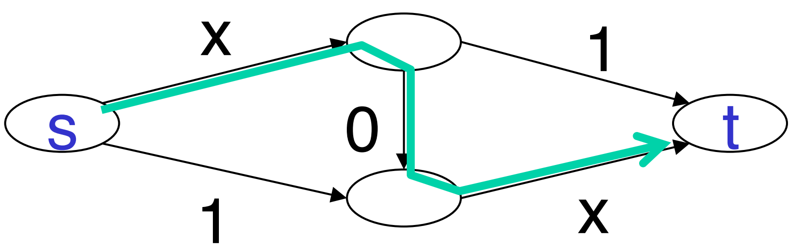

Cost = 2

All traffic incurs more cost! [Braess 68]

- also has physical analogs [Cohen/Horowitz 91]

High-Level Overview

Motivation: equilibria of noncooperative network games typically inefficient

- e.g., Pigou's example + Braess's Paradox
- don't optimize natural objective functions

Price of anarchy: quantify inefficiency with respect to an objective function

Our goal: when is the price of anarchy small?

- when does competition approximate cooperation?
- benefit of centralized control is small

Nonatomic Selfish Routing

- directed graph $G = (V,E)$
- source-destination pairs $(s_1,t_1), \ldots, (s_k,t_k)$
- r_i = amount of traffic going from s_i to t_i
- for each edge e, a cost function $c_e(\bullet)$
  - assumed continuous and nondecreasing

Defn: a multicommodity flow is an *equilibrium* if all traffic routed on shortest paths.

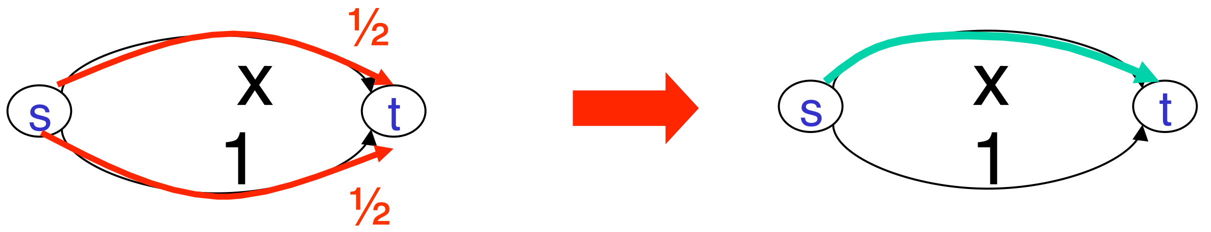

The Price of Anarchy

Defn: price of anarchy of a game = $\dfrac{\text{obj fn value of worst equilibrium}}{\text{optimal obj fn value}}$

– definition from [Koutsoupias/Papadimitriou 99]

Example: POA = 4/3 in Pigou's example

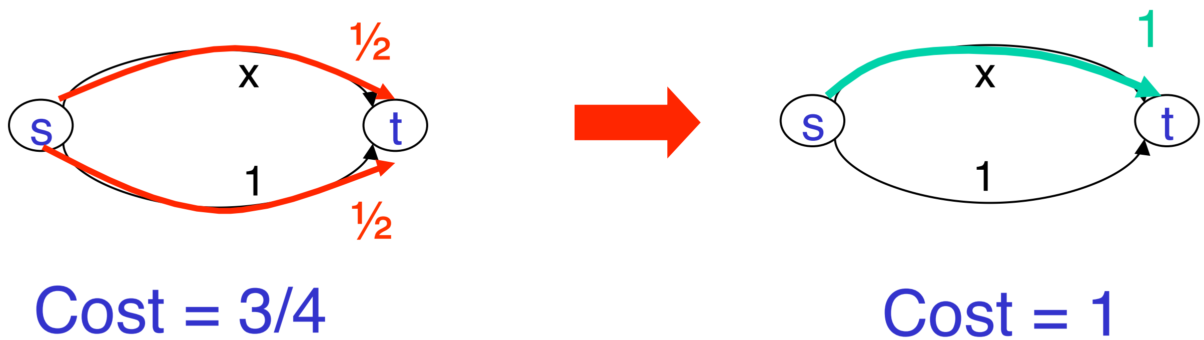

A Nonlinear Pigou Network

Bad Example: (d large)

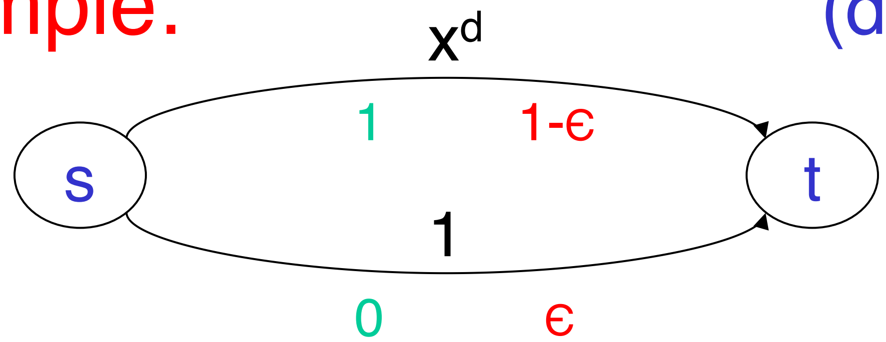

equilibrium has cost 1, min cost $\to$ 0

A Nonlinear Pigou Network

Bad Example: (d large)

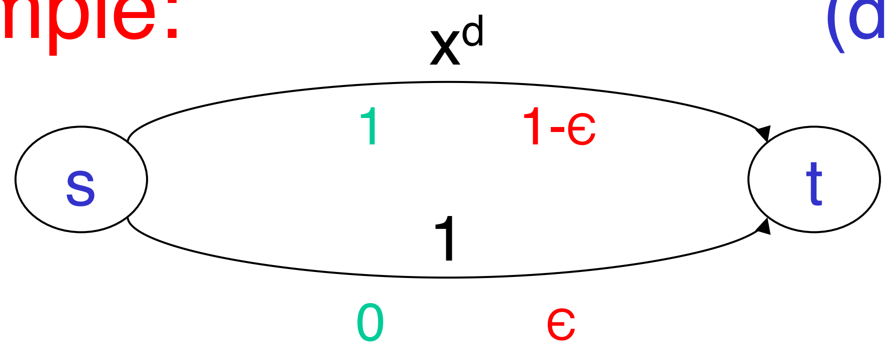

equilibrium has cost 1, min cost → 0

⇒ price of anarchy unbounded as d -> infinity

Goal: weakest-possible conditions under which the price of anarchy is small.

When Is the Price of Anarchy Bounded?

Examples so far:

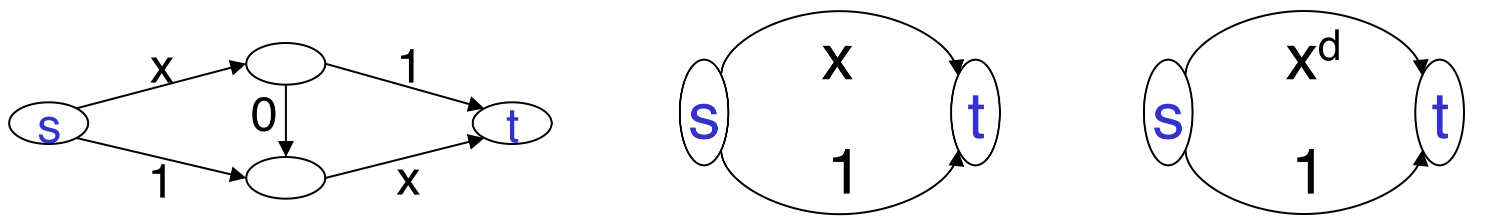

Hope: imposing additional structure on the cost functions helps

– worry: bad things happen in larger networks

Polynomial Cost Functions

Defn: linear cost function is of form $c_e(x)=a_e x+b_e$

Theorem: [Roughgarden/Tardos 00] for every network with linear cost functions:

cost of Nash flow $\leq$ 4/3 × cost of opt flow

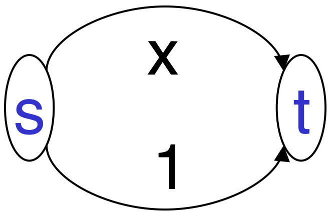

Polynomial Cost Functions

Defn: linear cost function is of form $c_e(x)=a_ex+b_e$

Theorem: [Roughgarden/Tardos 00] for every network with linear cost functions:

cost of Nash flow $\leq$ 4/3 × cost of opt flow

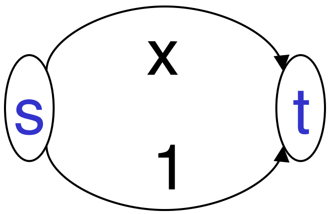

Bounded-degree polynomials: replace 4/3 by $\approx$ d/ln d

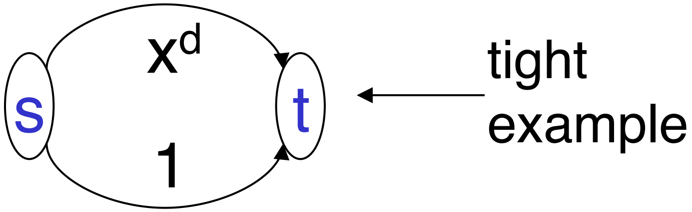

A General Theorem

Thm: [Roughgarden 02], [Correa/Schulz/Stier Moses 03] fix any set of cost functions. Then, a Pigou-like example --- 2 nodes, 2 links, 1 link w/a constant cost function --- achieves the worst P.O.A.

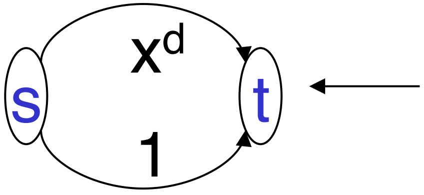

tight example

Interpretation

Bad news: inefficiency of selfish routing grows as cost functions become "more nonlinear".

- think of "nonlinear" as "heavily congested"
- recall nonlinear Pigou's example

Good news: inefficiency does not grow with network size or # of source-destination pairs.

- in lightly loaded networks, no matter how large, selfish routing is nearly optimal

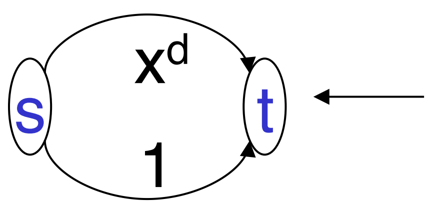

tight example

Benefit of Overprovisioning

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

Then: Price of anarchy is at most $\frac{1}{2}(1+1/\sqrt{\beta})$.

- arbitrary network size/topology, traffic matrix

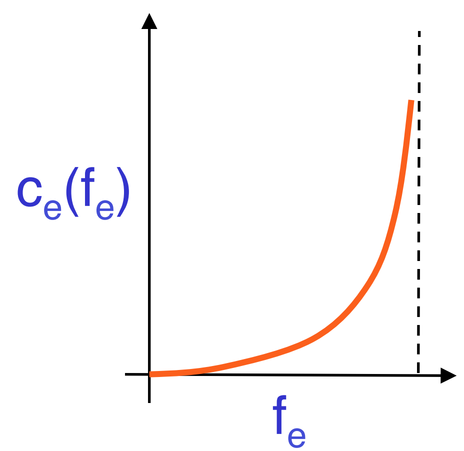

Moral: Even modest (10%) over-provisioning sufficient for near-optimal routing.

Potential Functions

- potential games: equilibria are actually optima of a related optimization problem
  - has immediate consequences for existence, uniqueness, and inefficiency of equilibria
  - see [Beckmann/McGuire/Winsten 56], [Rosenthal 73], [Monderer/Shapley 96], for original references
  - see [Roughgarden ICM 06] for survey

The Potential Function

Key fact: [BMV 56] Nash flows minimize “potential function” $\sum_e \int_0^{f_e} c_e(x)dx$ (over all flows).

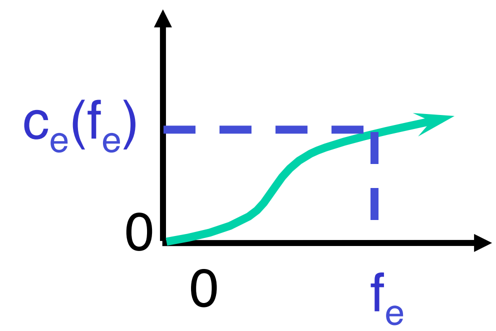

The Potential Function

Key fact: [BMV 56] Nash flows minimize "potential function" $\sum_e \int_0^{f_e} c_e(x)dx$ (over all flows).

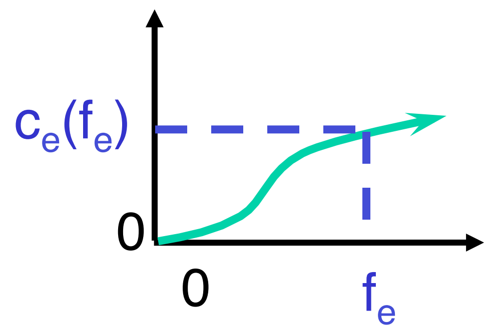

Lemma 1: locally optimal solutions are precisely the Nash flows (derivative test).

Lemma 2: all locally optimal solutions are also globally optimal (convexity).

Corollary: Nash flows exist, are unique.

Consequences for the Price of Anarchy

Example: linear cost functions.

Compare cost and potential functions:

$$C(f) = \sum_e f_e \cdot c_e(f_e) = \sum_e [a_e f_e^2 + b_e f_e]$$

$$PF(f) = \sum_e \int_0^{f_e} c_e(x)dx = \sum_e [(a_e f_e^2)/2 + b_e f_e]$$

Consequences for the Price of Anarchy

Example: linear cost functions.

Compare cost and potential functions:

$$C(f) = \sum_e f_e \cdot c_e(f_e) = \sum_e [a_e f_e^2 + b_e f_e]$$

$$PF(f) = \sum_e \int_0^{f_e} c_e(x)dx = \sum_e [(a_e f_e^2)/2 + b_e f_e]$$

- cost, potential functions differ by factor of ≤ 2
- gives upper bound of 2 on price on anarchy
  - $C(f) \le 2 \times PF(f) \le 2 \times PF(f^*) \le 2 \times C(f^*)$

Better Bounds?

Similarly: proves bound of d+1 for degree-d polynomials (w/nonnegative coefficients).

- not tight, but qualitatively accurate
  - e.g., price of anarchy goes to infinity with degree bound, but only linearly
- to get tight bounds, need "variational inequalities"
  - see my ICM survey for details

PRICE OF ANARCHY (2G)

POA Bounds Without Convergence

Meaning of a POA bound: *if* the game is at an equilibrium, *then* outcome is near-optimal.

Problem: what if can't reach an equilibrium?

- non-existence (pure Nash equilibria)
- intractability (mixed Nash equilibria) [Daskalakis/Goldberg/Papadimitriou 06], [Chen/Deng/Teng 06]

Worry: are our POA bounds "meaningless"?

POA Bounds Without Convergence

Theorem: [Roughgarden STOC 2009] most known POA bounds hold *even if the game is not at Nash equilibrium!*

- e.g., if game is played repeatedly, no-regret conditions or a few myopic best responses are enough

Concluding Remarks

- lens of approximation gives new insights into fundamental mathematical models
- good bounds for many games of interest, even out-of-Nash-equilibrium
  - refutes non-existence/intractability critiques
- routing games: worst-case price of anarchy depends only on "nonlinearity" of cost functions
  - parameterize POA bounds via worst-case examples
  - equilibria inadvertently optimize a potential function