How Much Can Taxes Help Selfish Routing?

Tim Roughgarden (Cornell)

Joint with Richard Cole (NYU)
and Yevgeniy Dodis (NYU)

Selfish Routing

- a directed graph $G = (V,E)$
- a source s and a destination t
- one unit of traffic from s to t
- for each edge e, a latency function $\ell_e(\cdot)$
  - assumed continuous, nondecreasing

Example:

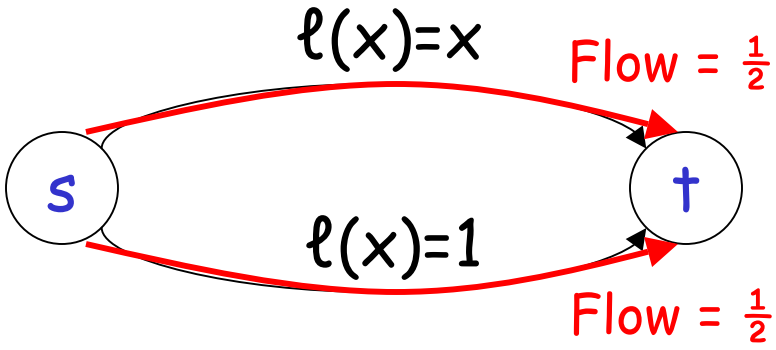

Routings of Traffic

Traffic and Flows:

- f_P = fraction of traffic routed on s-t path P
- flow vector f ⇔ routing of traffic

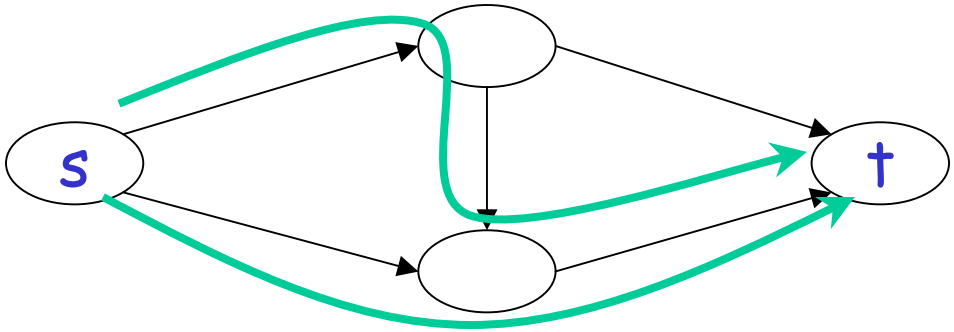

Selfish routing: what flows arise as the routes chosen by many noncooperative agents?

Nash Flows

Some assumptions:

- agents small relative to network
- want to minimize personal latency

Def: A flow is at Nash equilibrium (or is a Nash flow) if all flow is routed on min-latency paths [given current edge congestion]

  - have existence, uniqueness [Wardrop, Beckmann et al 50s]

Example:

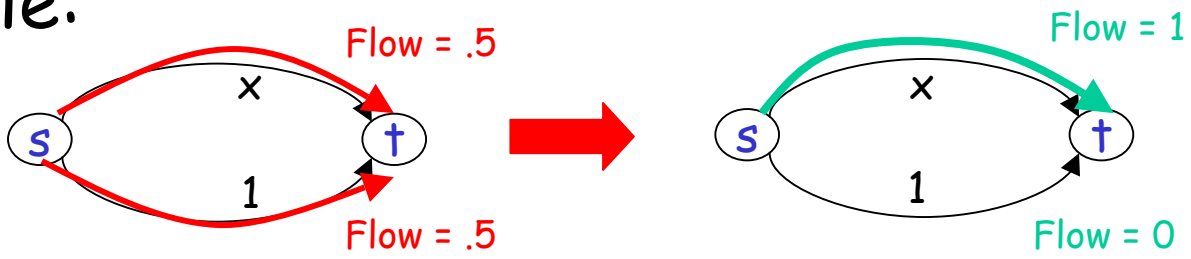

Inefficiency of Nash Flows

Our objective function: average latency

- $\Rightarrow$ Nash flows need not be optimal
- observed informally by [Pigou 1920]

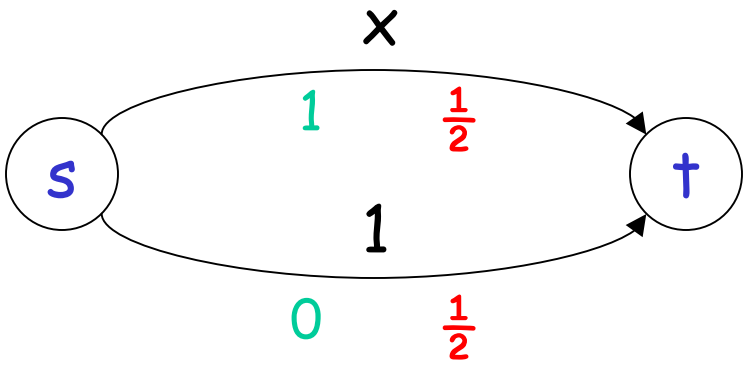

- Average latency of Nash flow = $1 \cdot 1 + 0 \cdot 1 = 1$
- of optimal flow = $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = \frac{3}{4}$

Braess's Paradox

Initial Network:

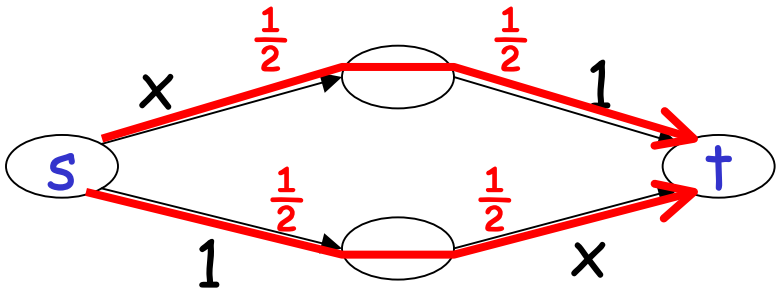

Delay = 1.5

Braess's Paradox

Initial Network:

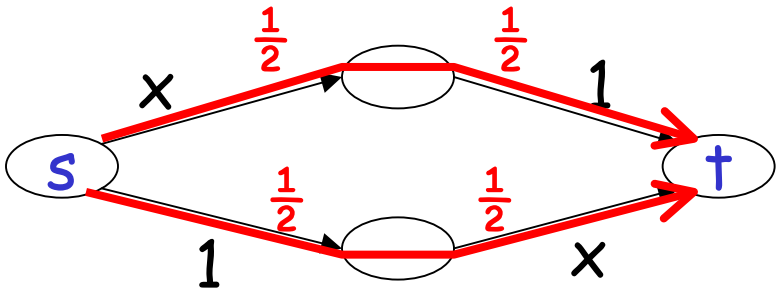

Delay = 1.5

Augmented Network:

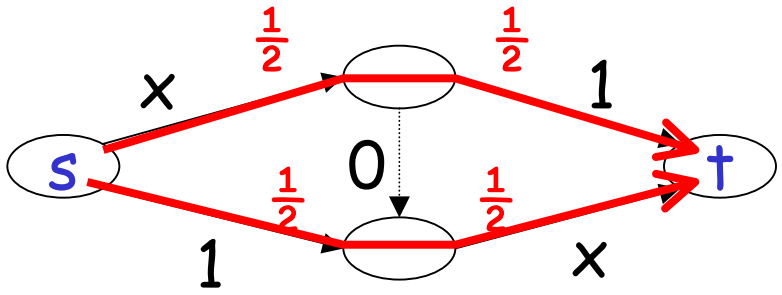

Now what?

Braess's Paradox

Initial Network:

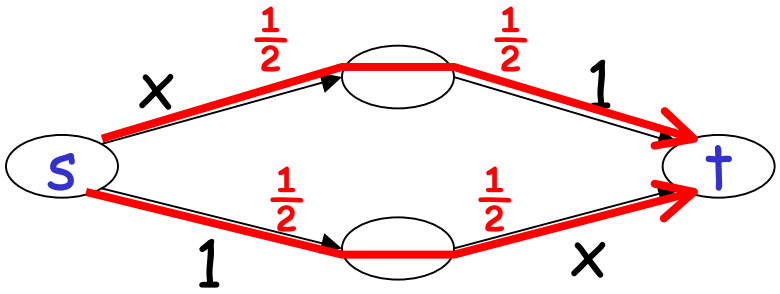

Delay = 1.5

Augmented Network:

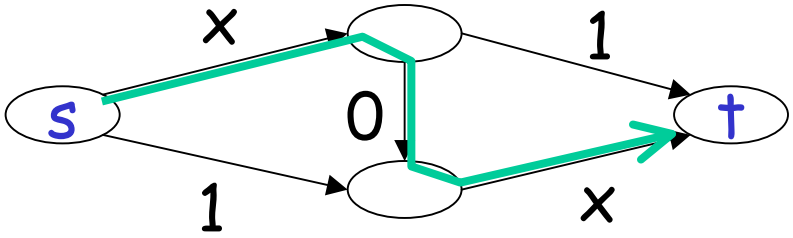

Delay = 2

All traffic incurs more delay! [Braess 68]

Marginal Cost Taxes

Goal: do better with taxes (one per edge)

- not addressing implementation

Marginal Cost Taxes

Goal: do better with taxes (one per edge)

- not addressing implementation

Assume: all traffic minimizes time + money

- see STOC '03 paper for relaxing this

Def: the marginal cost tax of an edge (w.r.t. a flow) is the extra delay to existing traffic caused by a marginal increase in traffic

Marginal Cost Taxes

Goal: do better with taxes (one per edge)

- not addressing implementation

Assume: all traffic minimizes time + money

- see STOC '03 paper for relaxing this

Def: the marginal cost tax of an edge (w.r.t. a flow) is the extra delay to existing traffic caused by a marginal increase in traffic

Thm: [folklore] marginal cost taxes w.r.t. the opt flow induce the opt flow as a Nash eq.

Are Taxes a Social Loss?

- Problem with MCT: min delay is holy grail; exorbitant taxes ignored

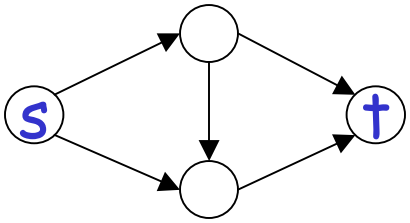

Are Taxes a Social Loss?

- Problem with MCT: min delay is holy grail; exorbitant taxes ignored
- Ever reasonable?: yes, iff taxes can be refunded (directly or indirectly)

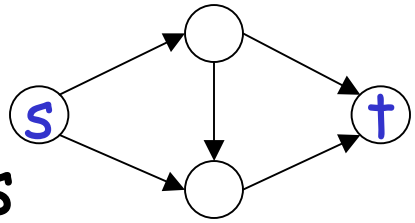

Are Taxes a Social Loss?

- Problem with MCT: min delay is holy grail; exorbitant taxes ignored

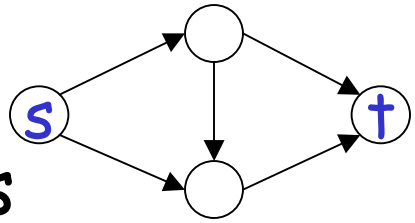

- Ever reasonable?: yes, iff taxes can be refunded (directly or indirectly)
- New Goal: minimize total disutility with nonrefundable taxes (delay + taxes paid)
  - call new objective fn the cost
  - marginal cost taxes now not a good idea, e.g.:
  - Thm: w/linear latency fns, MCT never help.

Taxes vs. Edge Removal

Note: taxes at least as good as edge removal

- can effect edge deletion with large tax
- are they strictly more powerful?

Taxes vs. Edge Removal

Note: taxes at least as good as edge removal

- can effect edge deletion with large tax
- are they strictly more powerful?

Thm: taxes can improve cost by a factor of $n/2$ ($n = |V|$), but no more.

- same for edge removal [Roughgarden FOCS '01]
- also same as edge removal for restricted classes of latency fns

Taxes vs. Edge Removal

Question: taxes no better than edge removal in best case, how about in specific networks?

Taxes vs. Edge Removal

Question: taxes no better than edge removal in best case, how about in specific networks?

Thm:

- (a) taxes can improve the Nash flow cost by an $n/2$ factor more than edge removal
  - uses step function-like latency fns
  - variation of Braess graphs from [Roughgarden FOCS '01]
- (b) taxes are never more powerful than edge removal in networks w/linear latency fns

Taxes vs. Edge Removal

General Latency Fns

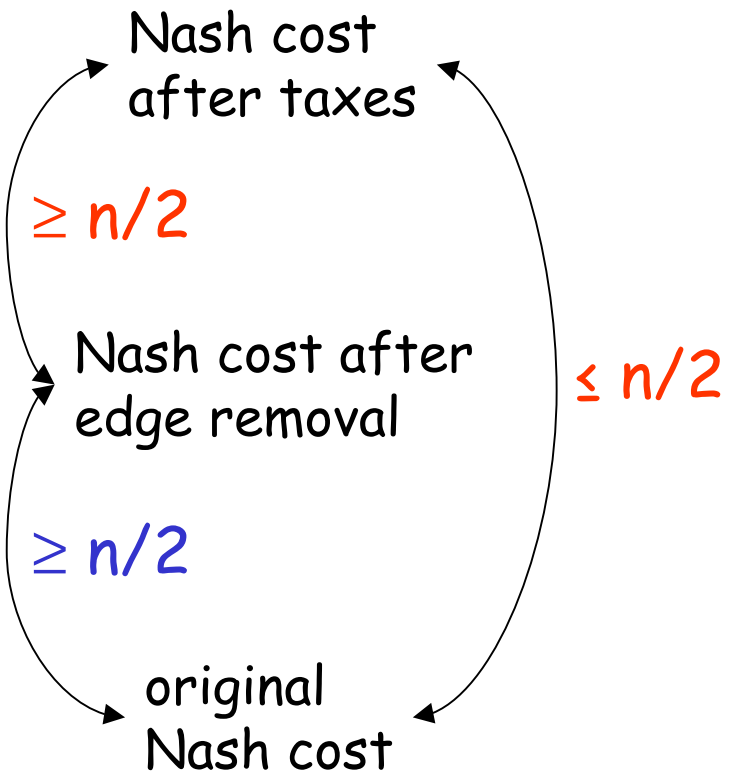

Linear Latency Fns

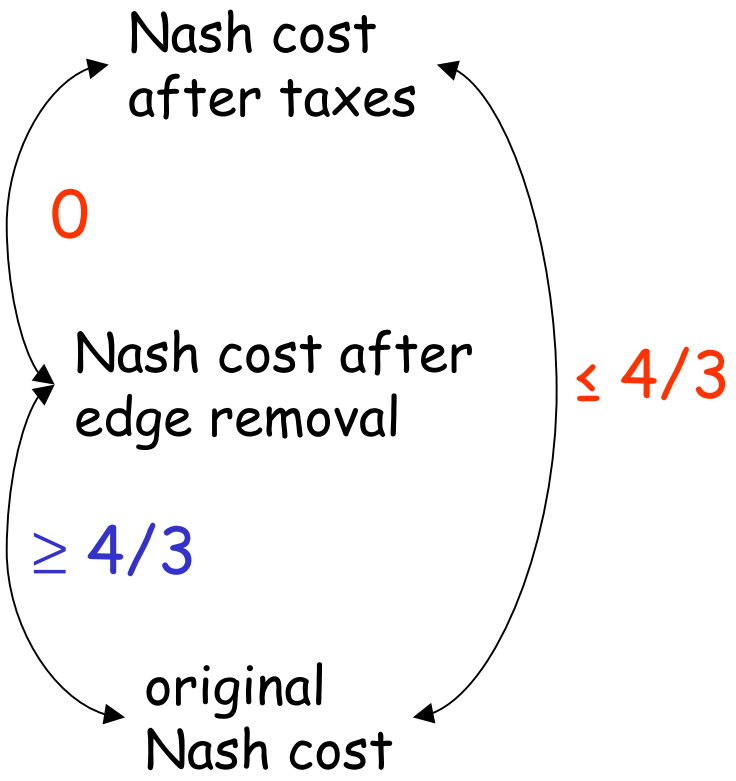

Proof Sketch for Linear Case

- First: assume false, look at minimal counterexample.
- Look at counterexample tax on this network that minimizes cost and has smallest sum.
  - Technical Lemma: this minimum exists (use minimality).
- Understand how Nash flow changes under local perturbations of the tax (minimality, linearity).
- Perturbing to a smaller tax must increase cost.
- Opposite perturbation lowers cost (contradiction).

Taxes Are Powerful but Elusive

Recall: taxes can improve cost by a factor of $n/2$ ($n = |V|$), but no more.

- powerful, but can we compute them?

Thm: optimal taxes NP-hard to approximate within factor of $o(n/\log n)$.

- complexity casts doubt on potential for taxes that minimize cost
- based on [Roughgarden FOCS '01]

Some Future Directions

- Improve model
  - convergence issues, imperfect info
  - other notions of incentive-compatibility
    - e.g., robust to malicious users
  - other objective fns
- Better results in this model
  - multicommodity flow networks