Games in Networks

and connections to algorithms

Éva Tardos
Cornell University

Why care about Games?

Users with a multitude of diverse economic interests sharing a Network (Internet)

- browsers
- routers
- servers

Selfishness:

Parties deviate from their protocol if it is in their interest

Why care about Games?

Users with a multitude of diverse economic interests sharing a Network (Internet)

- browsers
- routers
- servers

Selfishness:

Parties deviate from their protocol if it is in their interest

Model Resulting Issues as

Games on Networks

Two Related Issues

Price of Anarchy

Measure degradation of performance caused by lack of cooperation, or limited cooperation (selfishness)

Two Related Issues

Price of Anarchy

Measure degradation of performance caused by lack of cooperation, or limited cooperation (selfishness)

Mechanism Design

How to design/modify games so that selfish behavior leads to desired outcome

Two Related Issues

Price of Anarchy

Measure degradation of performance caused by lack of cooperation, or limited cooperation (selfishness)

Mechanism Design

How to design/modify games so that selfish behavior leads to desired outcome

Through Examples

- Service Provider Game
- Routing and Network Design
- Bandwidth Allocation

Price of Anarchy

$$\frac{\text{cost of selfish outcome}}{\text{"socially optimum" cost}}$$

Papadimitriou-Koutsoupias '99

Closely related to approximation algorithms

$$\text{approx ratio} = \frac{\text{cost of algorithm}}{\text{"optimum" cost}}$$

What is Selfish Outcome?

We will use: Nash equilibrium

- Current strategy “best response” for all players (no incentive to deviate)

Theorem [Nash 1952]:

- Always exists if we allow randomized strategies

We Use Nash equilibrium

- Current strategy “best response” for all players (no incentive to deviate)

Plus:

- “Stable” design
- [Nash 1952] Always exists (randomized)

We Use Nash equilibrium

- Current strategy "best response" for all players (no incentive to deviate)

Plus:

- "Stable" design
- [Nash 1952] Always exists (randomized)

Critique:

- can users learn the best behavior? [Greenwald, Friedman and Shenker, 2001]
- There can be many Nash equilibria
- Groups of users can have incentive to deviate

Nash $\leftrightarrow$ Local Optimization

Nash equilibrium:

- Current strategy "best response" for all players
- No user acting alone can improve his own solution

Local Optimum:

- No "local move" can improve the overall solution

Connection:

- Local move $\leftrightarrow$ user acting alone

Example: Service Provider Game
[Vetta FOCS'02]

Set **D** of clients, and set **F** of possible facility locations

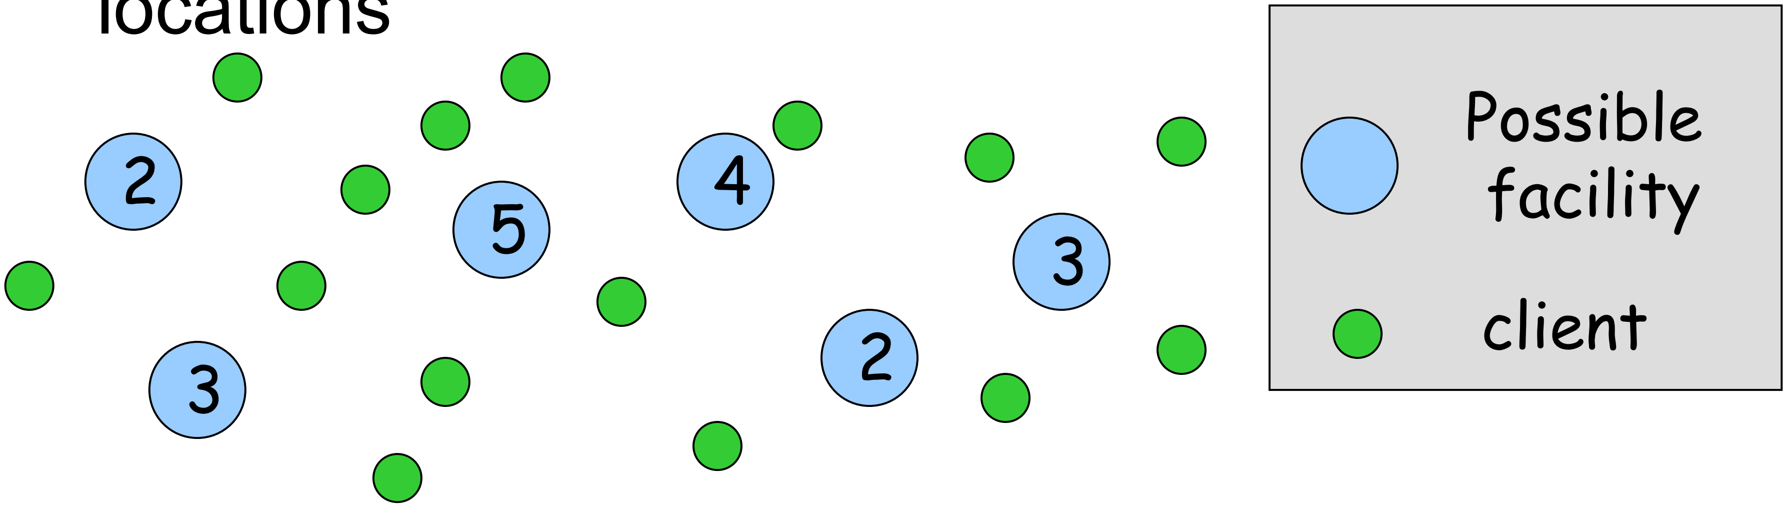

Example: Service Provider Game:

[Vetta FOCS'02]

Set **D** of clients, and set **F** of possible facility locations

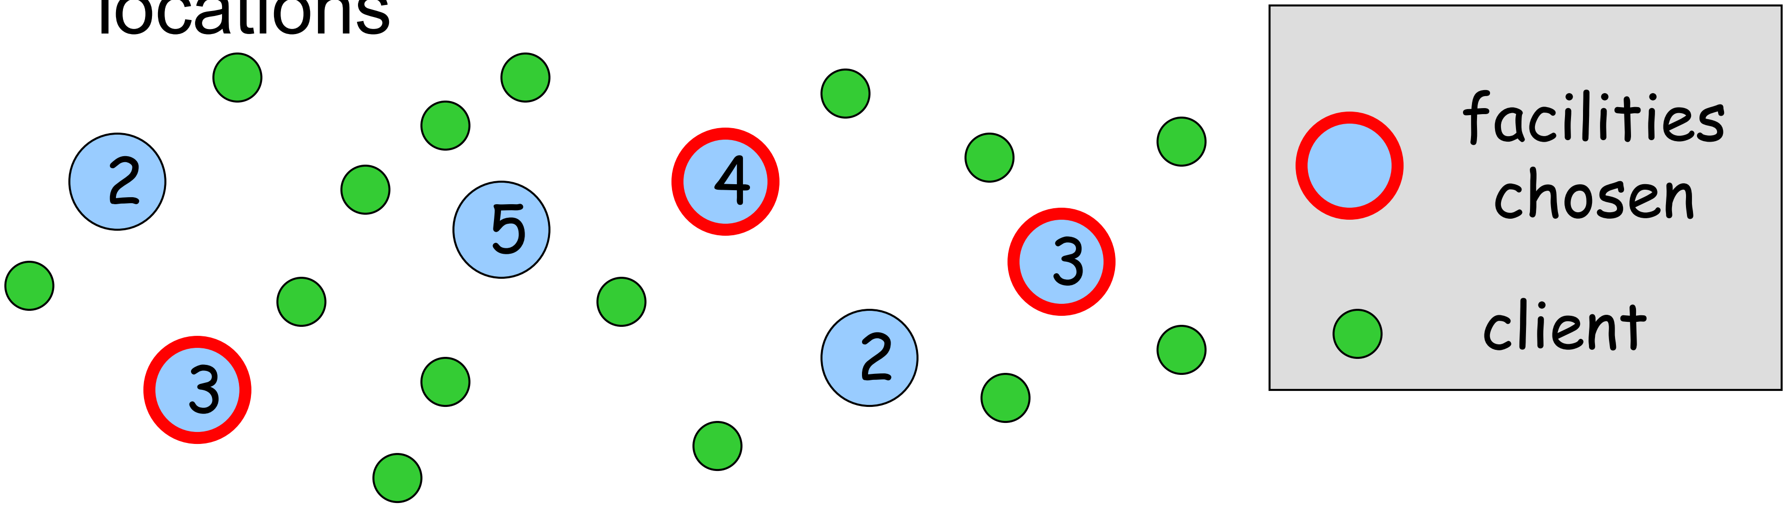

Players choose a facility to locate at

Example: Service Provider Game:

[Vetta FOCS'02]

Set **D** of clients, and set **F** of possible facility locations

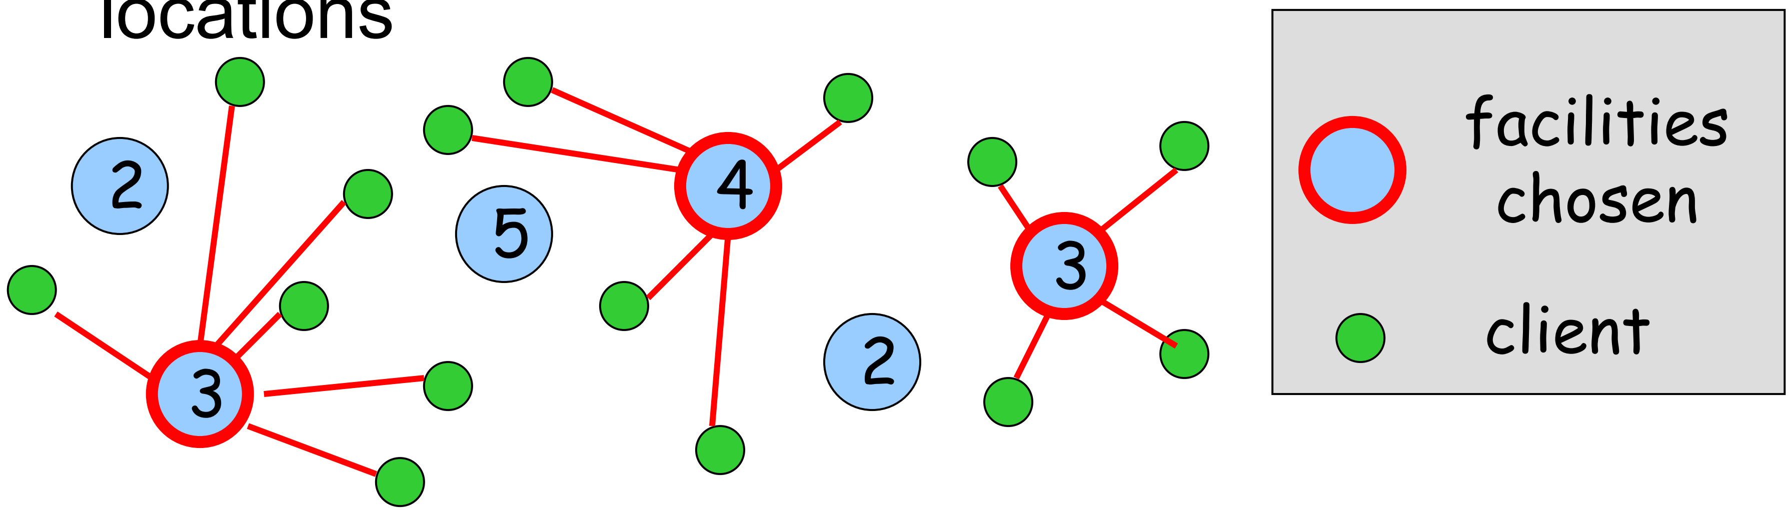

Players choose a facility to locate at

Each client connects to closest open facility

Service Provider Game: costs and income

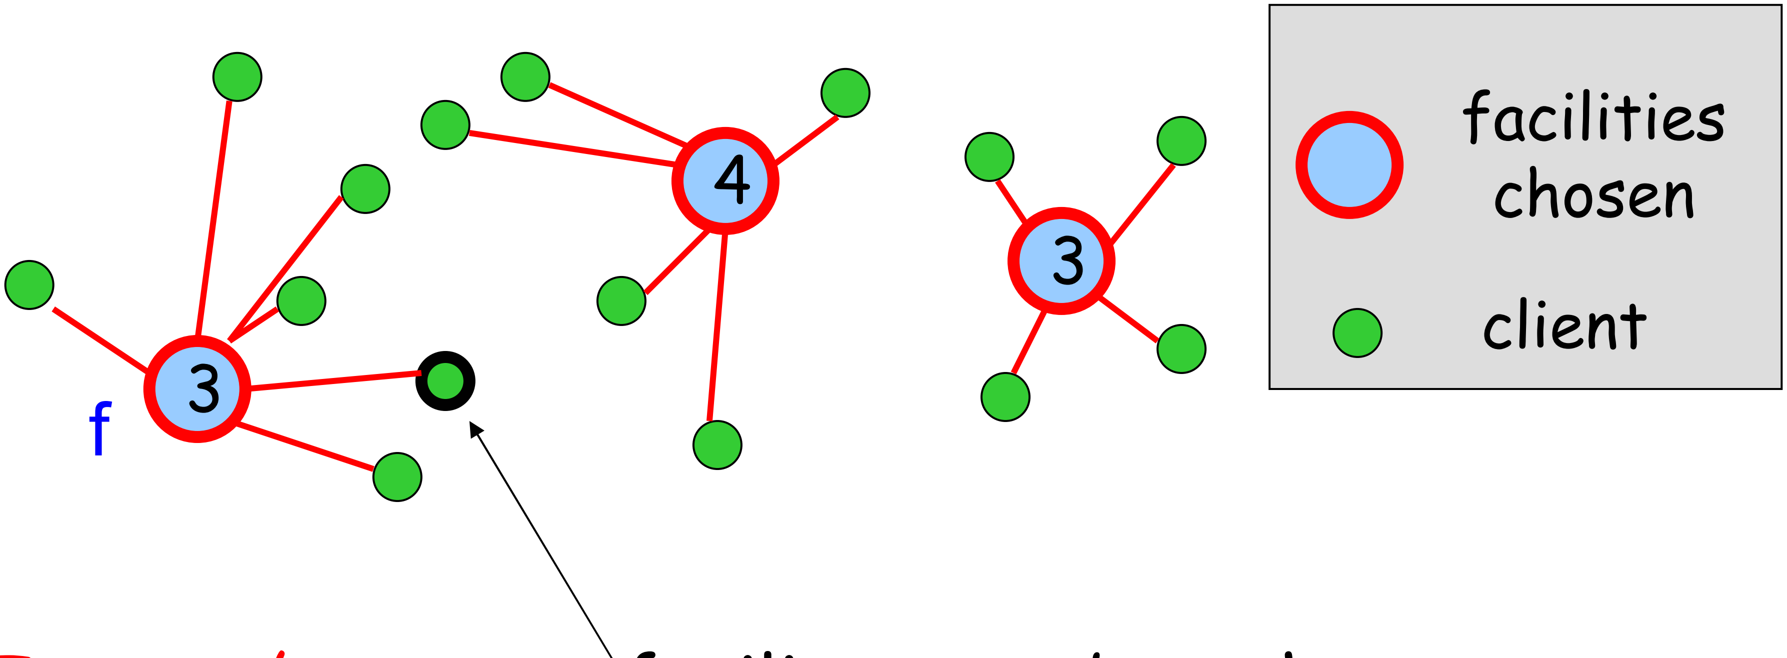

Income/cost: pay facility cost, but what to charge clients ??

Total cost: cost of facility f + $\Sigma_{p \to f}$ dist(p,f).

Service Provider Game: costs and income

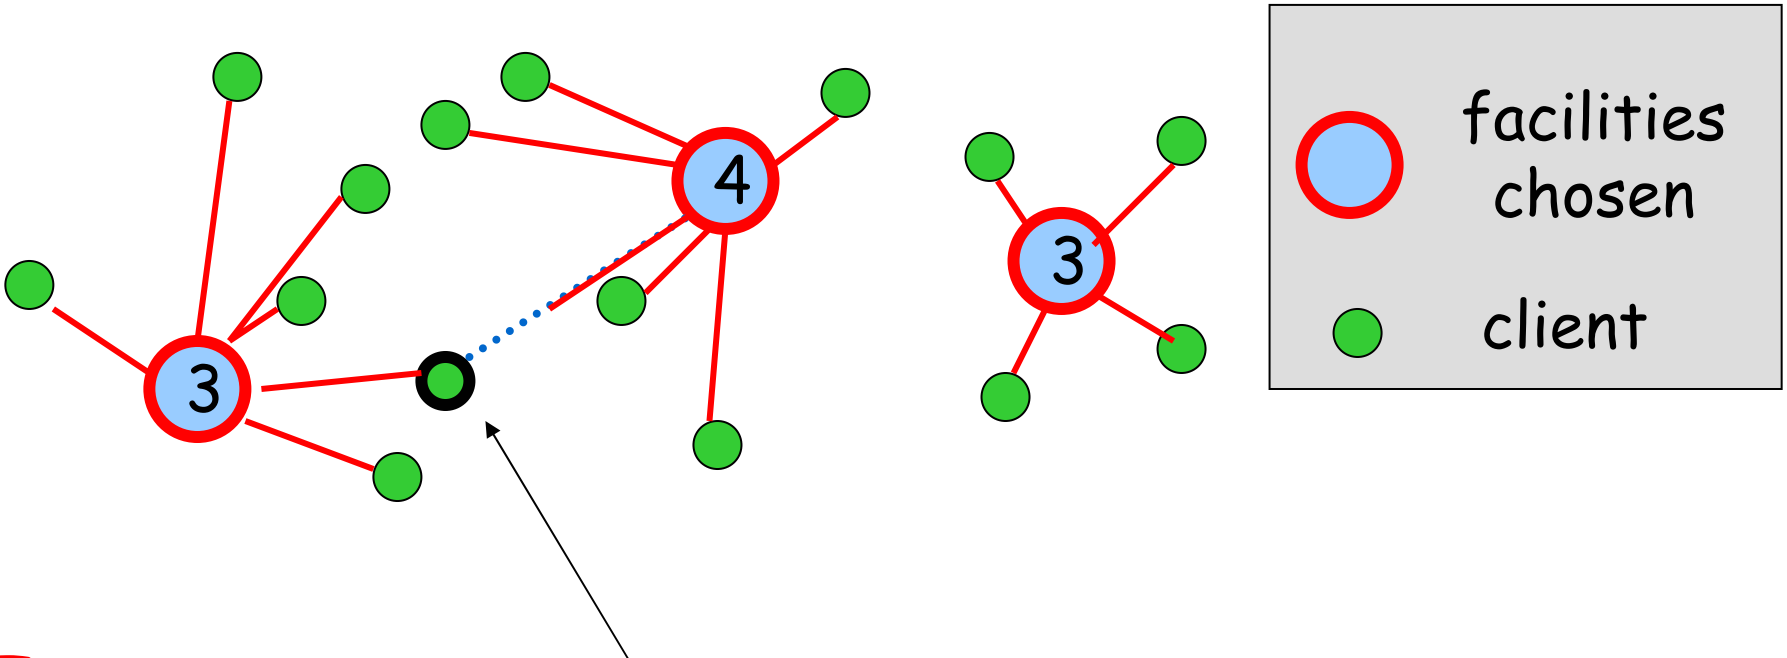

Income: charge clients for distance to second closest open facility (competitor, Vickrey price)

Provider benefit: income - cost of

Service Provider Game: selfish moves

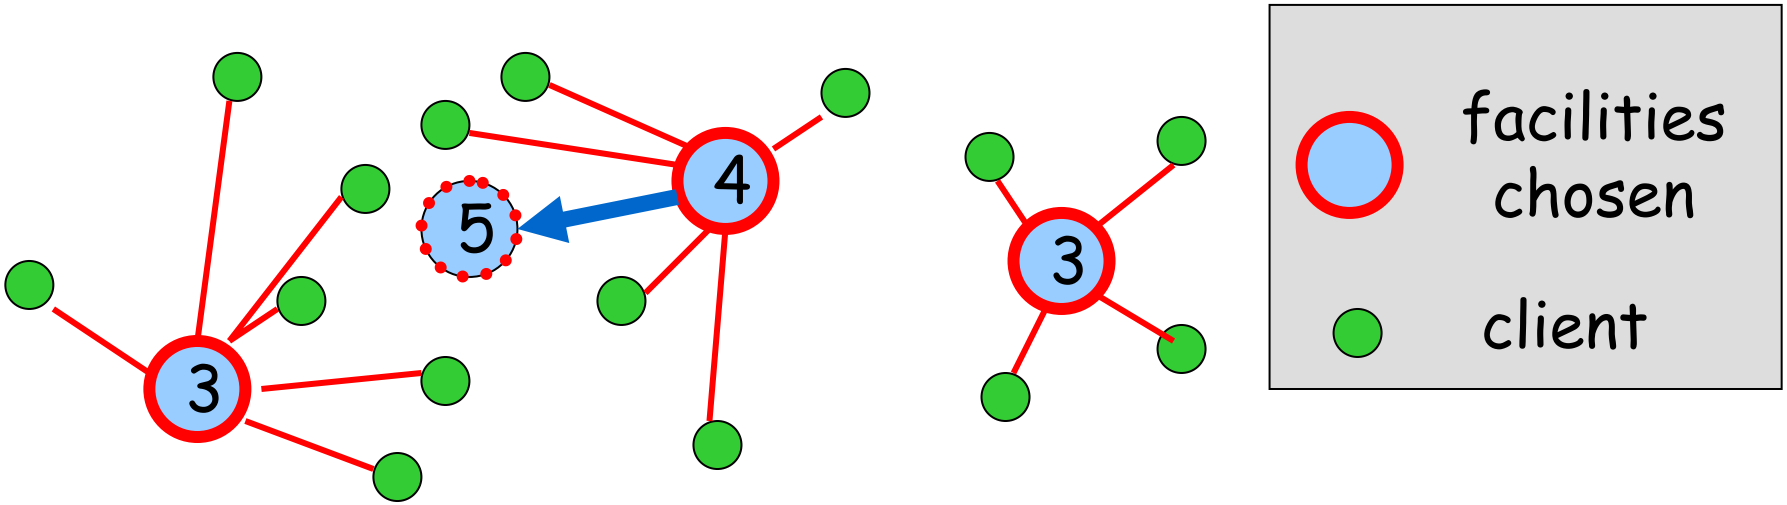

Service Provider Game:
selfish moves

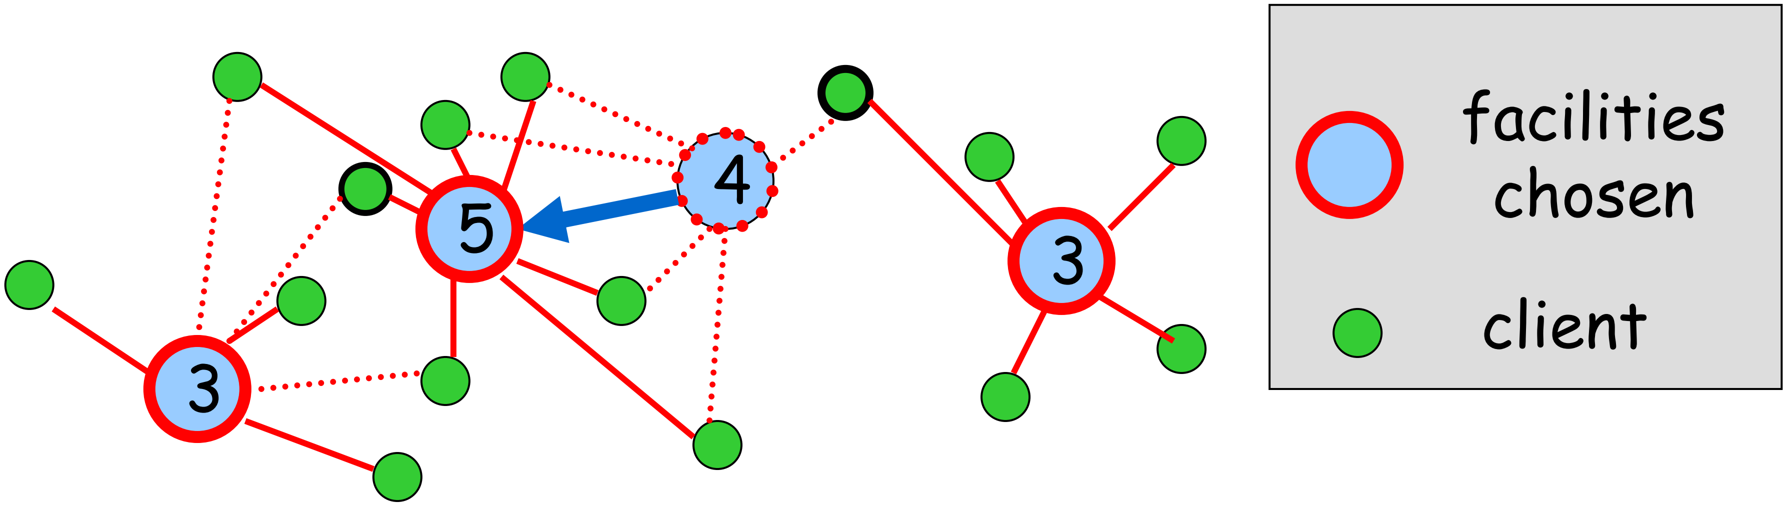

Service Provider Game: selfish moves

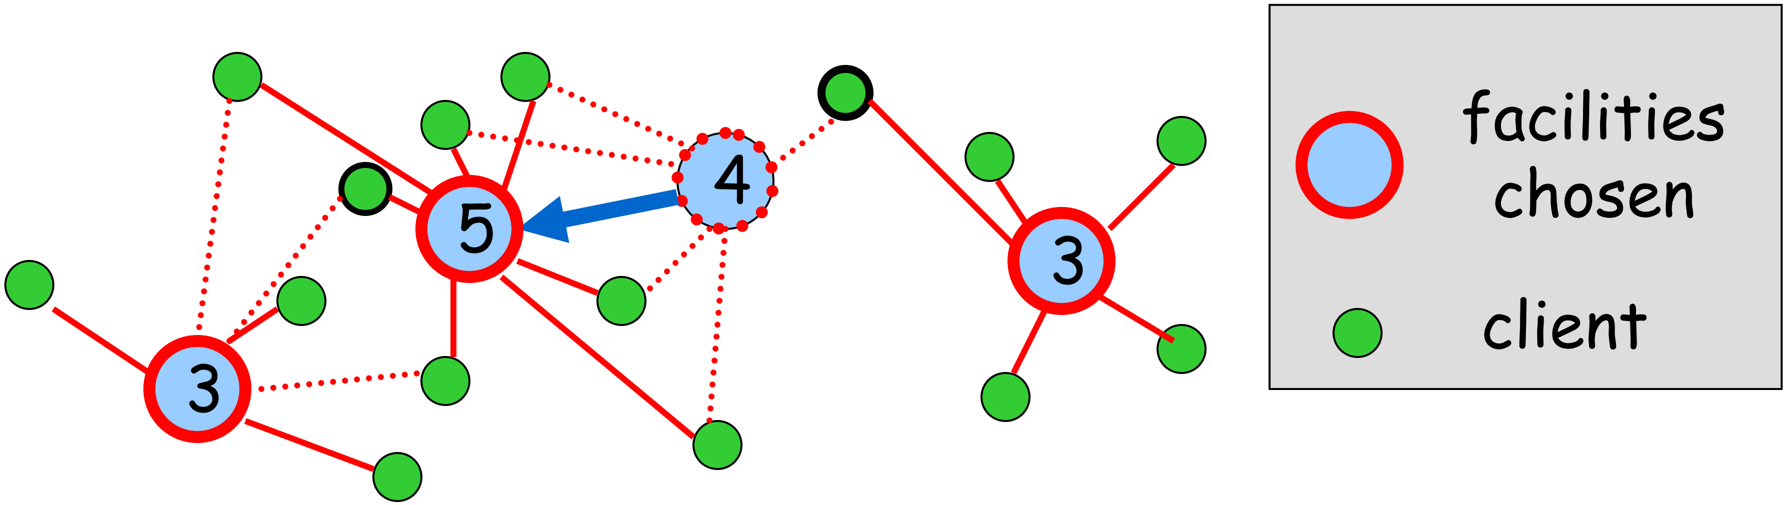

Fact: selfish improvement = better solution

- Δ in provider's income = Δ in overall solution cost

Service Provider Game:
selfish moves

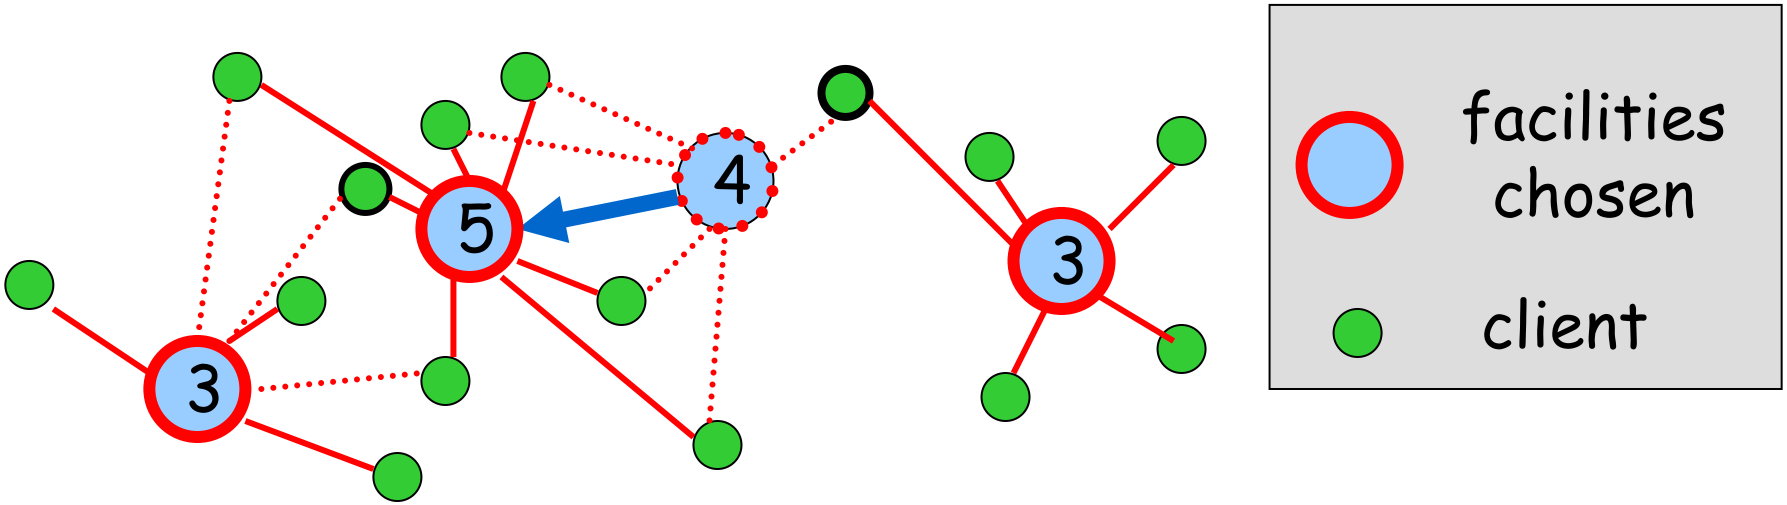

Fact: selfish improvement = better solution

- Δ in provider's income = Δ in overall solution cost
- **Why:** users switch to their best alternate
  - we charged second highest price

Service Provider Game

Selfish improvement = overall improvement

Corollary: Deterministic Nash exists

How Good is a Nash solution?

Service Provider Game

Selfish improvement = overall improvement

Corollary: Deterministic Nash exists

How Good is a Nash solution?

Vetta FOCS'2002:

- Any Nash 2-approximates social happiness of no facility costs (valuations-cost)

Devanur, et al. 2004

- Any Nash 5-approximates total cost

Service Provider Game

Selfish improvement = overall improvement

Corollary: Deterministic Nash exists

How Good is a Nash solution?

Vetta FOCS'2002:

- Any Nash 2-approximates social happiness of no facility costs (valuations-cost)

Devanur, et al. 2004

- Any Nash 5-approximates total cost

Why? all local optima are approximately optimal

What we have seen so far:

Service Provider Game

Nash = local optima

Next: routing

weaker connection to local optimization

Example: Routing and Load Balancing Games

Routing:

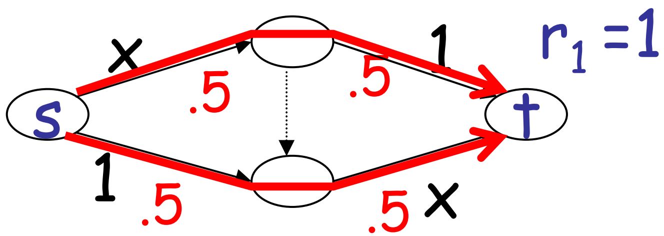

Load balancing:

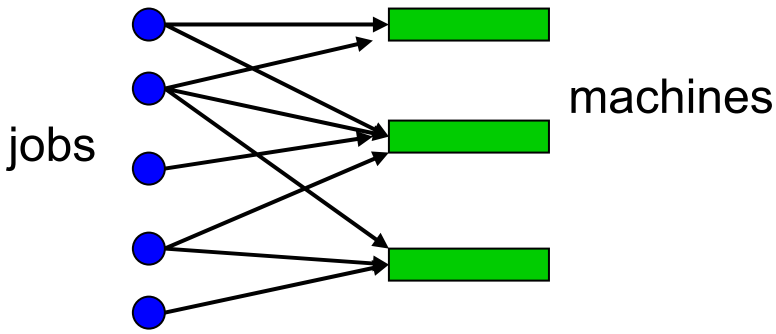

Example: Routing and Load Balancing Games

Routing:

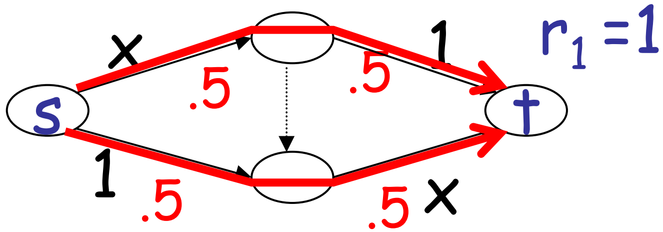

Load balancing:

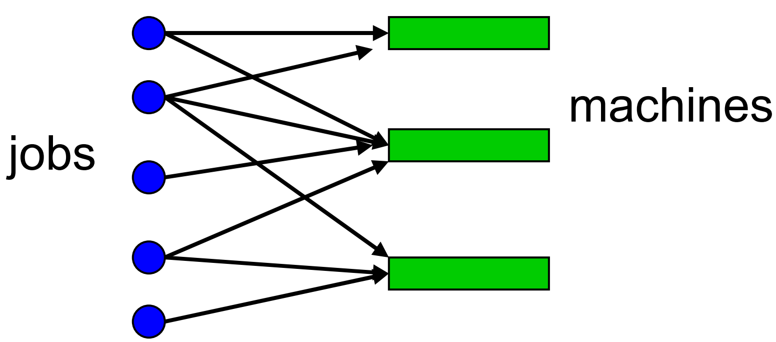

- A directed graph $G = (V,E)$
- source-sink pairs s_i,t_i for $i=1,..,k$
- rate $r_i \geq 0$ of traffic between s_i and t_i for each $i=1,..,k$
- For each edge e or machine, a latency function $\ell_e(\cdot)$

Example: Braess paradox

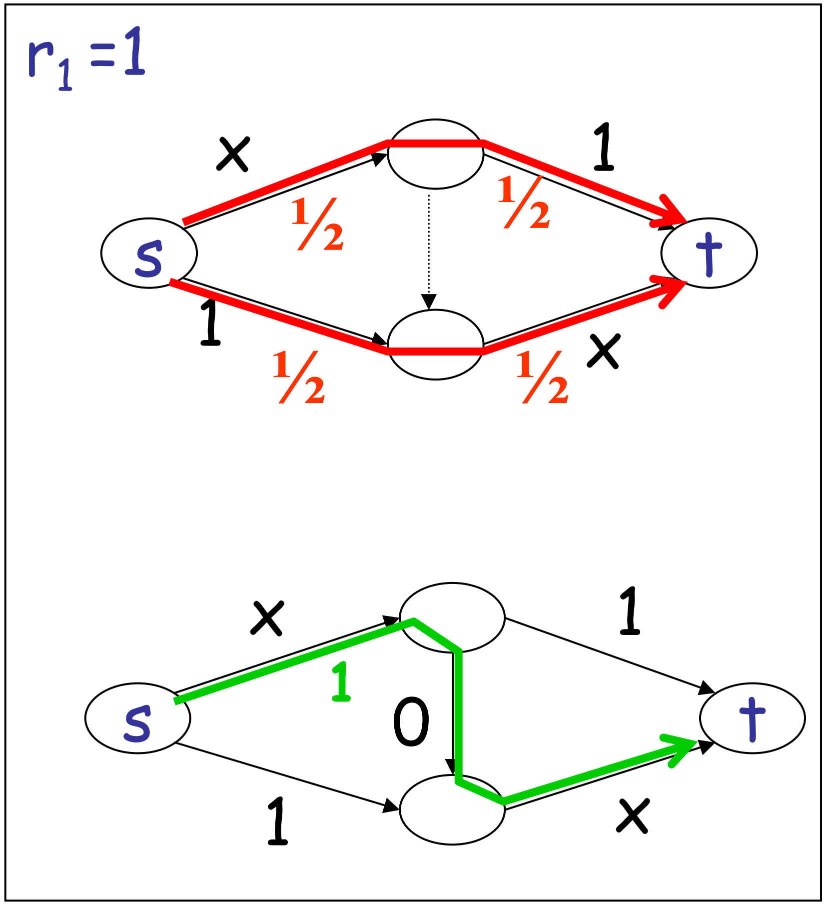

Braess paradox:

- Delay increases from 1.5 to 2
  - due to selfish routing
  - or due to added edge

Selfish Move ≠ Overall Improvement

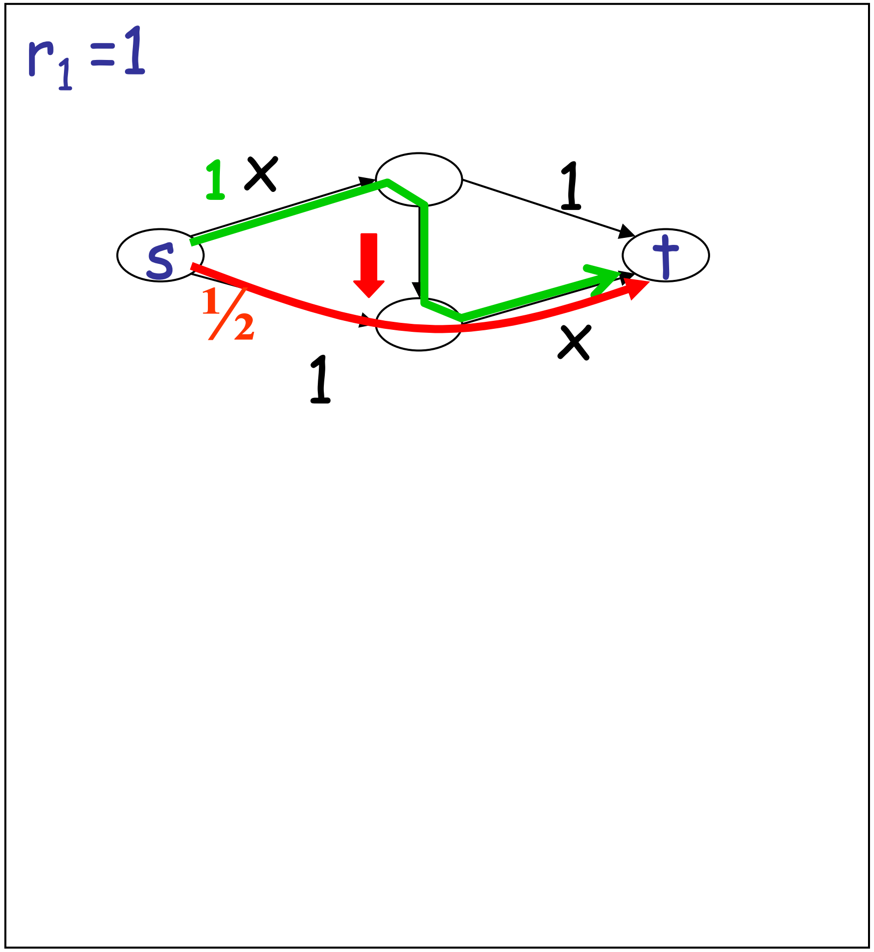

Moving ½ the flow
does not help
selfishly

but move improves
delay for remaining
flow

Selfish Move ≠ Overall Improvement

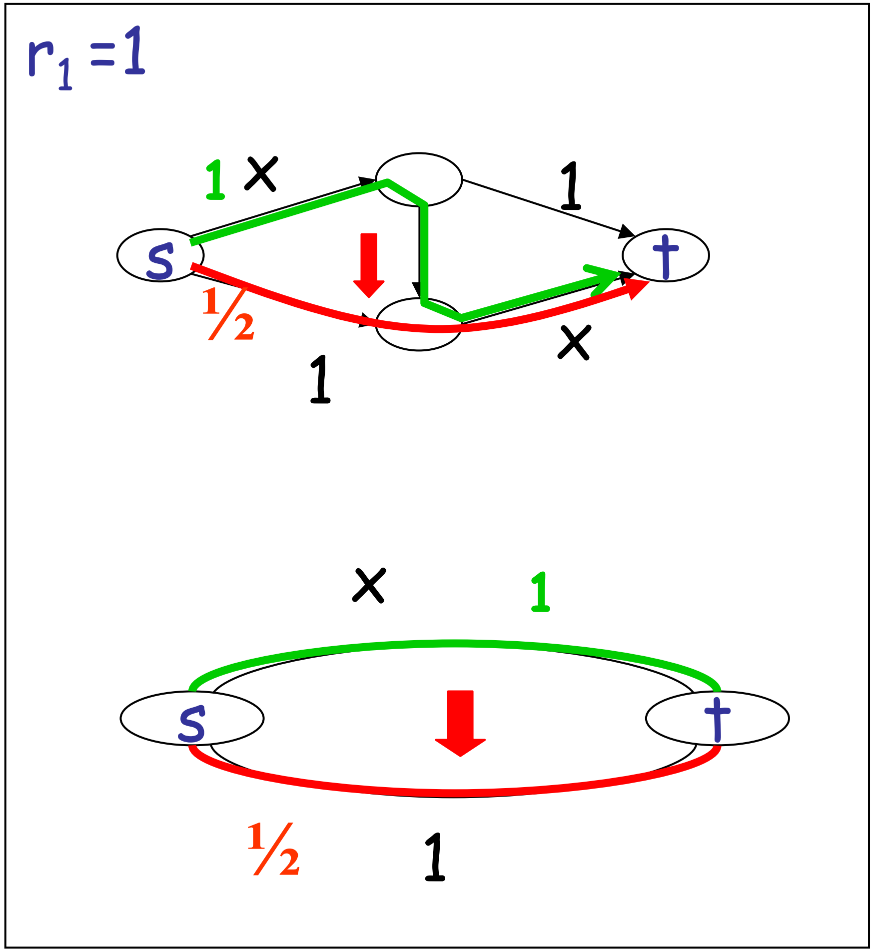

Moving ½ the flow does not help selfishly

but move improves delay for remaining flow

Selfish Move ≠ Overall Improvement

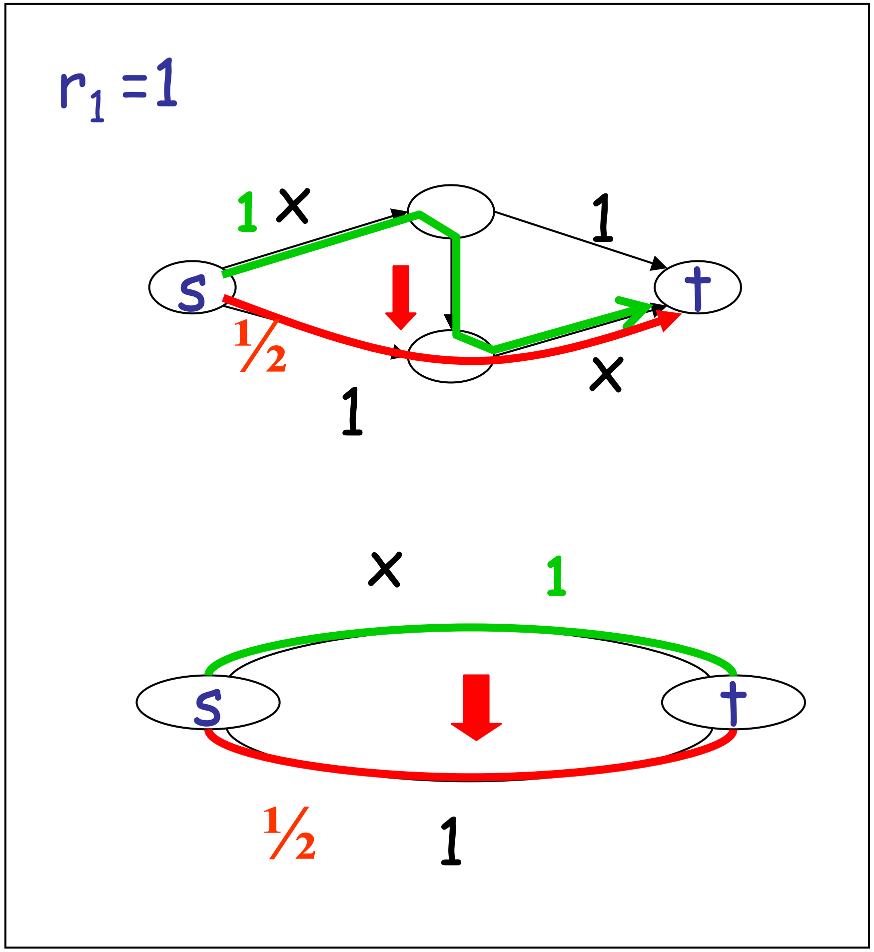

Personal objective:

- $\ell_P(f)$ = sum of latencies of edges along P (w.r.t. flow f)

Assumption: ℓ monotone increasing

Overall objective:

- $C(f)$ = cost or total latency of a flow f:

$$= \Sigma_P f_P \cdot \ell_P(f)$$

can also be written as

$$= \Sigma_e f_e \cdot \ell_e(f_e)$$

Nonatomic Game

Users control an infinitesimally small amount of flow

Nash: all flow path carrying flow are shortest

Facts: Nash is unique.

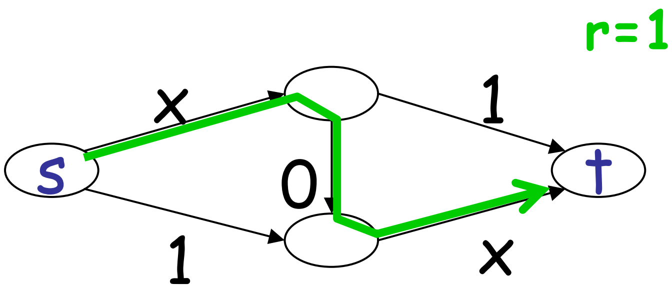

Nonatomic Game

Users control an infinitesimally small amount of flow

Nash: all flow path carrying flow are shortest

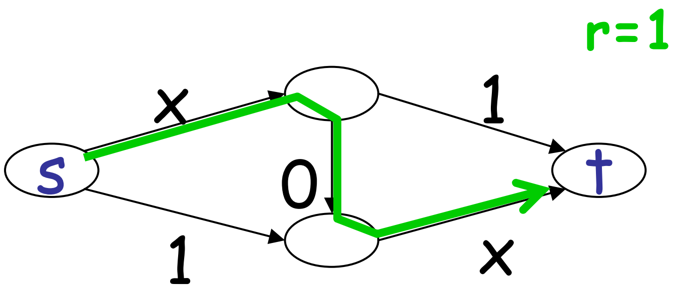

Facts: Nash is unique.

Results e.g.:

– bound Nash/Opt Ratio [Roughgarden '02],

– bicriteria bound [Roughgarden-Tardos '00].

Here: Atomic Game

- Each user controls one unit of flow, and
- selects a single path

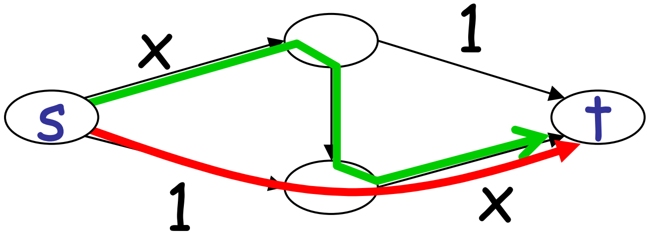

Atomic Game

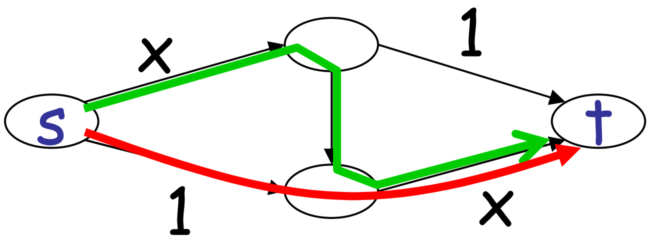

- Each user controls one unit of flow, and
- selects a single path

Selfish moves improve a different function [Monderer Shapley’96, potential game]

$$\Phi(f) = \Sigma_e (\ell_e(1)+\ldots+ \ell_e(f_e))$$

Atomic Game

- Each user controls one unit of flow, and
- selects a single path

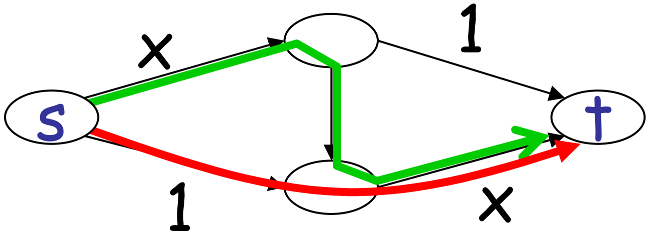

Selfish moves improve a different function [Monderer Shapley'96, potential game]

$$\Phi(f) = \Sigma_e (\ell_e(1)+...+ \ell_e(f_e))$$

Recall global objective is $C(f) = \Sigma_e f_e \cdot \ell_e(f_e)$

Atomic Game Results

Fabrikant, Papadimitriou, Talwar STOC'04

- Single source game: Nash can be found in poly time: minimizing $\Phi(f)$ via min-cost flow
- Multi-source game: Finding Nash is polynomial local search (PLS) complete

Atomic Game Results

Fabrikant, Papadimitriou, Talwar STOC'04

- Single source game: Nash can be found in poly time: minimizing $\Phi(f)$ via min-cost flow
- Multi-source game: Finding Nash is polynomial local search (PLS) complete

Anshelevich-D-K-T-W-R 2004:

- Extend to Network Design Games
  - Congestion (cost ℓ) decreases user cost
- use Φ to bound Best Nash/Opt ratio by O(log k)
  - k=number of players

Why care about

Finding a Nash?

Best Nash/Opt ratio?

Papadimitriou-Koutsoupias '99

Nash = outcome of selfish behavior

$\Rightarrow$ worst Nash/Opt ratio: Price of Anarchy

Nash as Stable Design

Need to Find a Nash equilibrium

- Stable: as no user finds it in their interest to deviate

Need to find a “good” Nash

- Best Nash/Opt ratio? = Price of Stability [ADKTWR 2004]

Games so far are nice…

Service provider game:

- Nash ≈ local optimization

Potential game

- Nash ≈ local optima for a related function

All have deterministic Nash.

No Deterministic Nash

Matching Pennies:

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

No Deterministic Nash

Network Design

[Anshelevich-DTW STOC'03]

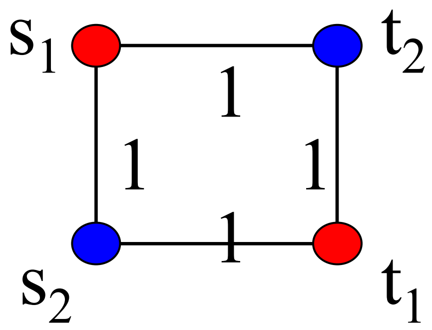

Users bid contribution on individual edges.

- Single source game: Price of Anarchy = 1

No Deterministic Nash

Network Design

[Anshelevich-DTW STOC'03]

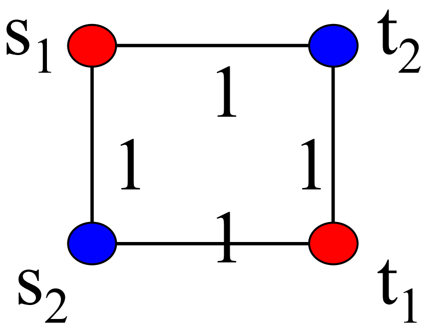

Users bid contribution on individual edges.

- Single source game: Price of Anarchy = 1

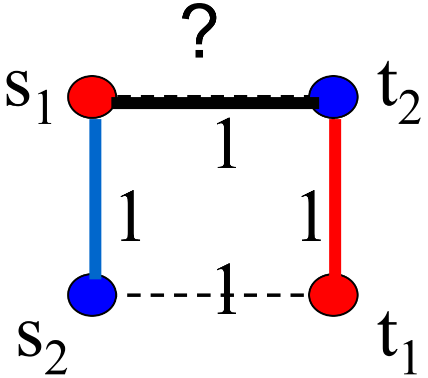

- Multi source: no Nash

Mechanism Design:
How to Design "Nice" Games?

Mechanism Design:
How to Design Games well

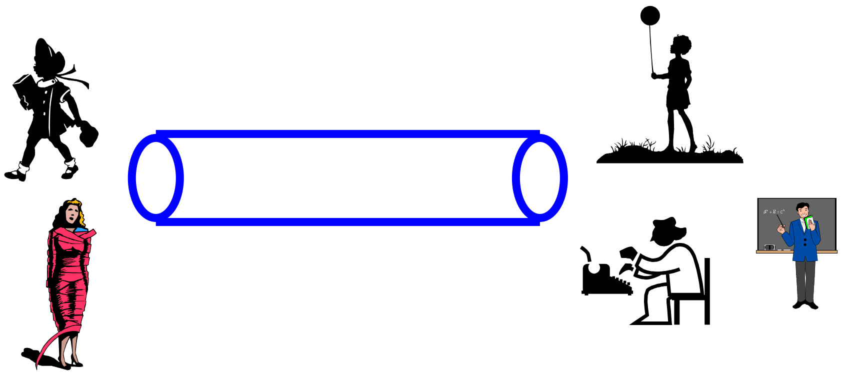

Many Users with diverse utilities for bandwidth.

How should we share a given B bandwidth?

Mechanism Design:
How to Design Games well

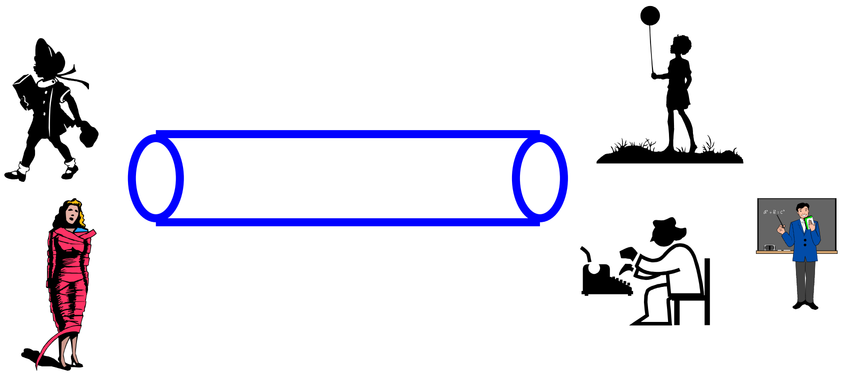

Many Users with diverse utilities for bandwidth.

How should we share a given B bandwidth?

Traditional Mechanism Design (VCG):

- use payments to induce all players to tell us his utility for bandwidth
- Assign bandwidth to maximize social welfare (total utility obtained)

Mechanism Design:
How to Design Games well

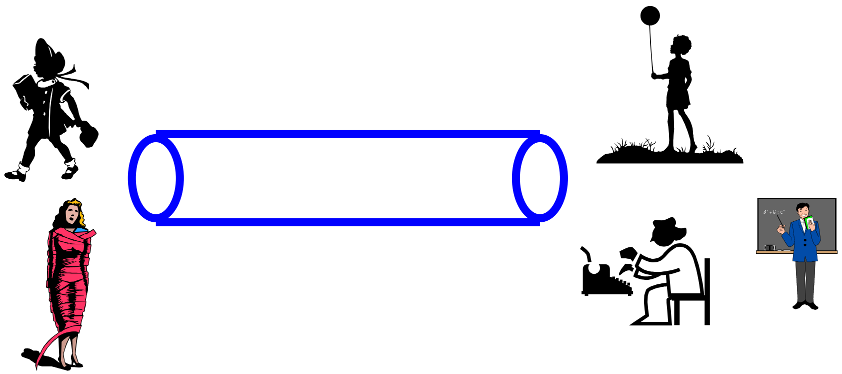

Many Users with diverse utilities for bandwidth.

How should we share a given B bandwidth?

Traditional Mechanism Design (VCG):

- use payments to induce all players to tell us his utility for bandwidth
- Assign bandwidth to maximize social welfare (total utility obtained)

Cost lot of money; lots of information to share

Mechanism Design:
How to Design Games well

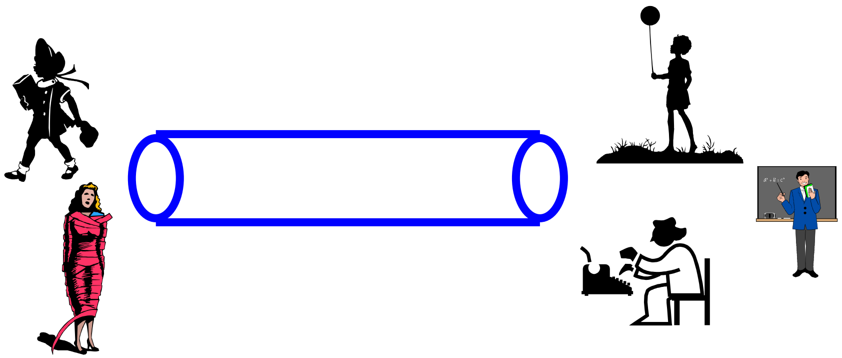

Many Users with diverse utilities for bandwidth.

How should we share a given B bandwidth?

Here:

- design a simple/natural Nash game that divides bandwidth and
- analyze the Prize of Anarchy

Kelly, Johari-Tsitsikis, 2004

Bandwidth Sharing Game

Users have a utility function $U_i(x)$ for receiving x bandwidth.

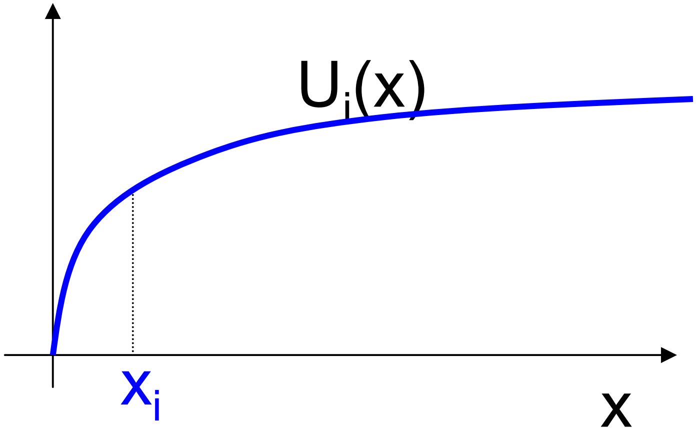

Assume elastic users (concave utility functions)

Bandwidth Sharing Game

Users have a utility function $U_i(x)$ for receiving x bandwidth.

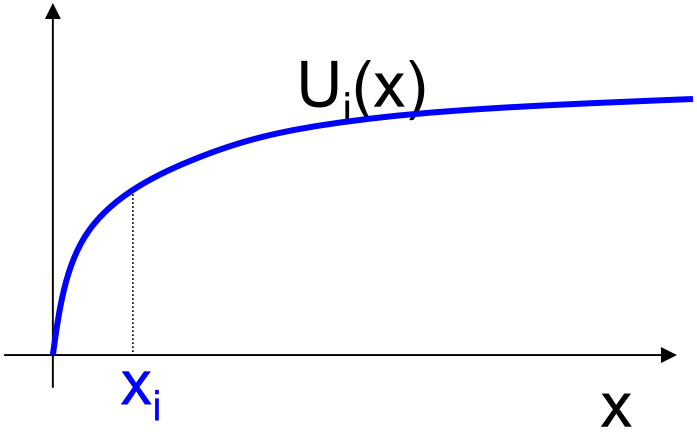

Assume elastic users (concave utility functions)

Given a prize p, users maximize

$\max_x U(x) - px \quad \rightarrow x_i$

User Optimization?

What do users do?

For a given price p:

- solve
- $\max_x U(x) - p \cdot x$
  - $U'(x) = p$

- offer to pay $w = x$

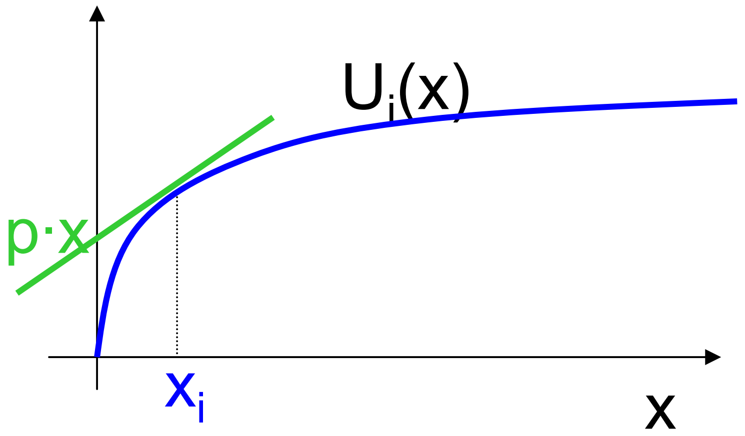

Bandwidth Sharing Game

Users have a utility function $U_i(x)$ for receiving x bandwidth.

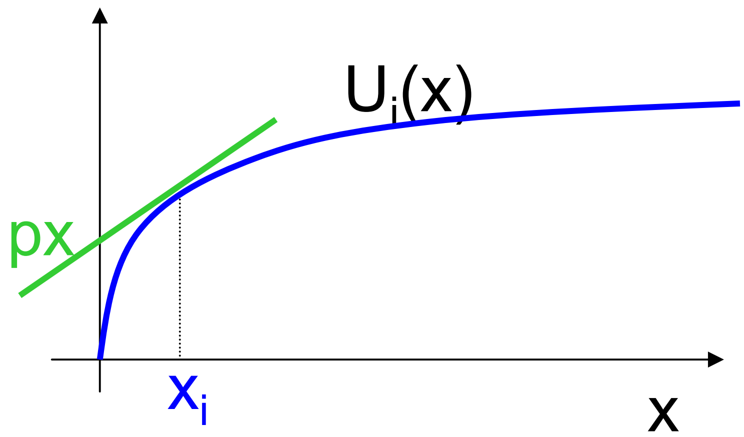

Assume elastic users (concave utility functions)

Given a prize p, users maximize

$\max_x U(x) - px \quad \rightarrow x_i$

Fact: there is a prize p at which selfish user optimization results in optimal allocation.

[convex optimization]

Bandwidth Sharing Game

Users have a utility function $U_i(x)$ for receiving x bandwidth.

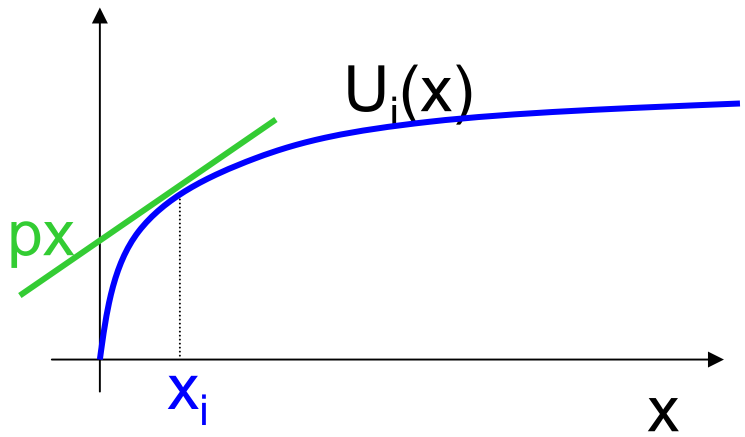

Assume elastic users (concave utility functions)

Given a prize p, users maximize

$\max_x U(x) - px \quad \rightarrow x_i$

Fact: there is a prize p at which selfish user optimization results in optimal allocation.

[convex optimization]

But what is the right prize to charge?

Proportional Allocation Mechanism

Allocation game:

- Players offer money w_i for bandwidth.
- Bandwidth allocated proportional to payments:
  - effective price $p = (\Sigma_i w_i)/B$
  - player allocation $x_i = w_i / p$

Proportional Allocation Mechanism

Allocation game:

- Players offer money w_i for bandwidth.
- Bandwidth allocated proportional to payments:
  - effective price $p = (\Sigma_i w_i)/B$
  - player allocation $x_i = w_i / p$

- Does this game have a Nash equilibrium?
- What is the Price of Anarchy?

User Optimization?

What do users do?

For a fixed price p:

- solve
- $\max_x U(x) - p \cdot x$
  - $U'(x) = p$

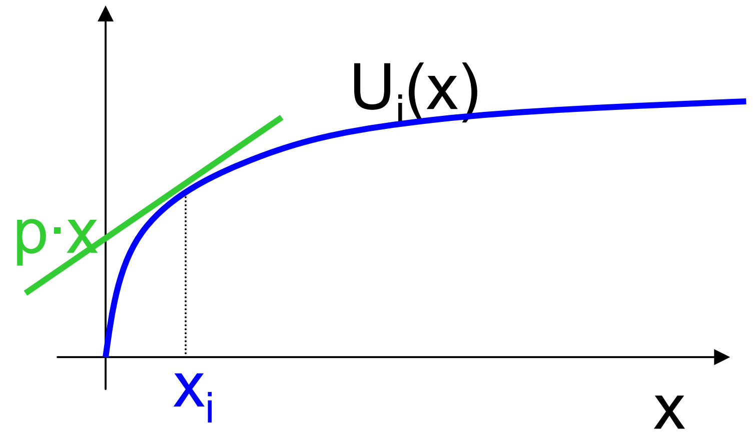

- offer to pay $w = x$

User Optimization?

What do users do?

For a fixed price p:

- solve
- $\max_x U(x) - p \cdot x$
  - $U'(x) = p$

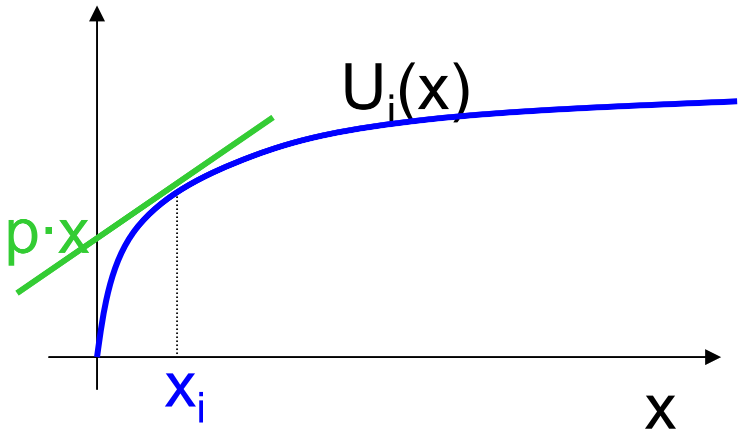

- offer to pay $w = x$

But their bid effects the price!

User Optimization?

Allocation game:

- Players offer money w_i for bandwidth.
- Bandwidth allocated:
  - price $p= (\Sigma_i w_i)/B$
  - allocation $x_i = w_i /p$

Solve optimization problem for best w_i

- given all other players bids w_j

User Optimization?

Allocation game:

- Players offer money w_i for bandwidth.
- Bandwidth allocated:
  - price $p = (\Sigma_i w_i)/B$
  - allocation $x_i = w_i / p$

Solve optimization problem for best w_i

- given all other players bids w_j

Result allocation and price such that

- $U'(x)(1-x/B)=p$

Bandwidth Allocation Results

Kelly'97

For users that are price takers:

- Nash allocation is socially optimal

[Convex optimization]

Johari-Tsitsiklis'04

For users that anticipate the price:

Bandwidth Allocation Results

Kelly'97

For users that are price takers:

- Nash allocation is socially optimal

[Convex optimization]

Johari-Tsitsiklis'04

For users that anticipate the price:

- Nash allocation exists, and
- Worst Price of Anarchy is 4/3.

Bandwidth Allocation: Worst case

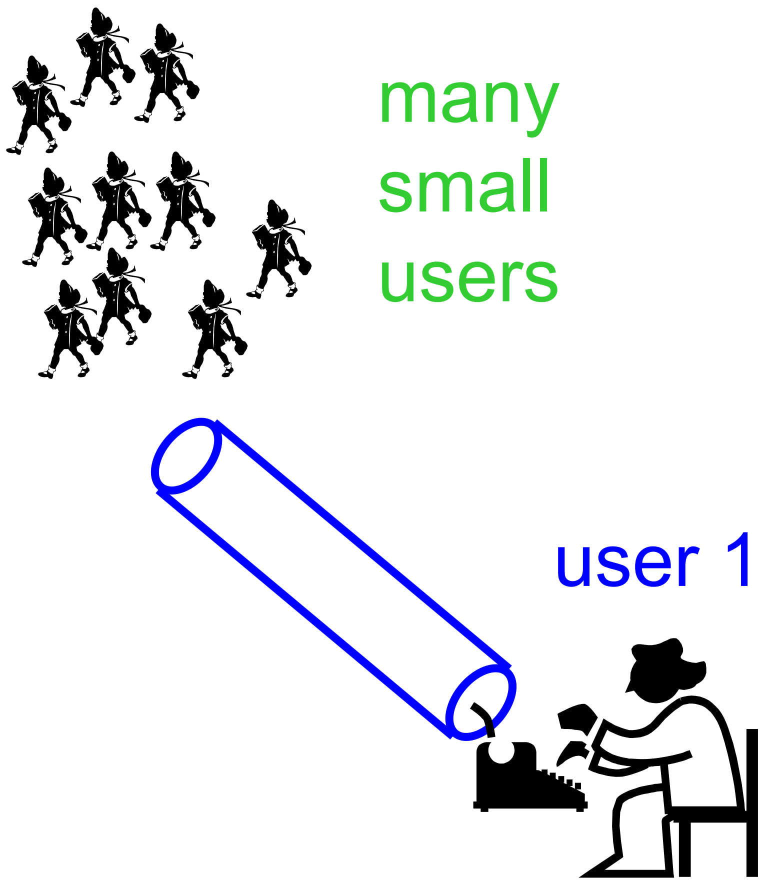

Linear utilities:

User 1 has utility $U_1(x)=x$ for

Many-many other users have utility $U_i(x)=x/2$.

- "Socially optimal" user 1 gets everything
- At Nash: user 1 gets ½ of the bandwidth
  - But gets it cheaper

Bandwidth Allocation: Networks

Both Kelly and Johari-Tsitsiklis results extend to networks

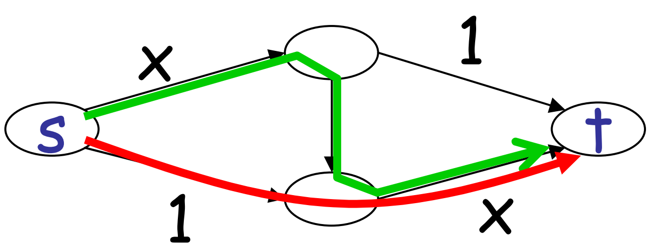

Users bid payments for each edge separately!

- fixed routing, or
- allowing users to split traffic between multiple paths

Summary

We talked about

- Some natural Network Games

  facility location, routing

Summary

We talked about

- Some natural Network Games

  facility location, routing

- Price of Anarchy, local search and potential games

Summary

We talked about

- Some natural Network Games

  facility location, routing
- Price of Anarchy, local search and potential games
- Designing games (mechanism design)

  bandwidth allocation

Connections to Algorithms, I

Local search

– Finding local opt $\leftrightarrow$ finding Nash
  - In STOC'04:

    [Fabrikant-Papadimitriou-Talwar] and [Gairing-Lücking-Mavronicolas-Monien]

– Quality of local search $\leftrightarrow$ Price of Anarchy
  - Approximation via local search (service provider game)
  - Connection to potential function Φ

Connections to Algorithms, II

Convex Optimization

– Price in Kelly's bandwidth allocation game

Pricing and primal-dual algorithms

– Market equilibrium

In STOC'04: [Garg-Kapoor], [Devanur-Vazirani]

Algorithmic Game Theory

- The main ingredients:
  - Lack of central control

    like distributed computing
  - Selfish participants

    game theory
- Common in many settings e.g., Internet
- Exciting area with many open problems:
  - Cost of anarchy in other network games
  - Design games with low cost of anarchy

See Proceedings for references