

A Game Theoretic View of Efficiency Loss in Network Resource Allocation

John N. Tsitsiklis
(w. Ramesh Johari and Shie Mannor)

Massachusetts Institute of Technology

March 2004

Agenda

Allocation of divisible resources
(e.g., bandwidth in a network)

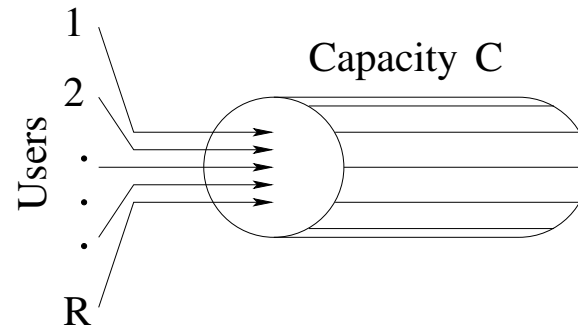
Simple “mechanisms”
(based on market clearing)

Study efficiency loss, w.r.t. social optimum
(in the presence of selfish users)

Outline


1. Single link, fixed capacity
2. Network case, fixed capacities
3. Elastic capacities

PART I: Single Link, Fixed Capacity



Rate d_r \longrightarrow utility $U_r(d_r)$

U_r : concave, strictly increasing, nonnegative



The Social Optimum

$$\begin{array}{ll} \text{maximize} & \sum_r U_r(d_r) \\ \text{subject to} & \sum_r d_r \leq C \\ & \mathbf{d} \geq 0 \end{array}$$

A Pricing Mechanism

User r submits a bid w_r .

Receives bandwidth: $\frac{w_r}{w_1 + \dots + w_R} C$

Example: $w_1 = 2$ $d_1 = 2C/5$
 $w_2 = 3$ $d_2 = 3C/5$


All bandwidth is allocated

Unit price of bandwidth: $\mu = \frac{1}{w_1 + \dots + w_R} C$

$$d_r = \frac{w_r}{\mu}$$

“Supply = Demand” Interpretation

User r submits w_r : same as submitting “demand curve”: $d_r = \frac{w_r}{\mu}$



Mechanism clears the market:

$$\frac{w_1}{\mu} + \dots + \frac{w_R}{\mu} = \text{total demand} = \text{supply} = C$$

Users as Price Takers

Given price μ , user r solves:

$$\max_{w_r \geq 0} U_r \left(\frac{w_r}{\mu} \right) - w_r$$

Theorem 1 (Existence of Competitive Equilibrium; Kelly, 1997)

There exist w and μ such that:

(a) *w_r is optimal for user r given the price μ .*

(b) *$(w_1 + \dots + w_R)/\mu = C$.*

The resulting allocation is socially optimal.

Users as Price Anticipators

Suppose users *know* the price setting procedure.

Given $(w_s, s \neq r)$, user r solves:

$$\max_{w_r \geq 0} U_r \left(\frac{w_r}{w_r + \sum_{s \neq r} w_s} C \right) - w_r$$

This is now a *game*, where the strategy of user r is the bid w_r .

Interested in Nash equilibria

Example

$$C = 1, \quad U_1(d_1) = 2d_1, \quad U_2(d_2) = d_2.$$

Social optimum: $d_1 = 1, d_2 = 0$

Price-taking equilibria: $\mu = 1$

Price-anticipating users:

(a) $\mu > 1 \implies w_2 = 0 \implies w_1 = ?$

(b) $\mu = 1, w_2 > 0$; user 2 will reduce w_2 , and reduce the price

(c) $\mu < 1, w_1 > 0, w_2 > 0$: inefficient


Nash Equilibrium

Theorem 2 (Hajek & Gopal, 2002) Assume $R > 1$.

There exists a *unique* Nash equilibrium w .

The resulting allocations d_r are the unique socially optimal solution for *modified utilities*:

$$\hat{U}_r(d_r) = \left(1 - \frac{d_r}{C}\right) U_r(d_r) + \left(\frac{d_r}{C}\right) \left(\frac{1}{d_r} \int_0^{d_r} U_r(z) dz\right)$$



Efficiency Loss

Theorem 3 *The efficiency loss is no more than 25%:*

$$\text{(Nash eq. utility)} \geq \frac{3}{4} \times \text{(socially optimal utility)}$$

Furthermore, this bound is tight.

Worst case:

Many users, linear utility functions, one “dominant” user

The Worst Case

$$C = 1$$

$$U_1(d_1) = d_1$$

$$U_r(d_r) \approx d_r/2$$

$$R \rightarrow \infty$$

$$d_1^G \rightarrow 1/2$$

$$\sum_{r>1} d_r^G \rightarrow 1/2$$

$$\mu \rightarrow 1/2$$

$$\sum_r U_r(d_r^S) = 1$$

$$\sum_r U_r(d_r^G) \rightarrow 3/4$$

Summary of Properties

1. Strategy = a “smooth” demand curve,
chosen from a 1-parameter family ($d = w/\mu$)
2. Price is set by market clearing (supply = total demand)
3. Players’ demand is always nonnegative.
4. Price taking behavior \implies full efficiency.
5. Players’ payoffs are **concave** when **price anticipating**.

Characterization theorem:

Out of all mechanisms with the above properties,
the one we have studied minimizes the worst-case efficiency loss!

PART II: Networks

Link j has capacity C_j

Each **user** is identified with a **path**

Social optimum:

$$\begin{array}{ll} \text{maximize} & \sum_r U_r(d_r) \\ \text{subject to} & \text{capacity constraints} \end{array}$$

The Pricing Mechanism

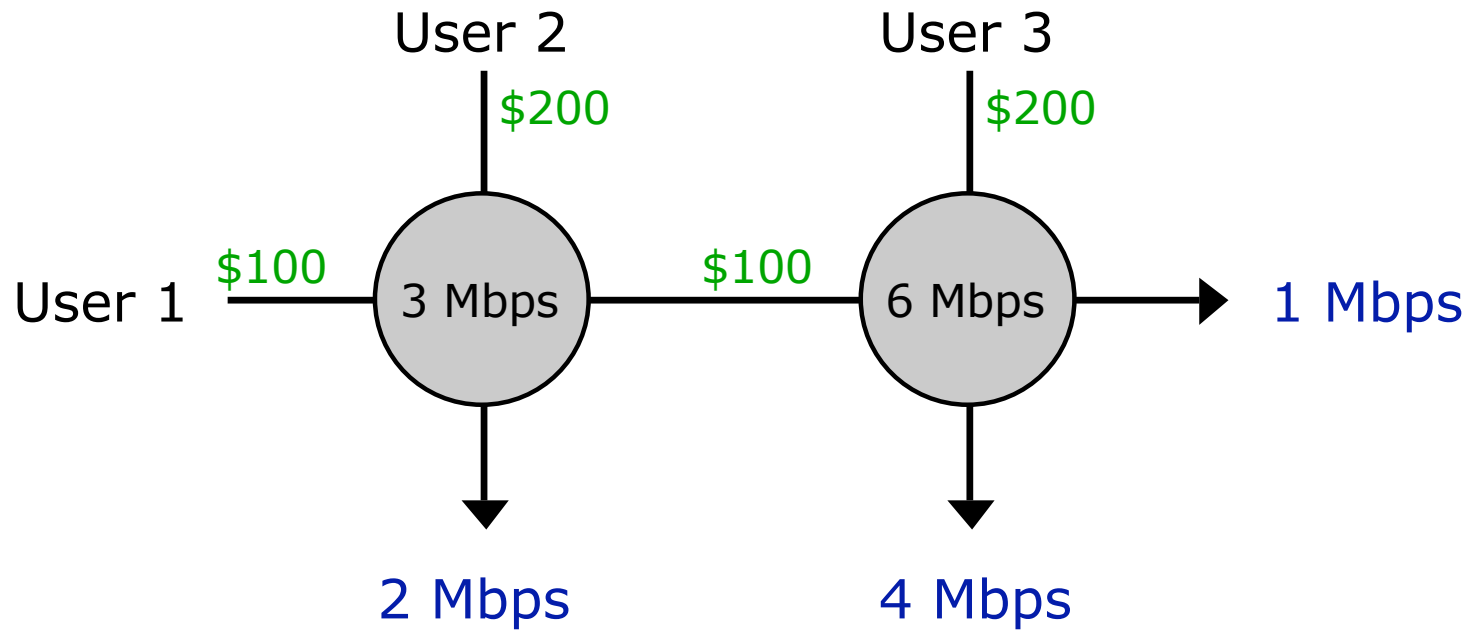
User r submits a bid w_r^j , at each link j

Receives bandwidth at that link: $x_r^j(\mathbf{w}) = \frac{w_r^j}{w_1^j + \dots + w_R^j} C_j$


User r sends as much as possible: $d_r = \min_{j \in \text{path}} x_r^j$

User r payoff: $U_r(d_r(\mathbf{x}_r(\mathbf{w}))) - \sum_j w_r^j$ **Concave!**

Example



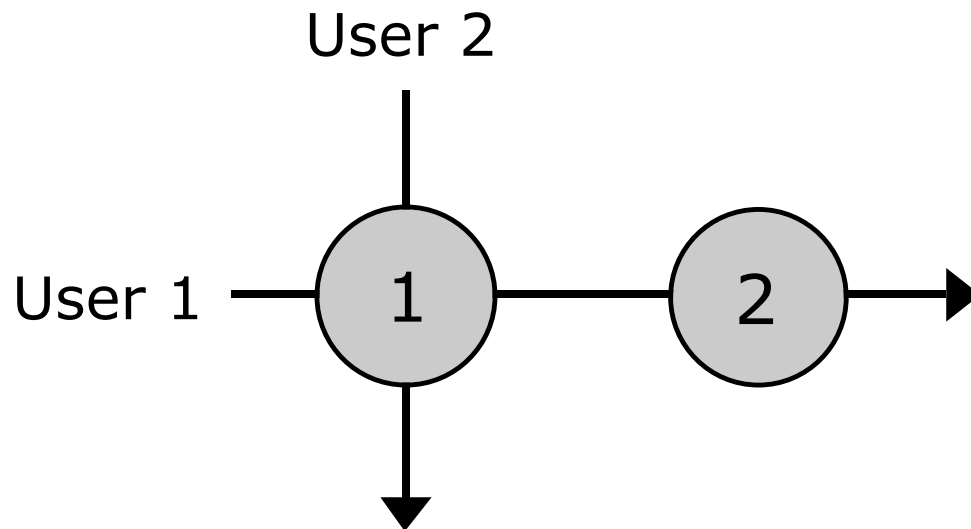
Nash Equilibrium?



User 1 would like to submit infinitesimally positive bid to link 2

No Nash equilibrium

Nash Equilibrium?



Solution: User 1 says: "Please give me C_2 for free"

An Extended Game

Each user r also submits a request ϕ_r^j for **free** rate at link j

Requests are:

considered: if all bids to that link are 0

granted: if sum of requests does not exceed the capacity

Existence of a Nash Equilibrium

Theorem 4 *With “price-anticipating” users, there exists a Nash equilibrium (w, ϕ) for the extended game.*

NE of original game maps to a NE of extended game

Proof idea:

- Introduce “virtual user” bidding $\epsilon > 0$ at each j .
- With $\epsilon > 0$, no discontinuities, NE exists (Rosen’s theorem)
- Take the limit as $\epsilon \rightarrow 0$
- Perturbed game allocations $x_r^j(\epsilon)$ lead to rate requests ϕ_r^j

General Networks: Efficiency Loss

Theorem 5 *The efficiency loss is no more than 25%:*

$$(\text{Nash eq. utility}) \geq \frac{3}{4} \times (\text{socially optimal utility})$$

Proof idea:

Reduce to analyzing multiple single link games, one at each $j \in J$.

Back to Congestion Control

Faster time scale

Fixed bids w : network protocol allocates rates

Slower time scale

User observes link prices and adjusts bids

Implicit in our model:

- (a) Observe individual link prices
- (b) Can anticipate the effect of a bid change on prices
- (c) Do not anticipate changes (reactions) in bids of other users

A Variation of Congestion Control (Kelly et al.)

Faster time scale:

Fixed **total** bid w_r from each user

Network sets link prices and allocates payments to links

Slower time scale:

(a) Observe **sum of link prices** along its path

(b) Anticipate the effect of total bid change

[If not, social optimality]

(c) Do not anticipate changes (reactions) in bids of other users

Requires less information on prices

Requires more intelligent congestion control algorithm

Anticipating the effects is nontrivial

Existence, uniqueness, and efficiency losses: unresolved

PART III: Elastic Capacities

Instead of a fixed capacity, suppose the link can allocate any rate.

f : total rate allocated

$C(f)$: monetary cost of allocating rate f ; $C(0) = 0$

Social optimization:

$$\begin{array}{ll} \text{maximize} & \sum_r U_r(d_r) - C\left(\sum_r d_r\right) \quad (\text{social welfare}) \\ \text{subject to} & \mathbf{d} \geq 0 \end{array}$$

Assume $C'(f)$ is strictly convex, increasing, with $C'(0) = 0$.

Socially optimal price μ : $C'(f) = \mu = U'_r(d_r)$, if $d_r > 0$
("marginal cost pricing")

The Pricing Mechanism

User r submits a bid w_r .

Network mgr. “clears the market” by choosing supply f and price μ so that:

$$\mu = C'(f) \quad (\text{marginal cost pricing})$$

$$\sum_r \frac{w_r}{\mu} = f \quad (\text{aggregate demand equals supply})$$

Theorem 6 (Kelly, Maulloo, Tan, 1998)

Assume *price taking* users.

There exists a *competitive equilibrium* w and μ .

Equilibrium implies social optimality

Price Anticipating Users

Suppose users anticipate the effect of their bid on the price

Theorem 7 *There exists a Nash equilibrium w .*

Theorem 8 *The efficiency loss is no more than approximately 34%:*

(social welfare at Nash eq.) $\geq (4\sqrt{2}-5)$ (optimal social welfare).

Furthermore, this bound is tight.

Worst case: Linear utility functions, piecewise linear marginal cost.

Additional Results

- If $C(f) = f^B$, then as $B \rightarrow \infty$, efficiency loss $\rightarrow 25\%$.
(Intuition: *as if* the resource is in inelastic supply.)
- Results generalize to network context.

Interesting Directions

Dynamics of convergence to equilibrium

Non-network contexts

- auctions for divisible goods
- games between suppliers (e.g., electric power)

“Complexity” versus efficiency guarantees