Game Theory for Resource Sharing in Large Distributed Systems

Corinne Touati, INRIA

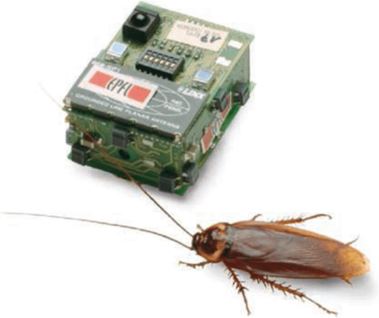

Robot Cockroach Tests Insect Decision-Making Behavior
(EPFL / ULB, Science 16 November 2007, Vol. 318. no. 5853, p. 1055)

Optimality of Multi-user system

- Optimality of a single user

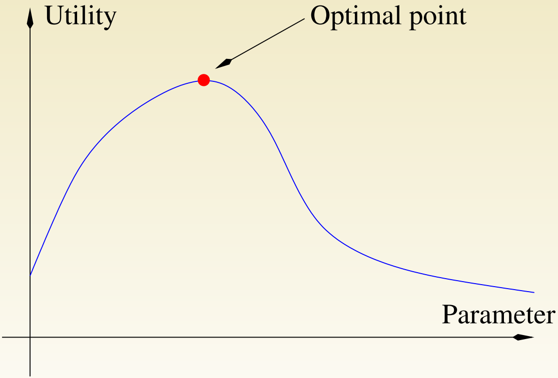

- Situation with multiple users

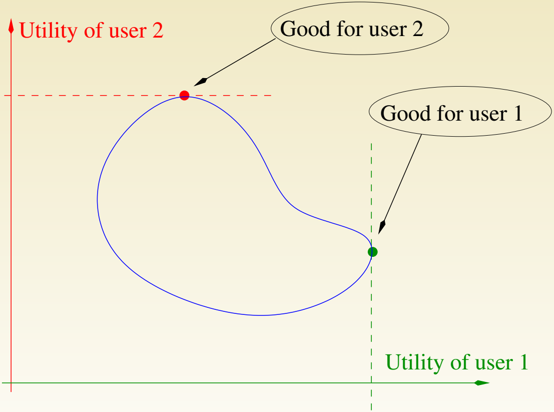

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

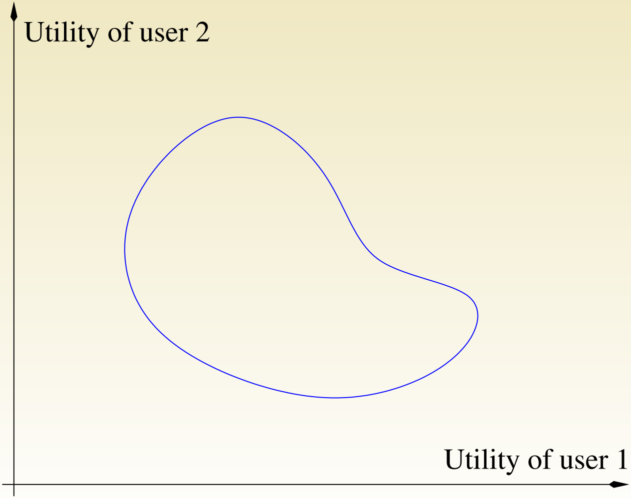

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

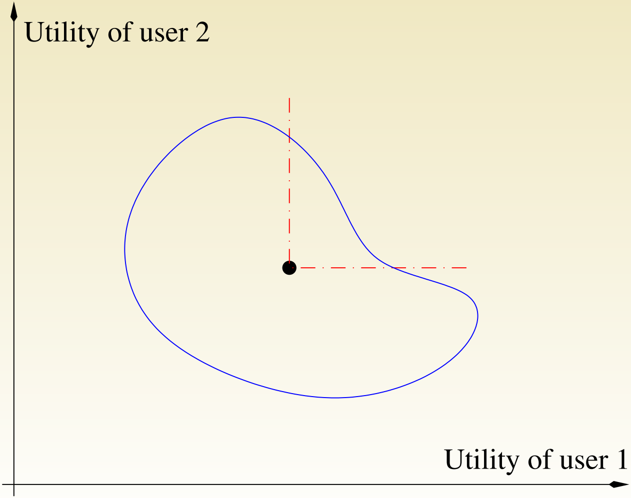

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

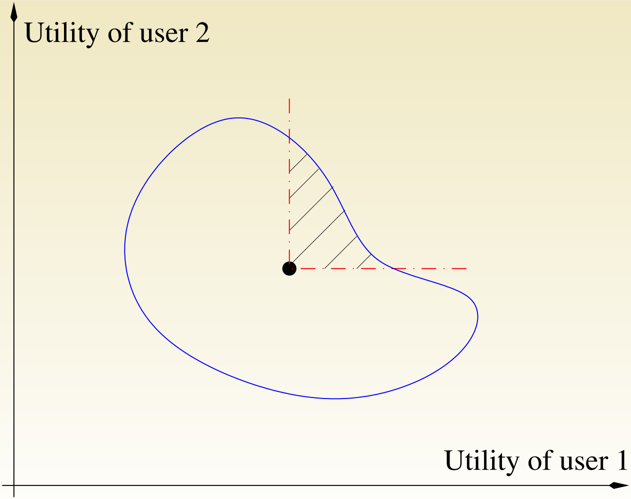

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

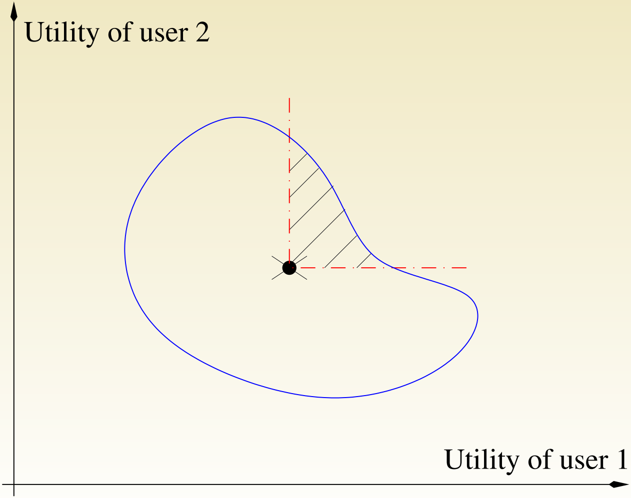

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

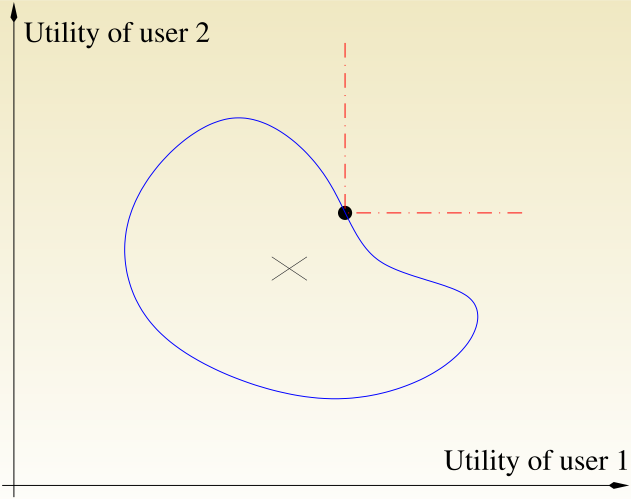

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

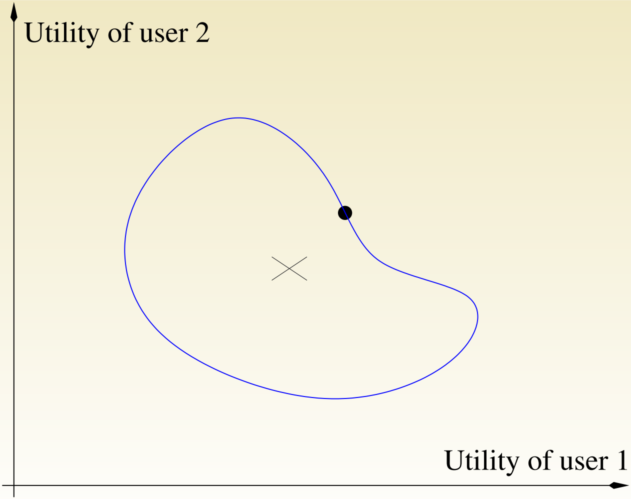

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

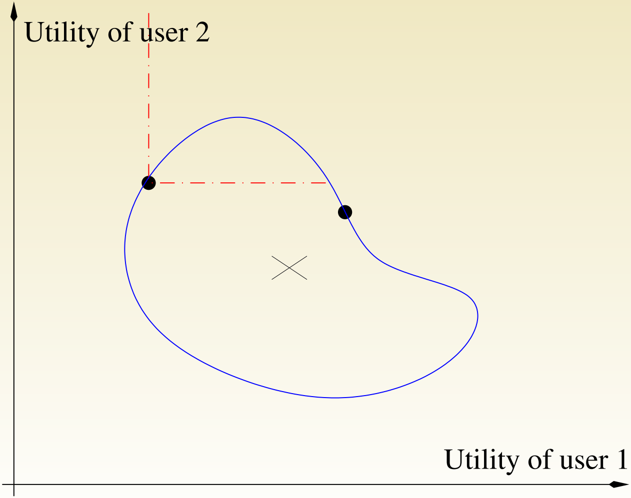

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

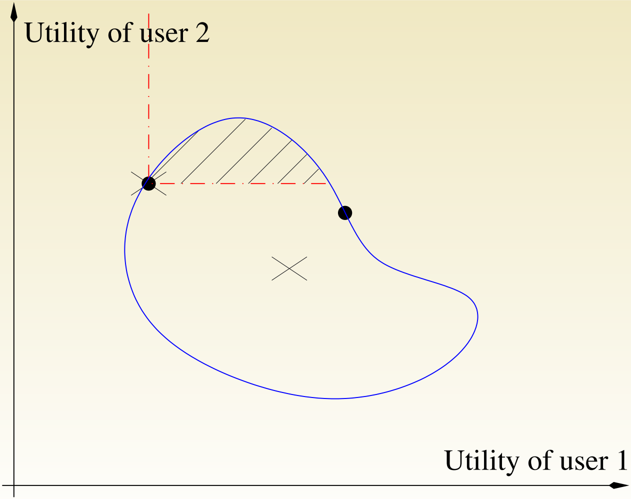

Definition.

A point is Pareto optimal if it cannot be strictly dominated

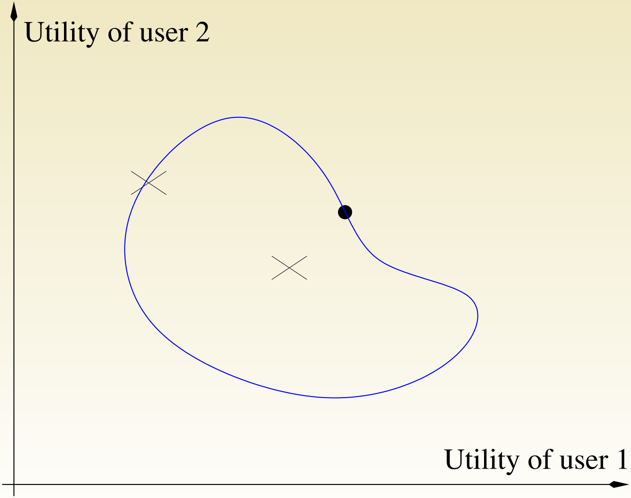

Optimality of Multi-user system

Definition.

A point is Pareto optimal if it cannot be strictly dominated

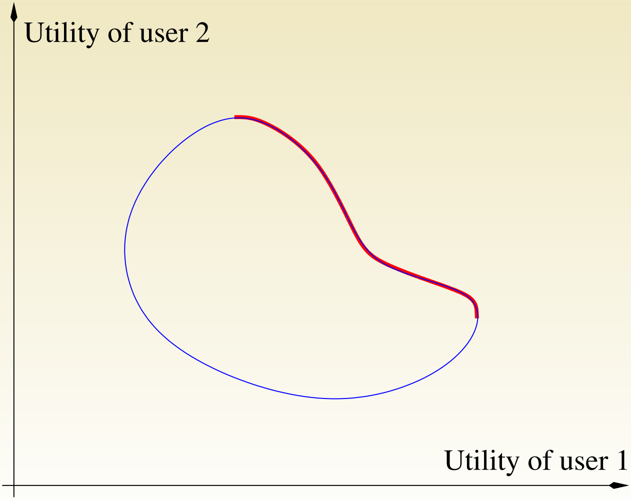

Cooperating or being selfish?

Cooperative games	Non-cooperative games
Institution setting rules and penalties to inforce them	Individual behavior converge (or not) to an equilibrium

Example: Routing intersection:

- Cooperative approach: set of roadsigns (traffic lights, "stop signs"...) inforced by the police
- Non-cooperative approach: everyone tries to cross it as quickly as possible

Outline

Outline

Nash equilibria : definition

Definition

In a non-cooperative setting, each player makes a decision so as to maximize its own return.

Nash equilibria

In a Nash equilibrium, no player has incentive to unilaterally modify his strategy.

strategy (choice)

utility

s^* is a Nash equilibrium iff:

$$\forall p, \forall s_p, u_p(s_1^*, \ldots, s_p^*, \ldots s_n^*) \geqslant u_p(s_1^*, \ldots, s_p, \ldots, s_n^*)$$

Nash Equilibria: definition (cont.)

Pros

- Intuitive

Cons

Nash Equilibria: definition (cont.)

Pros

- Intuitive
- Easy to implement

Cons

Nash Equilibria: definition (cont.)

Pros

- Intuitive
- Easy to implement

Cons

- No guaranty of existence / unicity

Nash Equilibria: definition (cont.)

Pros

- Intuitive
- Easy to implement

Cons

- No guaranty of existence / unicity
- difficult to compute analytically (fixed points)

Nash Equilibria: definition (cont.)

Pros

- Intuitive
- Easy to implement

Cons

- No guaranty of existence / unicity
- difficult to compute analytically (fixed points)
- usually not Pareto optimal

Various contexts:

- Load balancing systems
  Users decide which server to send their request so as to minimize their average delay.

Various contexts:

- Load balancing systems
  Users decide which server to send their request so as to minimize their average delay.
- Wireless systems
  Users decide what power to use so as to maximize a compromize between the transfer rate and the battery usage.

Various contexts:

- Load balancing systems
  Users decide which server to send their request so as to minimize their average delay.
- Wireless systems
  Users decide what power to use so as to maximize a compromize between the transfer rate and the battery usage.
- Pricing systems
  Providers choose their prices so as to maximize their revenue, which is a function of their charged price and their infrastructure cost and market share.

Nash equilibria: Applications

Various contexts:

- Load balancing systems
  Users decide which server to send their request so as to minimize their average delay.
- Wireless systems
  Users decide what power to use so as to maximize a compromize between the transfer rate and the battery usage.
- Pricing systems
  Providers choose their prices so as to maximize their revenue, which is a function of their charged price and their infrastructure cost and market share.
- Queuing systems
  Users optimize their "power" defined as the ratio of their throughput and their expected delay.

Nash equilibria: Application to scheduling of bg-of-task applications

Two computers /
two applications

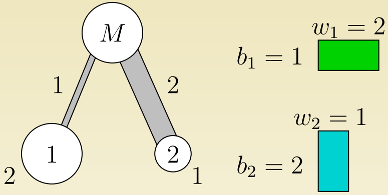

Nash equilibria: Application to scheduling of bg-of-task applications

Two computers / two applications

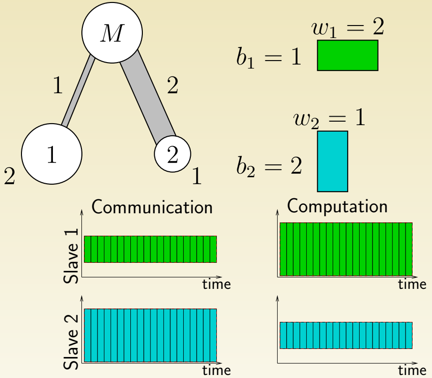

Cooperative Approach:

Application i is processed exclusively on computer i.

$\alpha_1^{(\text{coop})} = \alpha_2^{(\text{coop})} = 1.$

Nash equilibria: Application to scheduling of bg-of-task applications

Two computers / two applications

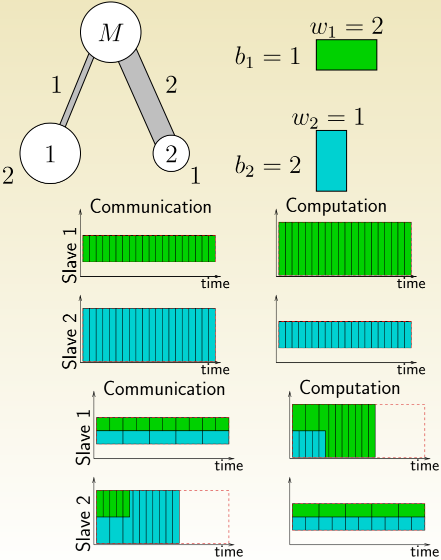

Cooperative Approach:

Application i is processed exclusively on computer i.

$\alpha_1^{(\text{coop})} = \alpha_2^{(\text{coop})} = 1.$

Non-Cooperative Approach:

$$\alpha_1^{(nc)} = \alpha_2^{(nc)} = \frac{3}{4}$$

Hypothesis: packets select their routes of travel from an origin to a destination so as to minimize their own travel cost.

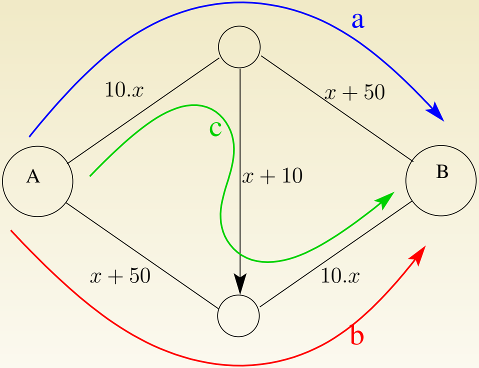

- 3 possible routes
- cost of links are proportional to the fraction of users x passing through it.

Difference with the previous example?

Hypothesis: packets select their routes of travel from an origin to a destination so as to minimize their own travel cost.

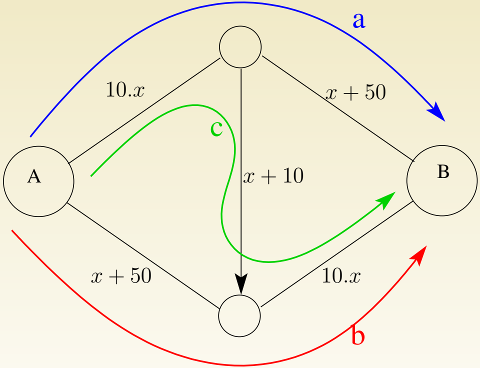

- The number of users is infinite
- Each of them has a negligible impact

Belongs to the class of "population games"

Population games

Definition: Population game.

- Q non atomic populations, each of them of mass $\hat{d}_q$
- A finite set of strategies for each population
- A strategy distribution $y = (y_1, ..., y_Q)$, where y_q is a vector containing the masses of the subset of population q adopting each possible strategy
- The marginal payoff per unit of class i of population q: $F_q^i(y)$

Definition: Wardrop equilibrium.

A state $\hat{y}$ is a Wardrop equilibrium if, for any population:

- All strategies being used by members of the population yelf the same marginal payoff: $\forall i, j, y_q^i \neq 0, y_q^j \neq 0, F_q^i(\hat{y}) = F_q^j(\hat{y})$
- The marginal payoff associated to all strategies actually used by members is lower than it would be with any of the strategies not chosen.

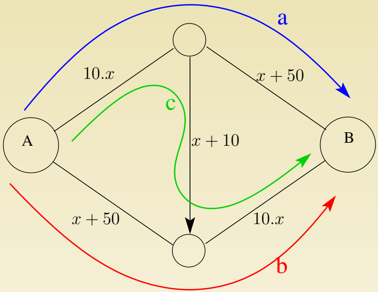

- 1 population ($Q = 1$), 3 possible strategies
- Strategy distribution $y = (y_1)$ with $y_1 = (m_1, m_2, m_3)$
- Marginal payoff per unit:

$$
\begin{aligned}
F_1^1(y) &= 10 * (m_1 + m_3) + (m_1 + 50) \\
&= 11.m_1 + 10.m_3 + 50 \\
F_1^2(y) &= (m_2 + 50) + 10 * (m_2 + m_3) \\
&= 11.m_2 + 10.m_3 + 50 \\
F_1^3(y) &= 10 * (m_1 + m_3) + (m_2 + 10) + 10 * (m_2 + m_3) \\
&= 10.m_1 + 20.m_3 + 11.m_2 + 10
\end{aligned}
$$

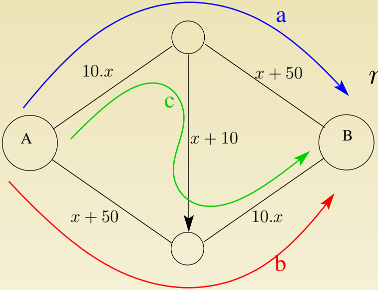

Total population mass $m_1 + m_2 + m_3 = 6$ and:

$$\begin{aligned} F_1^1(y) &= 11.m_1 + 10.m_3 + 50 \\ F_1^2(y) &= 11.m_2 + 10.m_3 + 50 \\ F_1^3(y) &= 10.m_1 + 11.m_2 + 20.m_3 + 10 \end{aligned}$$

Let $\hat{y}$ be the strategy distribution at the Wardrop equilibria. Then,

$$\forall i, j, m_i \neq 0, m_j \neq 0, F_1^i(y) = F_1^j(y),$$

and

$$\forall i, j, m_i \neq 0, m_j = 0, F_1^i(y) < F_1^j(y).$$

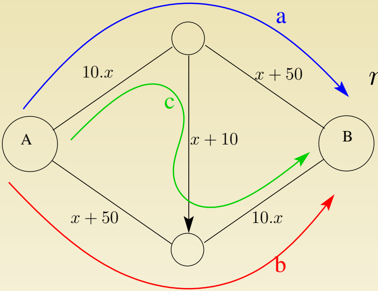

Total population mass $m_1 + m_2 + m_3 = 6$ and:

$$\begin{aligned} F_1^1(y) &= 11.m_1 + 10.m_3 + 50 \\ F_1^2(y) &= 11.m_2 + 10.m_3 + 50 \\ F_1^3(y) &= 10.m_1 + 11.m_2 + 20.m_3 + 10 \end{aligned}$$

Suppose only routes "a" and "b" are used ($m_3 = 0$), then

$$m_1 = 3 \text{ and } F_1^1(y) = F_1^2(y) = 83.$$

But the single cost of a packet going through path "c" would be

$$10.m_1 + 11.m_2 + 10 = 73 < F_1^1(y),$$

hence $(m_1.m_2 \neq 0) \Rightarrow m_3 \neq 0$.

Wardrop equilibria: application

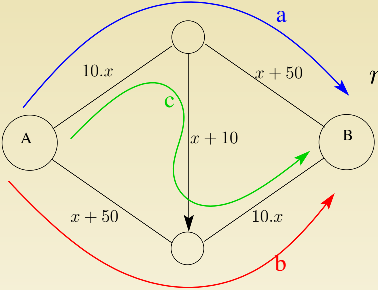

Total population mass $m_1 + m_2 + m_3 = 6$ and:

$$
\begin{aligned}
F_1^1(y) &= 11.m_1 + 10.m_3 + 50 \\
F_1^2(y) &= 11.m_2 + 10.m_3 + 50 \\
F_1^3(y) &= 10.m_1 + 11.m_2 + 20.m_3 + 10
\end{aligned}
$$

With similar arguments, we can show that $m_1.m_2.m_3 \neq 0$.
Hence $F_1^1(y) = F_1^2(y) = F_1^3(y)$.
Then $m_1 = m_2$
and $11.m_1 + 10.m_3 + 50 = 21.m_1 + 20.m_3 + 10$,
hence $40 = 10.m_3 + 10m_1$.
Finally $m_1 = m_2 = m_3 = 2$ and $F_1^1(\hat{y}) = F_1^2(\hat{y}) = F_1^3(\hat{y}) = 92$.

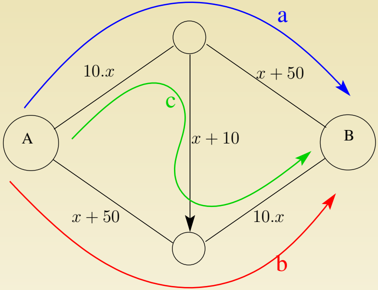

There is actually a simpler way :)

Potential games

Here, potential function $\Phi(m_1, m_2, m_3) = \sum_{l \text{ links}} \int_0^{\alpha_l} c_l(u)du$

with $\alpha_l = \sum_{p \text{ paths}} m_l \delta_{l,p}$, $\delta_{l,p} = \begin{cases} 1 \text{ if flow } l \text{ goes through path } p \\ 0 \text{ overwise} \end{cases}$,

and c_l the cost of crossing link l.

Then the Wardrop equilibria is the solution of:

$$\hat{m} = (\hat{m}_1, \hat{m}_2, \hat{m}_3), \text{argmin } \Phi(m) \text{ subject to } \sum m_i = 6.$$

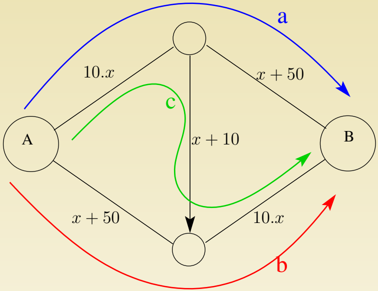

Important remark

We saw that $F_1^1(\hat{y}) = F_1^2(\hat{y}) = F_1^3(\hat{y}) = 92$.

But also that, if only routes "a" and "b" were used ($m_3 = 0$), then

$$m_1 = 3 \text{ and } F_1^1(y) = F_1^2(y) = 83.$$

(But the cost of a single packet going through path "c" would be $73 < F_1^1(y)$).

Outline

Context: urban transportation networks.
Hypothesis: travelers select their routes of travel from an origin to a destination so as to minimize their own travel cost or travel time.

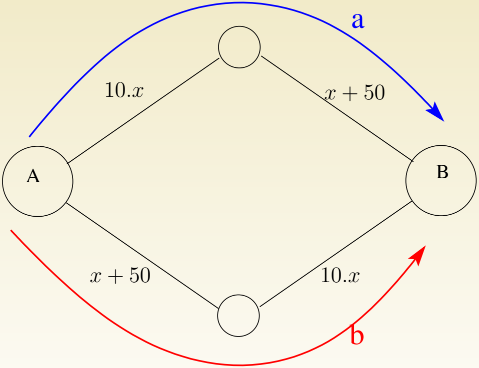

Rate: 6

With 2 roads,
$Cost_a = Cost_b = 83$

Context: urban transportation networks.
Hypothesis: travelers select their routes of travel from an origin to a destination so as to minimize their own travel cost or travel time.

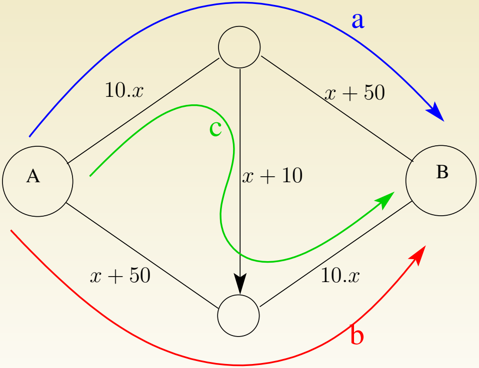

Rate: 6

With 2 roads,
$Cost_a = Cost_b = 83$

With 3 roads,
$Cost_a = Cost_b = Cost_c = 92$

Pareto-inefficient equilibria can exhibit unexpected behavior.

Definition: Braess Paradox.

There is a Braess Paradox if there exists two systems ini and aug such that

$$ini < aug \text{ and } \alpha^{(nc)}(ini) > \alpha^{(nc)}(aug).$$

i.e. adding resources to the system may reduce the performances of **ALL** players simulateously.

Braess Paradoxes: definition

From the New York Times, Dec 25, 1990, Page 38, What if They Closed 42d Street and Nobody Noticed?, By GINA KOLATA:

" ON Earth Day this year, New York City's Transportation Commissioner decided to close 42d Street, which as every New Yorker knows is always congested. "Many predicted it would be doomsday," said the Commissioner, Lucius J. Riccio. "You didn't need to be a rocket scientist or have a sophisticated computer queuing model to see that this could have been a major problem." But to everyone's surprise, Earth Day generated no historic traffic jam. Traffic flow actually improved when 42d Street was closed. "

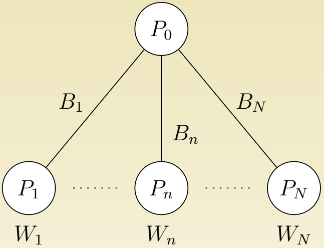

Hypothesis: the master can only send to 1 slave at a time.

Example

maître: $W = 2.55$

3 machines: $(B_i, W_i) = (4.12, 0.41)$, $(4.61, \mathbf{1.31})$, $(3.23, 4.76)$

2 applications: $b^1 = 1$, $w^1 = 2$, $b^2 = 2$, $w^2 = 1$

Equilibrium (ini): $a^1 = 0.173$, $a^2 = 0.0365$

Equilibrium ($W_2 = \mathbf{5.4}$): $a^1 = 0.127$, $a^2 = 0.0168$

Outline

Evolutionary games

- User strategies change with time as they adapt to the state
- Different possible dynamics:
  - Replicator dynamics:

$$\dot{y}_q^s = y_q^s \left(F_q^s(y) - \frac{1}{\hat{d}_q} \sum_{i=1}^{S_q} y_q^i F_q^i(y) \right).$$

  - Brown von Neumann Nash Dynamics (BNN):

$$\gamma_q^s = \max \left\{ F_q^s(y) - \frac{1}{\hat{d}_q} \sum_{i=1}^{S_q} y_q^i F_q^i(y), 0 \right\} \text{(excess payoff)}$$

$$\dot{y}_q^s = \hat{d}_q \gamma_q^s - y_q^s \sum_{j=1}^{S_q} \gamma_j^s.$$

(increase proportioanlly to the excess payoff / decrease proportionally to the sum of excess payoffs)

Evolutionary games

- Equilibria are called ESS (Evolutionary Stable Strategies) or
- Subset of Nash equilibria
- Stable by a deviation of a (small) fraction of users

Examples of applications:

- Power choice in ALOHA systems
  - Users can choose to transmit at high or low power (each packet)
  - High power has better chances of not being jammed
  - Low power save battery consumption
- Associations in wireless systems

Outline

We only scratched the surface...

Many other frameworks of games:

- Stackelberg equilibria: strategic game between 2 players: a leader and a follower (used in pricing mechanisms of e-services). Over types used in pricing Bertrand competition, Cournot competition.
- Stochastic games: a type of dynamic games (i.e. evolving over time) where the transitions are stochastic - the next state is determined by a probability distribution depending on the current state and the chosen actions (Markov Decision Processes) (used to choose efficient scheduling strategies)

How to improve non-cooperative performance?

No universal solution, but several options:

Correlated equilibria :

- A correlator give advises to each player
- (such that) the optimal strategy for each player is to follow the advice
- Nash equilibria $\subset$ Correlated equilibria

Interestingly, studies have shown that in certain cases, the correlator does not need to have any information on the system.

Pricing mechanisms :

- An entity gives money (reward) to players
- Each player strives to maximize its profit

Problem well studied in TCP-like networks (based on Lagrangian optimization)

Outline

Axiomatic definition VS optimization problem

1. Pareto optimality
2. Symmetry
3. Invariance towards linear transformations

$+$

- Independant to irrelevant alternatives
  Nash (NBS) / proportional fairness
  $\prod u_i$
- Monotonicity
  Raiffa-Kalai-Smorodinsky / max-min
  Recursively $\max\{u_i | \forall j, u_i \leqslant u_j\}$
- Inverse monotonicity
  Thomson / Social welfare
  $\max \sum u_i$

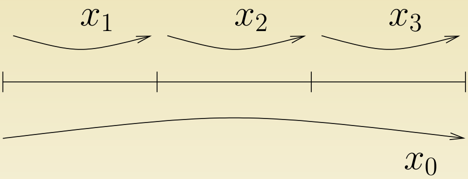

3 connections / 2 links.
Capacity constraints:

$$\begin{cases} x_1 + x_0 \leqslant 1, \\ x_2 + x_0 \leqslant 1, \\ x_3 + x_0 \leqslant 1. \end{cases}$$

4 unknowns and 3 (in)equations.

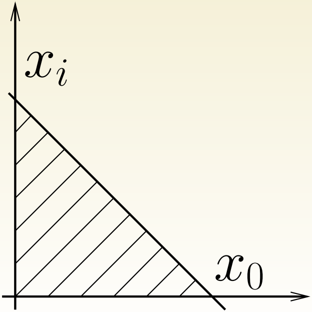

How to choose x_0 among the Pareto optimal points?

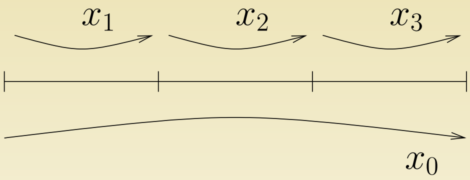

$$\begin{cases} x_1 + x_0 \leqslant 1, \\ x_2 + x_0 \leqslant 1, \\ x_3 + x_0 \leqslant 1. \end{cases}$$

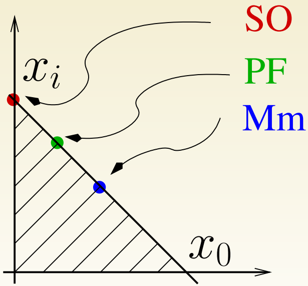

How to choose x_0 among the Pareto optimal points?

$$\begin{cases} x_0 = 0.5, \\ x_1 = x_2 = x_3 = 0.5 \end{cases}$$ Max-Min fairness

$$\begin{cases} x_0 = 0, \\ x_1 = x_2 = x_3 = 1 \end{cases}$$ Social Optimum

$$\begin{cases} x_0 = 0.25, \\ x_1 = x_2 = x_3 = 0.75 \end{cases}$$ Proportional fairness

Introduced by Mo and Walrand

utility

allocation

$$\max \sum_{n \in N} \frac{f_n(x_n)^{1-\alpha}}{1-\alpha}$$

player

Outline

Fairness family: example

The COST network (Prop. fairness)

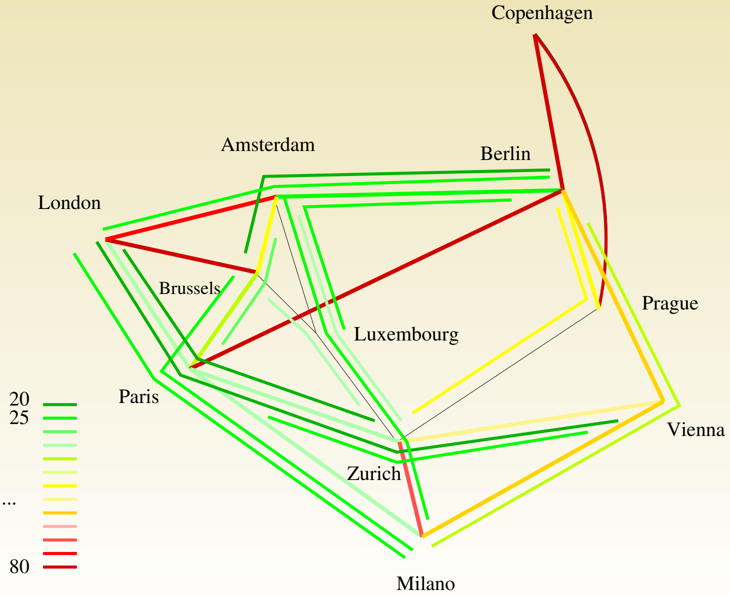

Fairness family: example

The COST network (Prop. fairness)

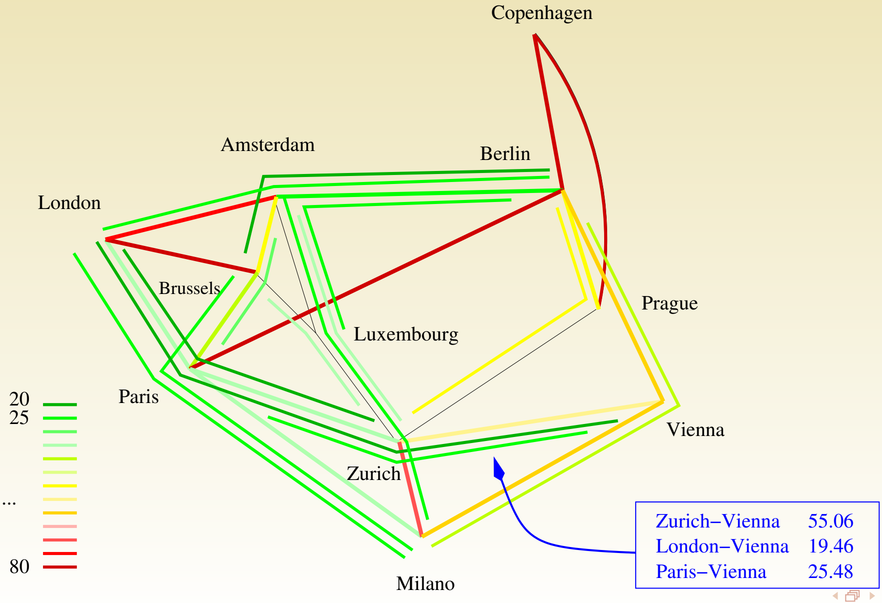

Other examples

- Bandwidth allocation: in TCP-like networks, in UMTS-like networks (joint bandwidth / power)
- Carriers allocation: between operators in satellite systems

Challenges:

- Non-convex systems: wireless communications, load balancing systems
- Developping distributed algorithms

Outline

Other hot topics in game theory

- Mechanism design: how to design rules of a game so as to achieve a specific outcome, even though each player is selfish. Done by setting up a structure in which each player has incentive to behave as the designer intends. (Leonid Hurwicz, Eric Maskin et Roger Myerson, Nobel 2007)
- Auctions: resource allocation in P2P, frequency allocation in wireless.
- Impact of non-cooperative players in a cooperative environment: free-riders of P2P, UDP clients in TCP networks.
- Fair division or cake cutting problem: how to divide resource such that all recipients believe that they have received their fair share (envy-free). (Steven Brams, Alan Taylor)

Even more topics in game theory

- Election: Plurality (traditional) voting systems are not necessarily fair.
- Stable marriages: Problem of finding a matching, where no element of the first set prefers an element of the other set that also prefers the first element, "Stable Marriage and its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms", Donald Knuth
- Super-modular games: utility functions are such that higher choices by one player make one's own strategy higher look relatively more desirable.
- Games with incomplete information or Bayesian games: some player have private information about something relevant to their decision making
- Games with imperfect information: players do not perfectly observe the actions of other players.

Outline

Example of inforced collaboration (set of rules inforced by the police)

While the purely non-cooperative approach would give...

Slides available at:
http://www-id.imag.fr/~touati/Talks/