Why using Game Theory in communication networks? Introduction of game theory concepts

Patrick Maillé and Bruno Tuffin

Institut Mines-Telecom/Telecom Bretagne, Inria Rennes Bretagne-Atlantique

IRISA Seminar on Network Economics, May 31, 2012

Outline

Outline

From centralization to decentralization

- Networking has switched from the centralized telephone network to the decentralized Internet (scalability reason).
- Decentralization (or deregulation) is a key factor.
- In such a situation:
  - From the decentralization, there is a general envisaged/advised behavior
  - But each *selfish* user can try to modify his behavior at his benefits and at the expense of the network performance.
    Example: TCP configuration
  - How to analyze this, and how to control and prevent such a thing?
- It is the purpose of non-cooperative game theory.

Competitive actors: not only users

- The Internet has also evolved from an academic to a commercial network with providers in competition for customers and services.
- As a consequence, users are not the only *competitive* actors, but also
  - network providers: several providers propose the same type of network access
  - applications/service providers/content providers: the same type of application can be proposed by several entities (ex: search engines...)
  - platforms/technologies: you may access the Internet from ADSL, WiFi, 3G, WiMAX, LTE...

All those interacting actors have to be considered.

Why changing the pricing scheme?

- Increase of Internet traffic due to
  - increasing number of subscribers
  - more and more demanding applications.
- Congestion is a consequence, with erratic QoS.
- Increasing capacity difficult if not impossible in access networks (last mile problem).
- Also, *flat rate* pricing is unfair and does not allow service differentiation.

Subject of debate...
But new contexts require new economic paradigms.

Convergence: requires new pricing offers

- Convergence of technologies and services: all services (web browsing, telephony, television) can be used on all technologies (Fixed, WiFi, 3G, LTE...).
  - How to charge fairly and efficiently those different technologies, with their different characteristics?
  - New technology: new issues to solve.
- Would it be possible to propose a pricing scheme involving several technologies at the same time?
- Marketing point of view from operators: propose grouped offers (bundles) to attract customers to services they would not consider otherwise.

A new issue: dealing with competition among providers

- Most works on pricing are dealing with a monopoly, but
  - competition forces providers to decide about prices and offers depending on competitors' ones.
  - some *a priori* promising pricing schemes may not resist to competition.
- Sometimes providers even operate on different technologies (Fixed, WiFi, 3G, LTE...), or on several simultaneously.
- Also, impact of competition on capacity investment? Do they have interest in investing (especially when congestion pricing is used)?

Illustration of an intricate competition model

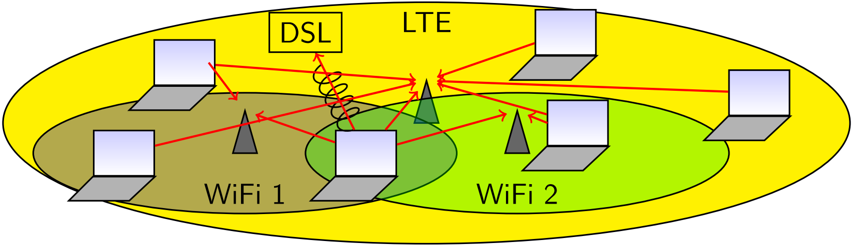

Sending end-to-end traffic through other (selfish) nodes/networks

- Here not a direct competition for customers, but providers have to pay other domains for forwarding their traffic and ensure end-to-end delivery (similar problems arise in ad hoc networks).
- How to design a self-managed network, with proper pricing incentives to forward traffic?
- Still unsolved: what are the optimal strategies of operators, in order to propose the best investments, in terms of:
  1. Investment on capacity: bandwidth for a domain or mobile network...
  2. Investment on products: new services.
  3. Investment on technology: new link between two domains, new base station, new WiFi hotspot...

Regulation: is it required?

- Free market can lead to an "inefficient" mechanism.
- Regulation can enforce providers to drive to the proper situation.
- Ex: to enforce providers to reduce retention time and authorize churn.
- New regulation/political issue: *network neutrality*
  - Network providers want to win on both sides: to charge users but also content providers, or degrade their services.
  - they do not want application providers not associated to them to congest their network.
  - Political debate: all players should be allowed the same access.
- Actors are then not free to do whatever we want.

Outline

What it changes

- While before optimization was the tool for routing, QoS provisionning, interactions between players has to be taken into account.
- Game theory: distributed optimization: individual users make their own decisions. "Easier" than to solve NP-hard problems (approximation).
- We need to look at a stable point (*Nash equilibrium*) for interactions.
- Tools used before in Economics, Transportation...
- and has recently appeared in telecommunications.
- We may have paradoxes (Braess paradox) that can be studied that way.
- A way to control things: to introduce pricing incentives/discouragements (TBC).

Typical networking applications

- P2P networks: a node tries to benefit from others, but limits its available resource (free riding)
- Grid computing: same issue, try to benefit from others' computing power, while limiting its own contribution.
- Routing games: each sending node tries to find the route minimizing delay, but intermediate links are shared with other flows (interactions).
- Ad hoc networks: what is the incentive of nodes to forward traffic of neighbors? If no one does, no traffic is successfully sent.
- Congestion control game (TCP...): why reducing your sending rate when congestion is detected?
- Power control in wireless networks: maximizing your power will induce a better QoS, but at the expense of others' interferences.
- Transmission games (WiFi...): if collision, when to resubmit packets?

Basic definitions

- Game theory: set of tools to understand the behavior of interacting *decision makers* or players.
- Classical assumption: players are rational: they have well-defined objectives, and they take into account the behavior of others.
- In this course: strategic or normal games, players play (simultaneously) once and for all.
- There are also branches called
  - extensive games, for which players play sequentially;
  - repeated games for which they can change their choices over time;
  - Bayesian games, evolutionary games...

General modelling tools

- Interactions of players through network performance. Tools:
  - queueing analysis or
  - signal processing.
- The action of a player has an impact on the output of other players, and therefore on their own strategies.
- They all have to play strategically.
- Each player i (user or provider) represented by its utility function $u_i(x)$ representing quantitatively its level of satisfaction (in monetary units for instance) when actions profile is $x = (x_i)_i$, where x_i denotes the action of player i.

Strategic Games

- A strategic game Γ consists of:
  - A finite set of players, N.
  - A set A_i of actions available to each player $i \in N$. and $A = \prod_{i \in N} A_i$.
  - For each player a utility function, (payoffs) $u_i : A \to \mathbb{R}$, characterizing the gain/utility from a state of the game.
- Players make decisions independently, without information about the choice of other players.
- We note $\Gamma = \{N, A_i, u_i\}$.
- For two players: description via a table, with payoffs corresponding to the strategic choices of users:

	C_1	C_2
R_1	$b_{11}\ c_{11}$	$b_{12}\ c_{12}$
R_2	$b_{21}\ c_{21}$	$b_{22}\ c_{22}$

$N = \{1,2\},\ A_1 = \{R_1, R_2\},\ A_2 = \{C_1, C_2\},\ u_1(R_j, C_k) = b_{jk},\ u_2(R_j, C_k) = c_{jk}.$

Example: association game

- Two users have the choice to connect to the Internet through WiFi and 3G
- If they both select the same technology, there will be interferences.
- They may get different throughput due to heterogeneous terminals and/or radio conditions
- Table of payoffs (obtained throughputs):

	3G	WiFi
3G	$3; 3$	$6; 4$
WiFi	$5; 6$	$1; 1$

- What is the best strategy for both players? Is there an "equilibrium" choice?

Nash equilibrium

- Most important equilibrium concept in game theory.
- Let $a \in A$ strategy profile, $a_i \in A_i$ player i's action, and a_{-i} denote the actions of the other players.
- Each player makes his own maximization.
- A Nash equilibrium is an action profile at which no user may gain by unilaterally deviating.

Definition

A Nash Equilibrium of a strategic game Γ is a profile $a^* \in A$ such that for every player $i \in N$:

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \;\; \forall a_i \in A_i$$

How to look for a Nash equilibrium?

- For each player i, look for the *best response* a_i in terms of a_{-i}.
- To find out a point such that no one can deviate (i.e. improve his utility): a strategy profile such that each player's action is a best response
- In a table with two players (can be generalized):
  1. Write in bold the best response of a player for each choice of the opponent;
  2. A Nash equilibrium is a profile where both actions are in bold.
  3. Example (blue is also used here):

	C_1	C_2
F_1	b_{11} $\mathbf{c_{11}}$	$\mathbf{b_{12}}$ c_{12}
F_2	$\mathbf{b_{21}}$ $\mathbf{c_{21}}$	b_{22} c_{22}

Classical illustration: The Battle of the Sexes

- *Bach or Stravisky ?* Married people want to go together to a concert of Bach or Stravisky. Their main concern is to go together, but one person prefers Stravisky and the other Bach.

	B	S
B	$2; 1$	$0; 0$
S	$0; 0$	$1; 2$

Classical illustration: The Battle of the Sexes

- *Bach or Stravisky ?* Married people want to go together to a concert of Bach or Stravisky. Their main concern is to go together, but one person prefers Stravisky and the other Bach.

	B	S
B	$2;1$	$0;0$
S	$0;0$	$1;2$

$\Rightarrow$

	B	S
B	$\mathbf{2;1}$	$0;0$
S	$0;0$	$\mathbf{1;2}$

- The game has two N.E.: (B, B) *and* (S, S).

Nash equilibrium in our association game

- Two users have the choice to connect to the Internet through WiFi and 3G
- If they both select the same technology, there will be interferences.
- They may get different throughput due to heterogeneous terminals and/or radio conditions
- Table of payoffs (obtained throughputs):

	3G	WiFi
3G	3; 3	6; 4
WiFi	5; 6	1; 1

Nash equilibrium in our association game

- Two users have the choice to connect to the Internet through WiFi and 3G
- If they both select the same technology, there will be interferences.
- They may get different throughput due to heterogeneous terminals and/or radio conditions
- Table of payoffs (obtained throughputs):

	3G	WiFi
3G	3; 3	6; 4
WiFi	5; 6	1; 1

$\Rightarrow$

	3G	WiFi
3G	3; 3	**6; 4**
WiFi	**5; 6**	1; 1

- Nash equilibria payoffs: (5; 6) and (6; 4).

Prisonner's Dilemma

- Suspects of a crime are in separate cells.
- If they both confess, each will be sentenced three years of prison.
- If only one confesses, he will be free and the other will be sentenced four years.
- If neither confess the sentence will be a year in prison for each one.
- Goal here: to minimize years in prison.
- Utility $u_i = 4-$number of year in jail.

	don't confess	*confess*
don't confess	$3; 3$	$0;$ **4**
confess	**4**$; 0$	**1**; **1**

- Best outcome: no one confesses, but this requires cooperation.
- But, (confess, confess) is the unique N.E.
- Not optimal!

Prisonner's Dilemma in wireless networks

Gaoning He PhD thesis, Eurecom, 2010

- Two players sending information at a base station.
- Two power levels: High or Normal.
- Payoff table:

	Normal	High
Normal	Win; Win	Lose much; Win much
High	Win much; Lose much	Lose; Lose

- Best outcome: Normal, but this requires cooperation.
- But, (High, High) is the unique N.E.
- Not optimal here too!

Pareto-optimum situation

Definition: Pareto optimum

An outcome of the game with player utilities $(u_i(a_1^*, \ldots, a_I^*))$ is Pareto-optimal if and only if for any action profile $(a_i) \in \prod A_i$,

$$\exists i : u_i(a_1, \ldots, a_I) > u_i(a_1^*, \ldots, a_I^*) \Rightarrow \exists j : u_j(a_1, \ldots, a_I) < u_j(a_1^*, \ldots, a_I^*)$$

At a Pareto optimum, there is no way of improving the utility of any player without deteriorating the utility of another one.

There can be a lot of Pareto-optimal situations!
Nash equilibria are not necessarily Pareto-optimal (cf the Prisoner's Dilemma)

A Nash equilibrium does not always exist

- Game where 2 players play odd and even:

	Odd	*Even*
Odd	**1**; -1	-1; **1**
Even	-1; **1**	**1**; -1

- This game does not have a N.E.
- So in general, games may have no, one, or several Nash equilibria...

Case of continuous set of actions

- In the case of a continuous set of strategies, simple derivation can be used to determine the Nash equilibrium (always simpler!).
- For two players 1 and 2: draw the best-response in terms

$$BR_1(x_2) = \text{argmax}_{x_1} u_1(x_1, x_2) \text{ and } BR_2(x_1) = \text{argmax}_{x_2} u_2(x_1, x_2).$$

A Nash equilibrium is an intersection point of the best-response curves:

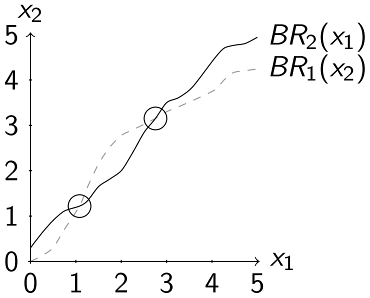

Application to power control

Saraydar, Mandayam & Goodman, 2002

- In CDMA-based networks each user can play on transmission power.
- QoS based on the signal-to-interference-and-noise ratio (SINR):

$$SINR_i = \gamma_i = \frac{W}{R} \frac{h_i p_i}{\sum_{j \neq i} h_j p_j + \sigma^2}$$

with W spread-spectrum bandwdith, R rate of transmission, p_i power transmission, h_i path gain, σ^2 background noise.

- Different utility functions found in the litterature. Ex: the number of bits transmitted per Joule

$$u_j(p_i, \gamma_i) = \frac{R}{p_i}(1 - 2BER(\gamma_i))^L = \frac{R}{p_i}(1 - e^{-\gamma_i/2})^L$$

where $BER(\gamma_i)$ bit error rate and L length of symbols (packets).

- Increasing *alone* your own power increases your QoS, but decreases the others'.
  $\Rightarrow$ Game theory.
- A Nash eq. exists, but its efficiency can be improved through pricing.

Mixed strategies

- Previous Nash equilibrium also called *pure Nash equilibrium*.
- A mixed strategy is a probability distribution over pure strategies: $\pi_i(a_i)\ \forall a_i \in A_i$.
- Player i utility function is the expected value over distributions

$$\mathbb{E}_\pi[u_i] = \sum_{a \in A} u_i(a) \left(\prod_i \pi_i(a_i) \right).$$

- A Nash equilibrium is a set of distribution functions $\pi^* = (\pi_i^*)_i$ such that no user i can unilaterally improve his expected utility by changing alone his distribution π_i.
  Formally,

$$\forall i, \forall \pi_i, \quad \mathbb{E}_{\pi^*}[u_i] \geq \mathbb{E}_{(\pi_i, \pi_{-i}^*)}[u_i].$$

Theorem

Advantage (proved by John Nash): for every finite game, there always exist a (Nash) equilibrium in mixed strategies.

Interpretation of mixed strategies

- Concept of mixed strategies known as “intuitively problematic”.
- Simplest and most direct view: randomization, from a ‘lottery”.
- Other interpretation: case of a large population of agents, where each of the agent chooses a pure strategy, and the payoff depends on the fraction of agents choosing each strategy. This represents the distribution of pure strategies (does not fit the case of individual agents).
- Or comes from the game being played several times *independently*.
- Other interpretation: randomization comes from the lack of knowledge of the agent’s information (purification).

Illustration of mixed strategies: jamming game

- Consider two mobiles wishing to transmit at a base station: a regular transmitter (1) and a jammer (2)
- Two channels, c_1 and c_2 for transmission, collision if they transmit on the same channel, success otherwise
- For the regular transmitter: reward for success 1, -1 if collision
- For the jammer: reward 1 if collision, -1 if missed jamming.
- payoff table

	c_1	c_2
c_1	$-1; \mathbf{1}$	$\mathbf{1}; -1$
c_2	$\mathbf{1}; -1$	$-1; \mathbf{1}$

- No pure Nash equilibrium.

Mixed strategy equilibrium for the jamming game

- The transmitter (resp. jammer) chooses a probability p_t (resp. p_j) to transmit on channel c_1.
- Utilities (average payoff values):

$$\begin{aligned} u_t(p_t, p_j) &= -1(p_t p_j + (1-p_t)(1-p_j)) + 1(p_t(1-p_j) + (1-p_t)p_j) \\ &= -1 + 2p_t + 2p_j - 4p_t p_j \\ u_j(p_t, p_j) &= 1(p_t p_j + (1-p_t)(1-p_j)) + -1(p_t(1-p_j) + (1-p_t)p_j) \\ &= 1 - 2p_t - 2p_j + 4p_t p_j \end{aligned}$$

- For finding the Nash equilibrium:

$$\begin{aligned} \frac{\partial u_t(p_t, p_j)}{\partial p_t} &= 2 - 4p_j = 0 \\ \frac{\partial u_j(p_t, p_j)}{\partial p_j} &= 2 - 4p_t = 0. \end{aligned}$$

- $(p_t = 1/2, p_j = 1/2)$ mixed Nash equilibrium (sufficient conditions verified too).

Other notion: Stackelberg game

- Decision maker (network administrator, designer, service provider...) wants to optimize a utility function.
- His utility depends on the reaction of users (who want to maximize their own utility, minimize their delay...)
- Hierarchical relationship: *leader-follower problem* called *Stackelberg game*.
  - For a set of parameters provided by the leader, followers (users) respond by seeking a new algorithm between them.
  - The leader has to find out the parameters that lead to the equilibrium yielding the best outcome for him.
- Typical application: the provider plays on prices, capacities, users react on traffic rates...

Stackelberg game: formal problem

- Say that there are N users
- Let $u(x) = (u_1(x), \ldots, u_N(x))$ the utility function vector for users for the set of parameters x set by the leader.
- Denote by $R(u(x), x)$ the utility of the leader.
- Define $u^*(x)$ as the (Nash) equilibrium (if any) corresponding to x.
- Goal: find x^* such that

$$R(u^*(x^*), x^*) = \max_x R(u^*(x), x).$$

- Works fine if $u^*(x)$ is unique
- If not, and if $U^*(x)$ is the set of equilibria, we may want to maximize the *worst case*: find x^* such that

$$x^* \in \arg\max_x \min_{u^*(x) \in U^*(x)} R(u^*(x), x).$$

Simple illustration of Stackelberg game

- leader: service provider fixing its price p
- followers: users, modeled by a demand function $D(p)$ representing the *equilibrium* population accepting the service for a given price.
- Equilibrium among users therefore already included in the model.
- The provider chooses the price p to maximize its revenue

$$R(p) = pD(p).$$

- Obtained by computing the derivative of $R(p)$.

The value of information

We consider a coordination game on an innovating service.
We assume that Operator 1 invests first in a service (A or B), then Operator 2 invests in A or B knowing the choice of Operator 1 (thus Operator 1 is a leader).

We assume that users will finally adopt only one service, each one with probability 1/2.

The payoff of each operator depends on the final choice of users, that is only known in probability.

The value of information

Given the time difference between both player choices, the game can be represented as a tree:

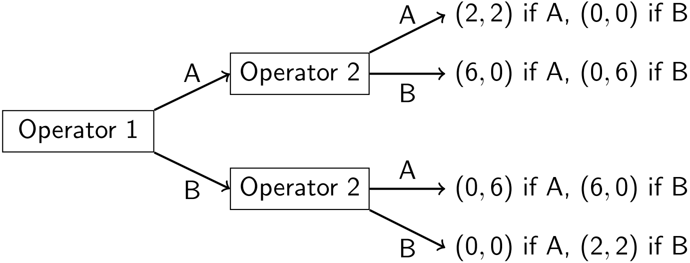

We consider that each operator maximizes his expected payoff.
What is the equilibrium?

The value of information

Given the time difference between both player choices, the game can be represented as a tree:

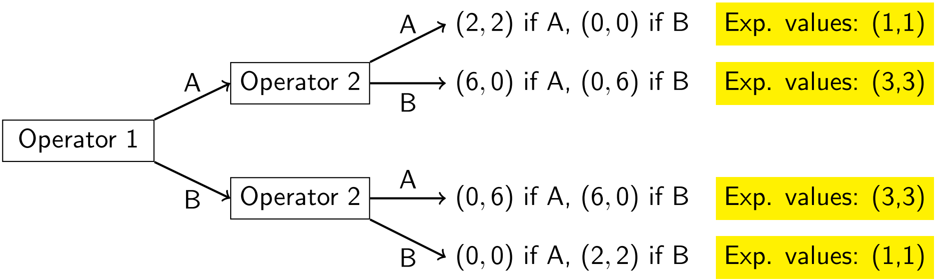

We consider that each operator maximizes his expected payoff.
What is the equilibrium? $\Rightarrow$Operator 2 maximizes his payoff by making a choice different from Operator 1.
Expected payoffs at equilibrium: $(3, 3)$

Incomplete information equilibrium

The same game is played, but Operator 1 knows the service that will be adopted, and Operator 2 knows that Operator 1 has that information:

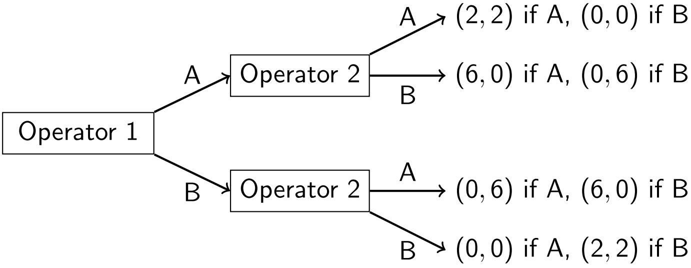

Expected payoffs at equilibrium:

Incomplete information equilibrium

The same game is played, but Operator 1 knows the service that will be adopted, and Operator 2 knows that Operator 1 has that information:

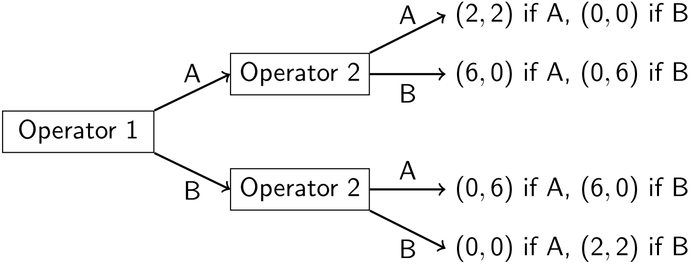

Expected payoffs at equilibrium: $(2, 2) < (3, 3)$

Incomplete information equilibrium

The same game is played, but Operator 1 knows the service that will be adopted, and Operator 2 knows that Operator 1 has that information:

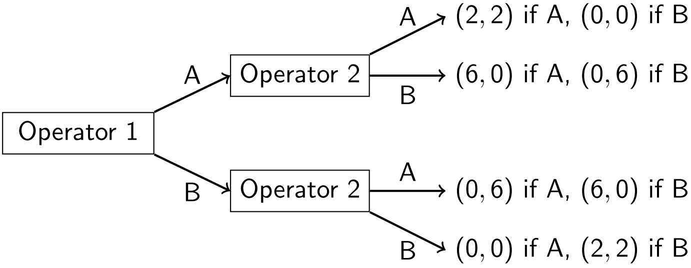

Expected payoffs at equilibrium: $(2, 2) < (3, 3)$
Particular case when information has a negative value!

Wardrop equilibrium

- Developped to analyze road traffic, to distribute traffic between available routes.
- Each user wants to minimize his transportation time (congestion-dependent), non-cooperatively.

Definition (Wardrop's first principle)

Time in all routes actually used are equal and less than those which would be experienced by a single vehicle on any unused route.

- Exactly the same idea as Nash equilibrium (with minimal transportation cost), except that each user is infinitesimal (large number of users), meaning that his own action does not have any impact on the equilibrium; only an aggregated number does.

An example: Pigou's instance (1/2)

Interpretation: imagine one unit (million, thousand) of commuters willing to go from the suburbs to the city center to work.

Two choices:

- take public transport⇒fixed commute time (1h)
- take one's car⇒commute time depends on the number x of people taking their car (congestion dependence), assume commute time is x.

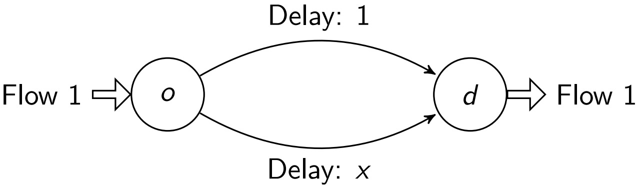

An example: Pigou's instance (2/2)

Only one equilibrium: everybody takes his car and experiences a commute time of 1h!

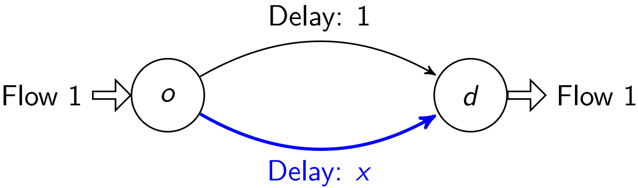

That outcome is **not Pareto-efficient**: we could strictly decrease the commute time of some users without increasing that of the others by making some users switch to public transport.

- Enforce people to take public transport⇒badly perceived
- Give *incentives* to take public transport instead of one's car: taxes on roads, subsidies on public transport.

The Braess paradox

We could reasonably expect that adding a resource (or reducing the cost of existing resources) always improves the total cost.

$\Rightarrow$It is false, due to participants' selfishness!

Definition

A situation where adding a resource increases total cost is called a *Braess paradox*.

Example (Braess, 1969):

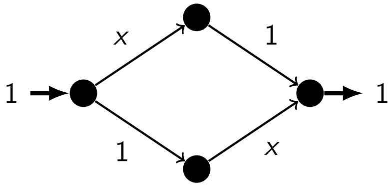

Without link between north and south nodes:

The Braess paradox

We could reasonably expect that adding a resource (or reducing the cost of existing resources) always improves the total cost.

$\Rightarrow$It is false, due to participants' selfishness!

Definition

A situation where adding a resource increases total cost is called a *Braess paradox*.

Example (Braess, 1969):

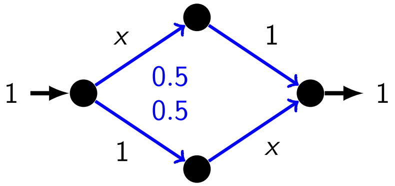

Without link between north and south nodes:

each user cost=3/2.

The Braess paradox

We could reasonably expect that adding a resource (or reducing the cost of existing resources) always improves the total cost.

$\Rightarrow$It is false, due to participants' selfishness!

Definition

A situation where adding a resource increases total cost is called a *Braess paradox*.

Example (Braess, 1969):

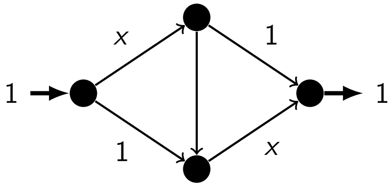

Without link between north and south nodes:

each user cost$=3/2$.

With a zero-cost link between north and south nodes:

The Braess paradox

We could reasonably expect that adding a resource (or reducing the cost of existing resources) always improves the total cost.

$\Rightarrow$It is false, due to participants' selfishness!

Definition

A situation where adding a resource increases total cost is called a *Braess paradox*.

Example (Braess, 1969):

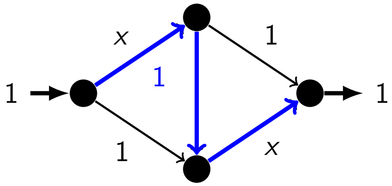

Without link between north and south nodes:

each user cost=3/2.

With a zero-cost link between north and south nodes:

each user cost=2.

The Braess paradox

We could reasonably expect that adding a resource (or reducing the cost of existing resources) always improves the total cost.

$\Rightarrow$It is false, due to participants' selfishness!

Definition

A situation where adding a resource increases total cost is called a *Braess paradox*.

Example (Braess, 1969):

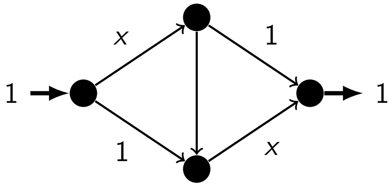

Without link between north and south nodes:

each user cost$=3/2$.

With a zero-cost link between north and south nodes:

each user cost$=2$.

$\Rightarrow$**Adding the link has worsened the cost for all users.**

Braess paradox in “real life”

In the New York Times, Dec. 25, 1990, p38, *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.*

Price of Anarchy

- The social utility function can be optimized when we have a single authority who dictates every agent what to do.
- When agents choose their own action, we should study their behavior and compare the obtained social utility with the optimal one.

Definition (Price of Anarchy)

Two possible definitions (depending on the type of objective functions):

- ratio of optimal social utility divided by the worst social utility at a Nash equilibrium.
- ratio of the worst social cost at a Nash equilibrium divided by the minimal social cost.

- A price of Anarchy of 1 corresponds to the optimal case where decentralization does not bring any loss of efficiency (that may happen).
- Research activity for computing bounds for the price of Anarchy in specific games.

Outline

Again, why pricing?

- Return on investment for providers
  - providers need to get their money back
  - if no revenue made, no network improvement possible
- Demand/congestion control
  - the higher the price, the smaller the demand, and the better the QoS
  - an "optimal" situation can be reached
- Why changing the current (flat) pricing scheme?
  - flat-rate pricing unfair, demand uncontrolled
  - service differentiation impossible to favor QoS-demanding applications otherwise
- Heterogeneity of technologies/applications
  - different services (telephony, web, email, TV) available through multiple medias (fixed, 3G, WiFi...)
  - appropriate and bundle contracts to be proposed.
- A lot of new contexts: MNO *vs* MVNO, cognitive networks...
  - adaptation of economic models to be realized for an optimal network use.

Other reasons for pricing

- Regulation issue
  - When no equilibrium, pricing can help to drive to such a point.
  - By playing on prices, a better situation can be obtained
- But, network neutrality problem: not everything can be proposed
  - current political debate
  - introduced because network providers wanted to differentiate among service providers
  - could limit the user-benefit-oriented service differentiation.

Illustration of pricing interest

Courcoubetis & Weber, 2003

- User i buying a service quantity x_i at unit price p.
- $u_i(x_i, y)$ utility for using quantity x_i, where $y = \sum_j x_j/k$ with k resource capacity.
- u_i assumed decreasing in y: negative externality because of congestion.
- Net benefit of user i:

$$u_i(x_i, y) - px_i$$

- Benefit of provider: $p\sum_i x_i - c(k)$.
- Social welfare: sum of benefits of all actors in the game (provider + users):

$$SW = \sum_i u_i(x_i, y) - c(k).$$

- Optimal SW determined by maximizing over $x_1, \ldots; x_n$. Leads to (by differentiating over each x_i)

$$\frac{\partial u_i(x_i^*, y^*)}{\partial x_i} + \frac{1}{k}\sum_j \frac{\partial u_j(x_j^*, y^*)}{\partial y} = 0 \quad \forall i.$$

Illustration of pricing interest (2) Courcoubetis & Weber, 2003

- Define the price as the marginal decrease in SW due to a marginal increase in congestion, at the SW optimum,

$$p_E = -\frac{1}{k} \sum_j \frac{\partial u_j(x_j^*, y^*)}{\partial y}$$

(positive thanks to the decreasingness of u_i in y)

- With this price, a user acting selfishly tries to optimize his net benefit

$$\max_{x_i} u_i(x_i, y) - p_E x_i.$$

- Differentiating with respect to x_i, this gives

$$\frac{\partial u_i}{\partial x_i} + \frac{1}{k}\frac{\partial u_i}{\partial y} - p_E = 0$$

- For a large n, assuming $\left|\frac{\partial u_i}{\partial y}\right| << \left|\sum_j \frac{\partial u_j}{\partial y}\right|$, we get approximately the same system of equations than when optimizing SW.
- Pricing can therefore help to drive to an optimal situation.

Proposed pricing schemes

- Pricing for guaranteed services through reservation and admission control.
  Drawback: scalability.
- *Paris Metro Pricing*: separate the network into logical subnetworks with different access charges.
  Advantage: simple. Drawback: does not work in a competitive market.
- *Cumulus pricing scheme*: +/- points awarded if predefined contract respected. Penalties and renegotiations.
  Advantage: easy to implement.
- Priority pricing: classes of traffic with different priority levels and access prices;
  - scheduling priority
  - rejection or dropping priority.

  Advantage: easy to implement.
- Auctioning, for priority at the packet level, or for bandwidth at the flow level.
- Pricing based on transfer rates and shadow prices.

Example: auctionning for bandwidth

The problem of resource allocation

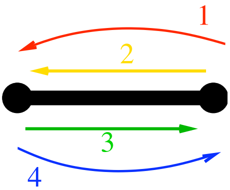

- Allocate bandwidth among users on a link with a capacity constraint Q
- More general results also obtained
- Allocation and pricing mechanism: determines the allocation a_i for each player i, and the price c_i he is charged.

Which allocation and pricing rule? Based on Vickrey-Clarke-Groves (VCG) auction mechanism.

General Vickrey-Clarke-Groves (VCG) auctions description

- Applicable to any problem where players (users) have a *quasi-linear* utility function.
- Utility of user i:

$$U_i(a, c_i) = \theta_i(a) - c_i,$$

  with
  - θ_i is called the *valuation* or *willingness-to-pay* function of user i
  - a outcome (say, the resource allocation vector), $a = (a_1, \ldots, a_n)$.
  - c_i total charge to i (can be non-positive).
- VCG asks users to declare their valuation function $\tilde{\theta}_i$

VCG allocation and pricing rules

- the mechanism computes an outcome $a(\tilde{\theta})$ that maximizes the declared social welfare:

$$a(\tilde{\theta}) \in \arg\max_x \sum_i \tilde{\theta}_i(x);$$

- the price paid by each user corresponds to the loss of declared welfare he imposes to the others through his presence:

$$c_i = \max_x \sum_{j \neq i} \tilde{\theta}_j(x) - \sum_{j \neq i} \tilde{\theta}_j(a(\tilde{\theta})).$$

VCG mechanism properties

The mechanism verifies three major properties:

- Incentive compatibility: for each user, bidding truthfully (i.e. declaring $\tilde{\theta}_i = \theta_i$) is a dominant strategy.
- Individual rationality: each truthful player obtains a non-negative utility.
- Efficiency: when players bid truthfully, social welfare ($\sum_i \theta_i$) is maximized.

Back to the auction for bandwidth issue N. Semret PhD thesis, 1999

For a link of capacity Q.

- Each player i submits bid $s_i = (q_i, p_i)$ with
  - q_i asked quantity
  - p_i associated price.
- Allocation a_i and total charge c_i such that
  - $\sum_i a_i \leq Q$: do not allocate more than the available capacity
  - $c_i \leq p_i q_i$: charge less than the declated total valuation.
- bid profile $s = (s_1, \ldots s_n)$ and s_{-i} bid profile excluding player i.
- Unused capacity for user i at price y:

$$Q_i(y; s_{-i}) = \left[Q - \sum_{j \neq i : p_j > y} q_j \right]^+ .$$

Allocation and pricing rule

- Allocation: priority to highest bids,

$$a_i(s) = \min\left(q_i, \frac{q_i}{\sum_{k:p_k=p_i} q_k} Q_i(p_i; s_{-i})\right)$$

  - you get 0 if nothing remains,
  - your quantity if still available at your bid and enough remains to serve all quantities at same unit price,
  - or you share proportionally what remains if not to serve to cover all bids at p_i.

- Charge

$$c_i(s) = \sum_{j \neq i} p_j [a_j(0; s_{-i}) - a_j(s_i; s_{-i})]$$

  - you pay the loss of valuation your presence creates on other players.

Numerical illustration

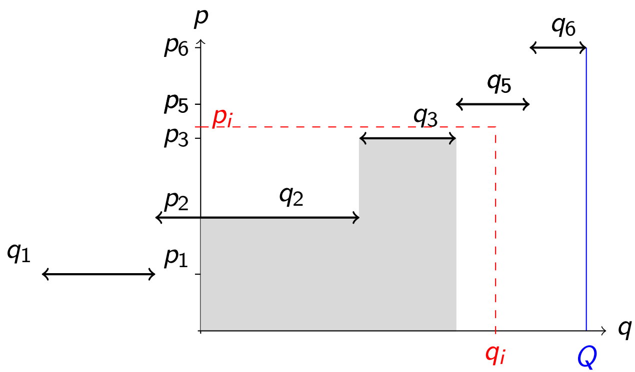

- bid (q_i, p_i) does not allows i to get the required quantity.
- Bids with higher price are allocated first.
- Player i gets *what remains*.
- Charge: loss declared by i's presence (here players 2 and 3); grey zone.

Algorithm and results

- Users' preferences: determined by their **utility function** $u_i(s) = \theta_i(a_i(s)) - c_i(s)$
- θ_i =player i's **valuation function**, assumed non-decreasing and concave
- User i's goal: maximizing his utility $\theta_i(a_i) - c_i$.
- Users play sequentially, optimizing their utility given s_{-i}, up to reaching an ϵ-Nash equilibrium where no user can improve his utility by more then ϵ.
- ϵ: bid fee. Avoids oscillations around the real Nash equilibrium.

Properties of the scheme

- *a) Incentive compatibility*: A player cannot do much better than simply revealing his valuation.
- *b) Individual rationality*: $U_i \geq 0$, whatever the other players bid.
- *c) Efficiency*: When players submit truthful bids, the allocation maximizes social welfare.

Issues:

1. requires a lot of signalling: at each round, users need to know the whole bid profile
2. takes time to reach an ϵ-Nash equilibrium
3. when users leave or enter: needs a new application of the sequential algorithm, with a loss of efficiency during the transient phase.

Those aspects solved by the next proposition.

Multi-bid auctions

Maillé & Tuffin, Infocom 04, IEEE/ACM ToN 06

Improvement in-between sending a single bid several times and sending a whole function (not practical).

- When entering the game, each player i submits M_i two-dimensional bids of the form $s_i^{m_i} = (q_i^{m_i}, p_i^{m_i})$ where
$$\begin{cases} q_i^j & = & \text{asked quantity of resource} \\ p_i^j & = & \text{corresponding proposed unit price} \end{cases}$$
- Allocations a_i and charges c_i computed based on s.

User behaviour

- Set $\mathcal{I}$ of users (players)
  - Users' preferences: determined by their **utility function** $u_i(s) = \theta_i(a_i(s)) - c_i(s)$
  - θ_i =player i's **valuation function**, assumed non-decreasing and concave
  - User i's goal: maximizing his utility $\theta_i(a_i) - c_i$.
- The auctioneer uses player i's multi-bid s_i to compute:
  - the pseudo-marginal valuation function $\bar{\theta}_i'$
  - the pseudo-demand function $\bar{d}_i$

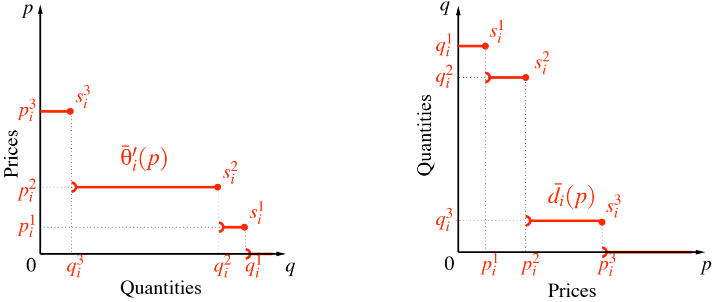

Pseudo-marginal valuation and pseudo-demand functions associated with the multi-bid s_i

$$\bar{\theta}_i'(q) \quad = \quad \max_{1\le m\le M_i}\{p_i^m : q_i^m \ge q\} \text{ if } q_i^1 \ge q, \qquad 0 \text{ otherwise.}$$

$$\bar{d}_i(p) \quad = \quad \max_{1\le m\le M_i}\{q_i^m : p_i^m \ge p\} \text{ if } p_i^{M_i} < p, \qquad 0 \text{ otherwise.}$$

Allocation and pricing rule

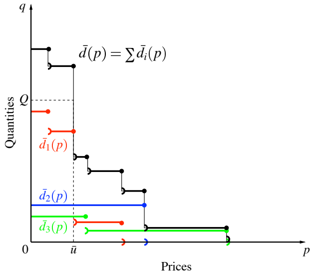

$\bar{u}$: pseudo market clearing price (highest unit price at which demand exceeds capacity).

- Multi-bid allocation: $a_i(s) = \bar{d}_i(\bar{u}^+) + \frac{\bar{d}_i(\bar{u}) - \bar{d}_i(\bar{u}^+)}{\bar{d}(\bar{u}) - \bar{d}(\bar{u}^+)}(Q - \bar{d}(\bar{u}^+))$
- Pricing principle : each user pays for the declared "social opportunity cost" he imposes on others
- If s denotes the bid profile,

$$c_i(s) = \sum_{j \in \mathcal{I} \cup \{0\}, j \neq i} \int_{a_j(s)}^{a_j(s_{-i})} \bar{\theta}'_j$$

Properties of the scheme

Here too, we have been able to prove the following properties are satisfied:

- *a) Incentive compatibility*;
- *b) Individual rationality*;
- *c) Efficiency* (in terms of social welfare).

Advantages:

- Bids given only once (when entering the game);
- No information required about network conditions and bid profile;
- No convergence phase needed: if network conditions change, new allocations and charges automatically computed (no associated loss of efficiency).

Other mechanisms since: double-sided auctions for instance...

Outline

Interdomain problem

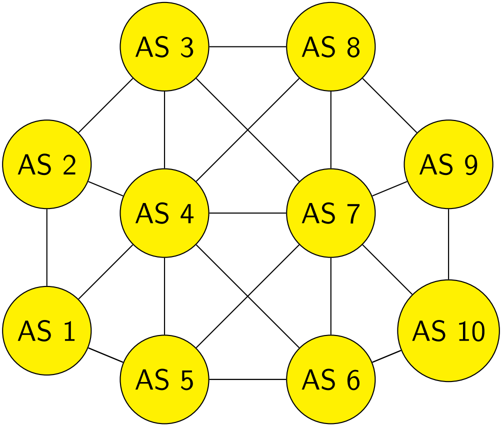

- Network made of Autonomous Systems (ASes) acting selfishly.

Interdomain problem

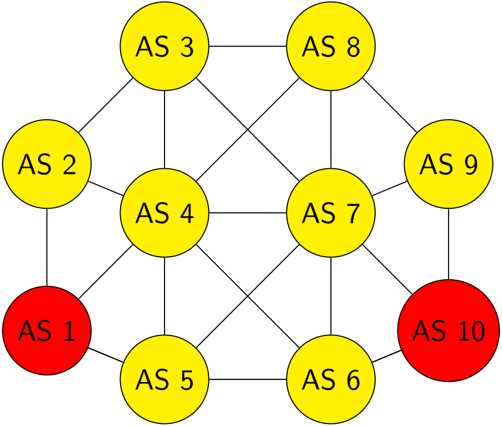

- Network made of Autonomous Systems (ASes) acting selfishly.
- A node (an AS) needs to send traffic from its own customers to other ASes.
- Introduce incentives for intermediate nodes to forward traffic, via pricing.

Interdomain problem

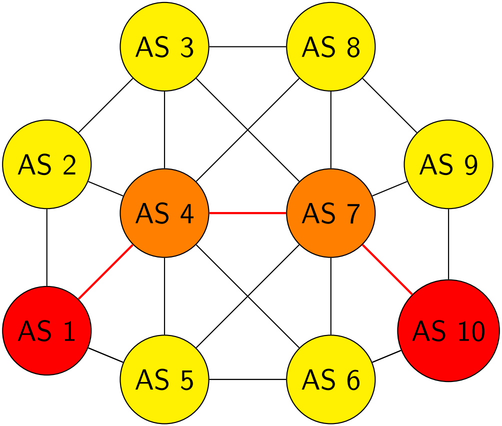

- Network made of Autonomous Systems (ASes) acting selfishly.
- A node (an AS) needs to send traffic from its own customers to other ASes.
- Introduce incentives for intermediate nodes to forward traffic, via pricing.
- What is the best path?

Interdomain problem

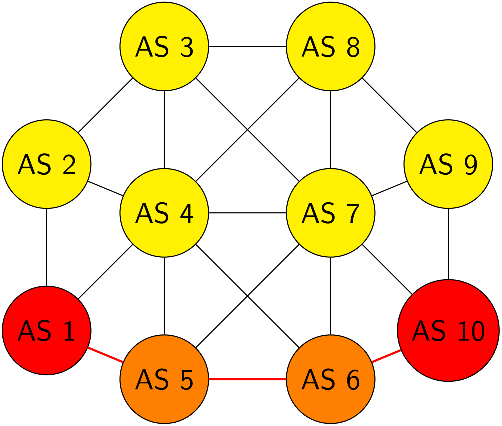

- Network made of Autonomous Systems (ASes) acting selfishly.
- A node (an AS) needs to send traffic from its own customers to other ASes.
- Introduce incentives for intermediate nodes to forward traffic, via pricing.
- What is the best path?

Interdomain issues

- similar problems in
  - ad-hoc networks: individual nodes should be rewarded for forwarding traffic (especially due to power use);
  - P2P systems: free riding can be avoided through pricing.
- How to implement it?
  - The AS can contact all potential ASes on a path to learn their costs, and then make its decision.
  - More likely: he contacts only its neighbors, which ask the cost to their own neighbors with a BGP-based algorithm.
    On the way back, declared costs are added.
- Two different mathematical problems
  - Finite capacity at each AS: it becomes similar to a knapsack problem.
  - Capacity assumed infinite if networks overprovisionned thanks to optic fiber (last mile problem, i.e., connection to users, not considered here).

Relevant (desirable) properties

- Individual rationality: ensures that participating to the game will give non-negative utility.
- Incentive compatibility: ASes' best interest is to declare their real costs.
- Efficiency: mechanism results in a maximized sum of utilities.
- Budget Balance: sum of money exchanged is null (or at least non-negative).
- Decentralized: decentralized implementation of the mechanism.
- Collusion robustness: no incentive to collusion among ASes.

Is there a pricing mechanism:

- verifying the whole set or a given set of properties?
- Or/and verifying *almost* all of them?

Interdomain pricing with no resource constraints

Feigenbaum et al. 2002

- Inter-domain routing handled by a simple modification of BGP.
- Amount of traffic T_{ij} from AS i to AS j, with per-unit cost c_k for forwarding for AS k.
- Valuation of intermediate domain k for a given allocation (a routing decision) is

$$\theta_k(\text{routing}) = -c_k \sum_{\{(i,j) \text{ routed trough } k\}} T_{ij}.$$

- Maximizing sum of utilities is equivalent to minimizing the total routing cost

$$\sum_{i,j} T_{ij} \sum_{k \in path(i,j)} c_k,$$

where

  - each AS declares its transit cost c_k
  - the least (declared) cost route $path(i,j)$ is computed for each origin-destination pair (i,j).

VCG auctions and drawback in interdomain context

- Payment rule to intermediate node k (opportunity cost-based):

$$p_k = c_k + \left(\sum_{\ell \text{ on } path^{-k}(i,j)} c_\ell - \sum_{\ell \text{ on } path(i,j)} c_\ell \right)$$

with $path^{-k}(i,j)$ the selected path when k declares an infinite cost.

- Subsequent properties
  - Efficiency
  - Incentive compatibility
  - Individual rationality
- Only pricing mechanism to provide the three properties at the same time.

VCG auctions and drawback in interdomain context

- Payment rule to intermediate node k (opportunity cost-based):

$$p_k = c_k + \left(\sum_{\ell \text{ on } path^{-k}(i,j)} c_\ell - \sum_{\ell \text{ on } path(i,j)} c_\ell \right)$$

  with $path^{-k}(i,j)$ the selected path when k declares an infinite cost.
- Subsequent properties
  - Efficiency
  - Incentive compatibility
  - Individual rationality
- Only pricing mechanism to provide the three properties at the same time.
- But who should pay the subsidies? Sender's willingness to pay not taken into account. That should be!
  The VCG payment from sender is the sum of declared costs if traffic is effectively sent: always below the sum of subsidies.
  Very unlikely to apply in practice: no central authority to permanently inject money.

Impossibility result and what is the good choice?

- General result: no mechanism can actually verify efficiency, incentive compatibility, individual rationality and budget balance.
- Current question: what set of properties to verify? Which mechanism to apply?
  - The "almost" property could be a more flexible choice.
  - Strict requirement: budget balance. Decentralization too if dealing with large topologies.

Outline

Things to remember

1. Selfishness **does** need to be taken into account in telecommunication networks.
2. The Nash equilibrium is a notion that helps predict the possible rational outcomes of a game.
3. It is often not Pareto-optimal, and different from the social optimum.
4. The outcome of the game strongly depends on the information that each player has.
5. The social optimum may be reached by changing the rules of the game (e.g., via additional payments).
6. However, designing a mechanism with a given set of desirable properties is not always doable.