Game Theory applied to Networking

Bruno Tuffin

INRIA Rennes - Bretagne Atlantique

PEV: Performance EValuation
M2RI - Networks and Systems Track
Rennes

Outline

References

1. M.J. Osborne and A. Rubinstein. *A course in Game Theory*. MIT Press, 1994.
2. E. Altman, T. Boulogne, R. El-Azouzi, T. Jimenez and L. Wynter. A Survey on networking games in telecommunications. *Computers & Operations Research*, 2004.
3. A.B. MacKenzie and S.B. Wicker. Game Theory in Communications: Motivation, Explanation and Application to Power Control. In *proceedings of IEEE Globecom*, 2001.
4. C. Courcoubetis, R.R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*, Wiley, 2003.
5. B. Tuffin. Charging the Internet without bandwidth reservation: an overview and bibliography of mathematical approaches. *Journal of Information Science and Engineering*, Vol. 19, pages 7-5-786, 2003.
6. P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing Economic Incentives in Peer-to-Peer Networks. *Special Issue on Network Economics, Computer Networks*, Elsevier, 45(1):133-146, 2004.

Outline

General context: 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.
- Illustration: "failure" of ATM networks.
- In such a situation:
  - From the decentralization, there is a general envisaged/advised behavior
  - But each *selfish* user can try to modify his behvior at his benefits and at the expense of the network performance.
  - How to analyze this, and how to control and prevent such a thing?
- It is the purpose of non-cooperative game theory.

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.
- Tool used befor in Economics, Transportation...
- and has recently appeared in telecommunications.
- We may have parodoxes (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 dealy, but intermediate links 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 resubmitting packets?

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/services 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...

All those interacting actors have to be considered.

Outline

1. Introduction and context
2. Basic concepts of game theory
3. Application to routing
4. Application to power control in 3G wireless networks
5. Application to P2P
6. Application to ad hoc networks
7. Application to grid computing
8. A way to control: pricing
9. Interdomain issues
10. Competition among providers
11. Concluding remarks

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, evolutionnary 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
F_1	b_{11} c_{11}	b_{12} c_{12}
F_2	b_{21} c_{21}	b_{22} c_{22}

$N = \{1, 2\},\ A_1 = \{F_1, F_2\},\ A_2 = \{C_1, C_2\},\ u_1(F_j, C_k) = b_{jk},\ u_2(F_j, C_k) = c_{jk}.$

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 N.E 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}, the.
- 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}

  4. Remark: on this example, *dominant strategies* so that the table can be simplified.

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	**2; 1**	0; 0
S	0; 0	**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: (5; 6) and (6; 4).

Prisonner's Dilemna

- Suspects in a crime are in separate cells.
- If they both confess, each will be sentenced to a 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 Dilemna 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!

A Nash equilibrium does not always exist

- Game where 2 players play odd and even:

	Odd	*Even*
Odd	$\mathbf{1}; -1$	$-1; \mathbf{1}$
Even	$-1; \mathbf{1}$	$\mathbf{1}; -1$

- 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:

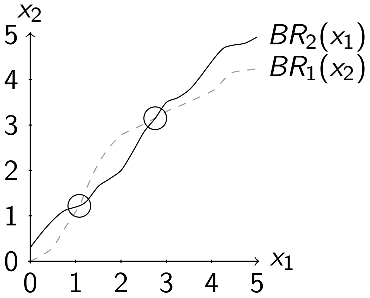

Mixed strategies

- 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: purification. Randomization comes from the lack of knowledge of the agent's information.

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) choose a probability p_t (resp. p_j) to transmit on channel c_1.
- Utilities (average payoff values):

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

- 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 adminsitrator, designer, service provider...) wants to optimize a utility function.
- His utility depends on the reaction of users (who want to maximize their own utility, minimiez 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

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

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 $R(p)$.

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 that 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.

Wardrop equilibrium illustration

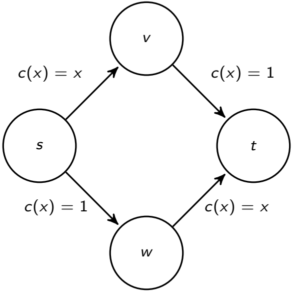

- Two disjoint routes from s to t. Volume of traffic to send: 1.
- Cost functions $c(x)$ on each link associated to traffic volume x.
- How infinitesimal selfish users distribute themselves?
- Wardrop's principle: the cost on each route is the same, otherwise some of them would switch to the other:
- if x_1 on route (s, v, t) and x_2 on route (s, w, t),
  - costs are equal: $1 + x_1 = 1 + x_2$.
  - Give that $x_1 + x_2 = 1$, this gives $x_1 = x_2 = 1/2$.
  - Cost on each route: $3/2$.

Price of Anarchy

- The optimal social utility function happens 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)

It is the ratio of optimal social utility divided by the worst social utility at a Nash equilibrium.

- 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 Anarcy in specific games.

Outline

Routing games

- Users choose their route to send traffic to destination.
- Goal: to minimize transportation cost: delay (pricing can be inserted, see later).
- Two types of games
  - nonatomic routing games, where each player controls a negligible fraction of the overall traffic. Wardrop equilibrium is the proper concept.
  - atomic routing games, where each player controls a nonnegligible amount of traffic. Nash equilibrium here.
- Existence of an equilibrium and uniqueness of cost at each edge proved in the case of nonatomic games.
- existence of an equilibrium proved in specific cases for the atomic case (common value to send; affine cost functions).
- Price of Anarchy can be studied. At most $(3+\sqrt{5})/2 \approx 2.618$ for nonatomic games with affine costs.
- See T. Roughgarden. Routing Games. `http://theory.stanford.edu/~tim/papers/rg.pdf`.

Braess paradox. Total traffic sent: 1

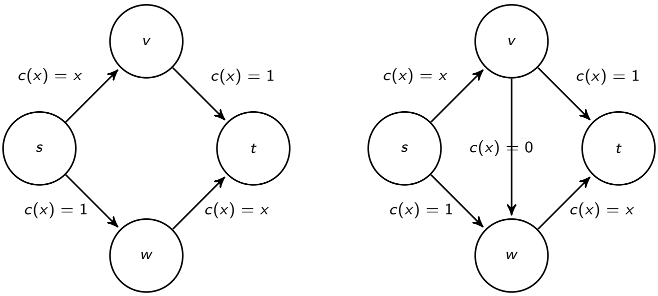

- Left: route costs $1 + x$, split equally at equilibrium, i.e. cost $3/2$.
- Right: expansion of the network, adding a route (cost 0).
- Right: at equilibrium everything on the new route (because never worse than along old routes): cost 2!
- Indeed cost $x + x$ less than $1 + x$ of any other route (since $x \geq 1$)

Outline

Application to power control in 3G networks

- In CDMA-based networks (*Code Division Multiple Access*), each user can play on transmission power.
- Quality of Service (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.

Game for power allocation

- In a one-shot game (strategic game), there is a unique Nash equilibrium.
- The equilibrium is *Pareto inefficient*.
- *Pareto efficiency:* no individual can be made better off without another being made worse off.
- Several proposals to cope with this and improve efficiency:
  - Pricing
  - Repeated games
  - ...
- Specific references:
  - C. Saraydar, N. Mandayam, and D. Goodman, Pricing and power control in a multicell wireless data network, *IEEE JSAC Wireless Series,* vol. 19, no. 2, p. 277-286, 2001.
  - T. Alpcan, T. Basar, R. Srikant, and E. Altman, CDMA uplink power control as a noncooperative game, *Wireless Networks*, 2002.
  - V. Siris, Resource control for elastic traffic in CDMA networks, in *Proc. of MOBICOM'02*, 2002.

Outline

Application to P2P

- Peer-To-Peer (P2P) networks are self-organizing, distributed systems, with no centralized authority or infrastructure.
- Typical candidate for game theory to study the interaction of strategic and rational peers.
- Ultimate goal: propose incentives or to improve the system's performance at the equilibrium of the game.
- In general, rational users are *free riders*: they contribute to little or nothing to the network.
- Different ways to enforce participation:
  - pricing incentives: money awarded when you share your files, and cost when dowloading files of others.
  - reputation incentives: the quality of your participation is dependent of your reputation, which is based on your participation.

P2P, some references

- Some specific references:
  - C. Buragohain, D. Agrawal, S. Suri. A Game Theoretic Framework for Incentives in P2P Systems. (google the title)
  - See also `http://nes.aueb.gr/p2p.html`

Outline

Application to ad hoc networks

- Ad hoc networks: networks without any infrastructure.
- Nodes send their one traffic, but also forward traffic of peers.
- Typical application: military ones, or emergency ones but aimed to be extended to commercial ones.
- Same problem than for P2P: what is the interest of forwarding the traffic of others?
  - Pricing or reputation can be used.
- Simular utility than in 3G networks, with a specificity: power battery.
- Therefore combines both characteristics.
- Some specific references:
  - Shen Zhong, Jiang Chen, Yang Richard Yang. Sprite : A Simple, Cheat- Proof, Credit-Based System for Mobile Ad Hoc Networks. In *Proceedings of IEEE Infocom 2003*. March 2003.
  - Levente Buttyan and Jean-Pierre Hubaux. Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks. *ACM Journal for Mobile Networks (MONET) special issue on Mobile Ad Hoc Networks*. 2002.
  - ...

Outline

Application to grid computing

- Problems similar to P2P: how to yield incentives to participate in grids?
- Some specific references:
  - Grid Economy project: `http://www.gridbus.org/ecogrid/`
  - J. Altmann and S. Routzounis, Economic Modeling of Grid Services, *e-Challenges2006*, Barcelona, Spain, October 2006.
    `http://it.i-u.de/schools/altmann/publications/Economic_Modeling_of_Grid_Services_v09.pdf`
  - Some references at `http://www.zurich.ibm.com/grideconomics/refs.html`

Outline

Pricing for producing incentives. Why changing?

- 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).
- **We also need to provide incentives to participate with a fair use of resources (see all above applications).**
- Properties to be verified:
  - Efficiency (provider's revenue or social welfare)
  - Incentive compatibility (truthful revelation of valuation)
  - Individual rationality (each user's best interest is to participate).

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 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 (fix, 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 $y = \sum_i x_i/k$ 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 increase 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 then 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 acces charges.
  Advantage: simple. Drawback: does not work in a competion market.
- *Cumulus pricing scheme*: +/- points awarded if predefined contract respected. Penalities and renegociations.
  Advantage: easy to implement.
- Priority pricing: classes of traffic with different priority levels and access prices;
  - schedulling priority
  - rejection or dropping priority.

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

Example: pricing and schedulling

- Goal of DiffServ architecture: to introduce differentiation of service by providing mutiple classes
  - introduced to deal with congestion
  - because applications are more or less stringent in terms of QoS.
- If no pricing associated to DiffServ, all users/applications will likely choose the "best" service class.
- DiffServ architecture deals with strict priority or generalized processor sharing.
  Which one is the "best" from an economical point of view?
- Questions to solve:
  - For eah schedulling policy, what are the prices maximizing the provider's benefits?
  - Which schedulling policy to implement? I.e., which one yields larger benefits (at optimal prices)?

Basic model

- Bottleneck node of the network represented by an M/M/1 queue with service rate μ.
- Infinite number of potential users, users being assumed *infinitesimal*
- Two types of flows: voice and data (voice more sensitive to delay than data), with rate λ_v, λ_d per user.
- Two classes of service with possible schedulling policies:
  - strict priority:
    - ★ class-1 always served before class-2
  - generalized processor sharing (GPS):
    - ★ a part of the server is dedicated to class-1, the other to class-2, except when no server in one class (full service then)
    - ★ FIFO scheduling within a class
  - or discriminatory processor sharing (DPS):
    - ★ a weight w_i (corresponding to its class) associated to a flow i
    - ★ a proportion $w_i/(\sum_j w_j)$ of the server is allocated to flow i.
- Cases of dedicated classes or open classes
  - we restrict ourselves to dedicated classes here.

User behaviour

- Utility depending on the *average delay* D and per-packet price p:

$$U_d(D) = D^{-\alpha_d} - p \text{ and } U_v(D) = D^{-\alpha_v} - p$$

  where $\alpha_d < \alpha_v$: voice users have preference for small delays.
- A user enter as soon as his utility is positive, or leaves if it is negative $\Rightarrow$ Game between classes on the steady state number of active connections (users).
  - The number of users N_d and N_v in one class may influence the number in the other class.
  - Prices influence that number too.
- At (Wardrop) equilibrium, $\forall j \in \{v, d\}$:
  - either $N_j > 0$ and $U_j(D) = 0$
  - or $N_j = 0$ and $U_j(D) \leq 0$.

Schedulling Policies considered: priority, GPS, DPS

- If only one class, average response time when N users: $D = 1/(\mu - N\lambda)$.
- Priority: closed form available for delay per class, with higher priority for voice users:

$$D_v = \frac{1}{\mu - N_v\lambda_v} \quad \text{and } D_d = \frac{\mu}{(\mu - N_v\lambda_v)(\mu - N_v\lambda_v - N_d\lambda_d)}.$$

- GPS: no closed-form formula. Though, under heavy load assumption, can be approximated by independent queues (similar to so-called Paris Metro Pricing). If γ_v and γ_d proportions allocated to v and d:

$$D_v = \frac{1}{\gamma_v\mu - N_v\lambda_v} \quad \text{and } D_d = \frac{1}{\gamma_d\mu - N_d\lambda_d}.$$

- DPS: closed-form formula also. If γ relative priority of data users,

$$D_v = \frac{\left(1 + \frac{\lambda_d N_d(2\gamma-1)}{\mu-(1-\gamma)\lambda_v N_v-\gamma\lambda_d N_d}\right)}{\mu - \lambda_v N_v - \lambda_d N_d} \quad \text{and } D_d = \frac{\left(1 - \frac{\lambda_v N_v(2\gamma-1)}{\mu-(1-\gamma)\lambda_v N_v-\gamma\lambda_d N_d}\right)}{\mu - \lambda_v N_v - \lambda_d N_d}.$$

Dedicated classes, strict priority

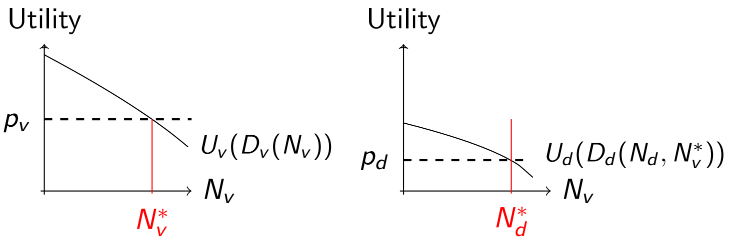

- High priority user demand N_v^* computed first:
  - N_v increases up to $U_v(D_v) = \left(\frac{1}{\mu - N_v\lambda_v}\right)^{-\alpha_v}$ decreases to p_v;
  - If N_v too large and $U_v(D_v) < p_v$, then N_v naturally decreases.
  - it gives $N_v^* = \frac{\mu - p_v^{-\alpha_v}}{\lambda_v}$.
- Next, with this value of N_v^*, N_d^* computed similarly, solution of

$$U_d(N_d, N_v^*) = \left(\frac{\mu}{(\mu - \lambda_v N_v^*)(\mu - \lambda_v N_v^* - \lambda_d N_d)}\right)^{\alpha^d} = p_d.$$

- User equilibrium easily *explicitely* characterized.

Dedicated classes, GPS

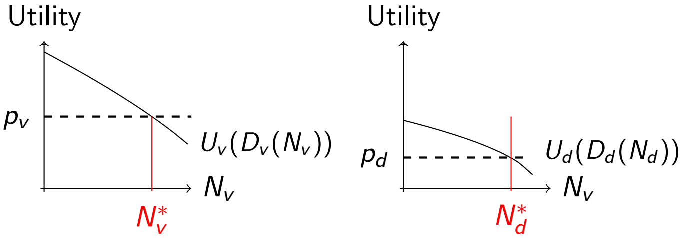

- Both queues considered *independently*. $\forall j \in \{v, d\}$,
  - N_j increases up to $U_j(D_j) = \left(\frac{1}{\gamma_j \mu - N_j \lambda_j}\right)^{-\alpha_j}$ decreases to p_j;
  - If N_j too large and $U_j(D_j) < p_j$, then N_v naturally decreases.
  - it gives

$$N_j^* = \frac{\mu - p_j^{-\alpha_j}}{\lambda_j}.$$

Open classes, strict priority

- For the high priority class 1
  - respective utilities $U_v = D^{-\alpha_v} - p_1$ and $U_v = D^{-\alpha_d} - p_1$.
  - If $p_1 > 1$, curve $U_v = 0$ always above $U_d = 0$;
  - If $p_1 < 1$ $U_v = 0$ always under $U_d = 0$.

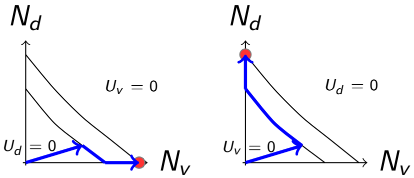

- Only voice (resp. date) users in class 1 if $p_1 > 1$ (resp. $p_1 > 1$).
- Similar results for low priority class.
- Four situations with easy chracterization of (N_v^*, N_d^*):
  - $p_1, p_2 > 1$: only voice users
  - $p_1, p_2 > 1$: only data users
  - $p_1 > 1$, $p_2 < 1$: voice users in class 1, data users in class 2
  - $p_1 < 1$, $p_2 > 1$ (strange!): data users in class 1, voice users in class 2.

Open classes, GPS

- Same analysis that with the highest queue with strict priority, cinsidering both queues separately.
- Four situations with easy explicit characterization of (N_v^*, N_d^*):
  - $p_1, p_2 > 1$: only voice users
  - $p_1, p_2 > 1$: only data users
  - $p_1 > 1$, $p_2 < 1$: voice users only in class 1, data users only in class 2
  - $p_1 < 1$, $p_2 > 1$: data users only in class 1, voice users only in class 2.

Economic issues

Results

Hayel, Ros & T., Infocom 04

- Prices optimizing the network revenue found for each policy using the user equilibrium:
  - Revenue defined as

$$\begin{aligned} R &= R_v + R_d \\ &= \lambda_v N_v^* p_v + \lambda_d N_d^* p_d \end{aligned}$$

  - simple derivation applied each time in terms of prices;
  - optimal revenue computed then.
- Policy that produces the best revenue: **strict priority**: $\gamma_1 \in \{0, 1\}$ optimal in terms of revenue for the GPS case.
  - for dedicated classes
  - and open classes as well.

Dedicated classes, DPS; dynamics

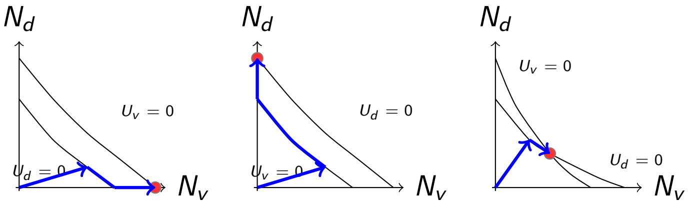

- The value of N_j influences directly the utility of the other class i.
- Three possible situations
  - One curve U_i is always below the other (two cases)
    - ★ The numbers of customers increase up to reaching the lowest curve $U_i = 0$
    - ★ but N_j still increases ($U_j > 0$), it slides on the curve to $N_i = 0$ on $U_i = 0$
    - ★ the on the axis to the equilibrium point $N_i = 0$ and $U_j = 0$.
  - The curves have an intersection point
    - ★ The number of customers increase up to reaching one curve;
    - ★ Then thit slides up to the intersection point.

Remark: DPS and TCP modelling

- DPS not applicable at the packet level.
- Though, DPS in an M/M/1 queue is a good approximation of interactions of TCP sessions in comptetion at the flow level.
- The results remain valid, but the λ are here for session lengths, and the number of sessions are considered *in average*.
- It therefore provides a pricing scheme for TCP sessions.

Example: auctionning for bandwidth

The problem of resource allocation

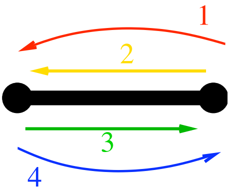

- Allocate bandwidth among users on a link with a capacity constraint Q
- More general results also obtained
- Allocation and pricing mechanism: determines the allocation a_i for each player i, and the price c_i he is charged.

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

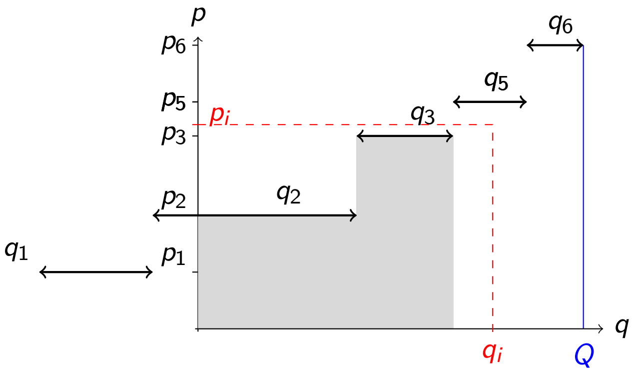

- bid (q_i, p_i) does not allows i to get the required quantity.
- Bids with higher price are allocated first.
- Player i gets *what remains*.
- Charge: loss declared by i's presence (here players 2 and 3); grey zone.

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 & T., 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 $\bar{d}_i$

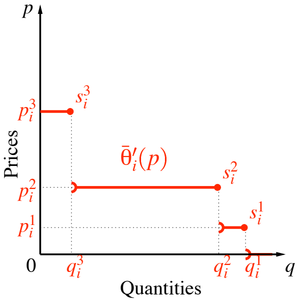

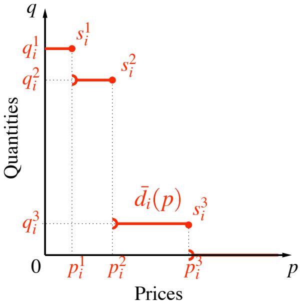

Pseudo-marginal valuation and pseudo-demand functions associated with the multi-bid s_i

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

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

Allocation and pricing rule

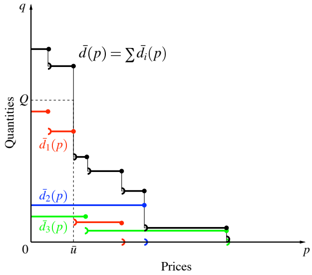

$\bar{u}$: pseudo market clearing price (highest unit price at which demand exceeds capacity).

- 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

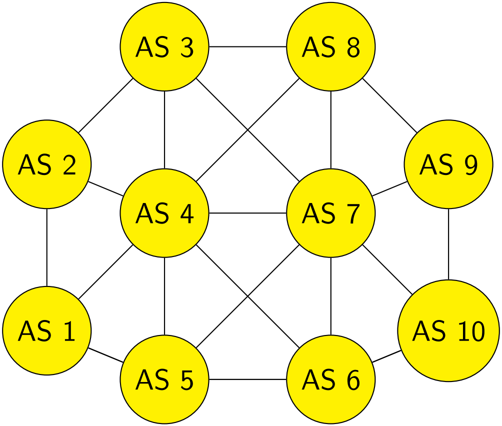

- Network made of Autonomous Systems (ASes) acting selfishly.

Interdomain problem

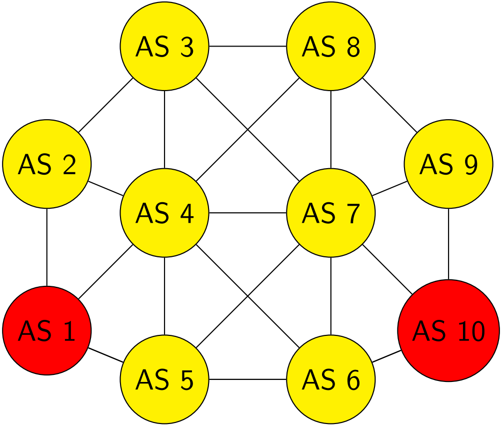

- Network made of Autonomous Systems (ASes) acting selfishly.
- A node (an AS) needs to send traffic from its own customers to other ASes.
- Introduce incentives for intermediate nodes to forward traffic , via pricing.

Interdomain problem

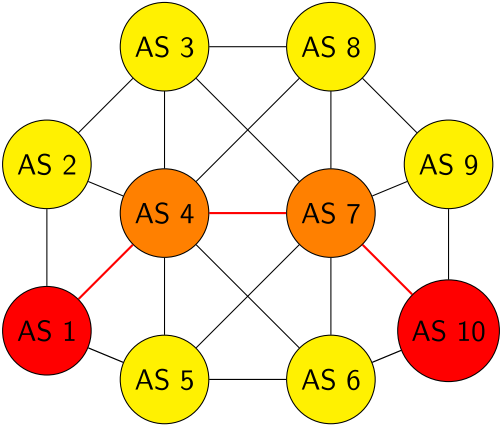

- Network made of Autonomous Systems (ASes) acting selfishly.
- A node (an AS) needs to send traffic from its own customers to other ASes.
- Introduce incentives for intermediate nodes to forward traffic , via pricing.
- What is the best path?

Interdomain problem

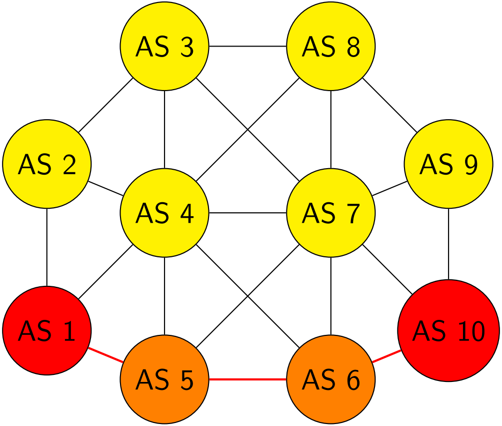

- Network made of Autonomous Systems (ASes) acting selfishly.
- A node (an AS) needs to send traffic from its own customers to other ASes.
- Introduce incentives for intermediate nodes to forward traffic , via pricing.
- What is the best path?

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 contacts all potential ASes on a path to learn their costs, and then make its decisions.
  - 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.
- 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 when 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 amore flexible choice.
  - Strict requirement: budget balance. Decentralization too if dealing with large topologies.

Outline

Specific model of competition among providers

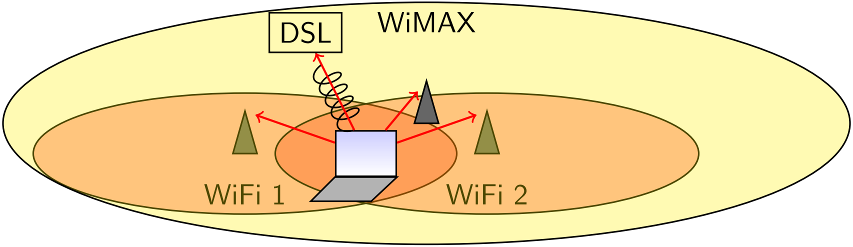

Specific model of competition among providers

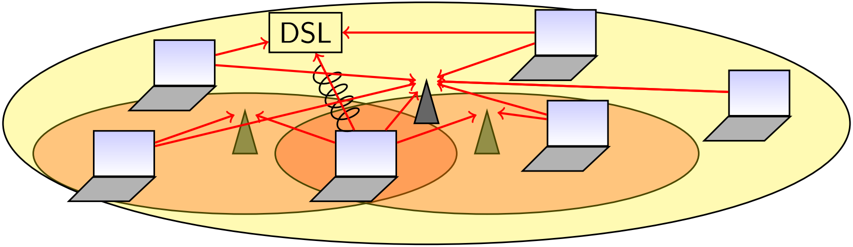

- Interactions among non-cooperative consumers: *game*
- Congested networks provide poorer quality (packet losses)

Specific model of competition among providers

But **providers** play first!

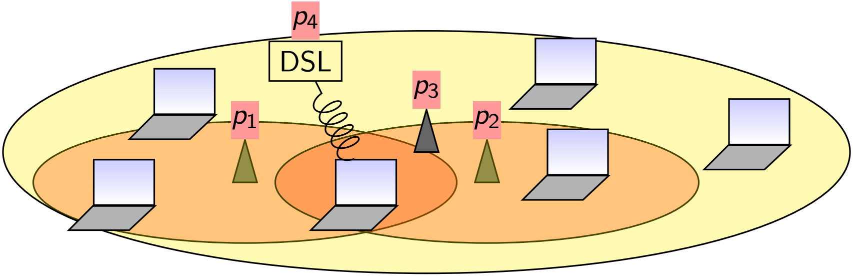

Specific model of competition among providers

But **providers** play first!

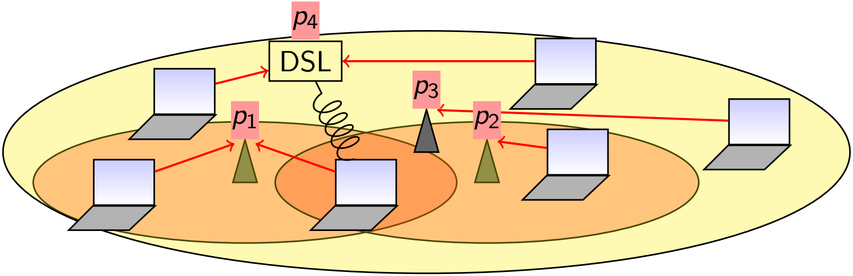

Study of the two-level noncooperative game.

1. *Higher level*: **providers** set prices to maximize revenue
2. *Lower level*: **consumers** choose their provider

Communication model: packet losses

- Time is slotted
- Each provider i has finite capacity C_i
- If total demand d_i at provider i exceeds C_i: exceeding packets are *randomly* lost

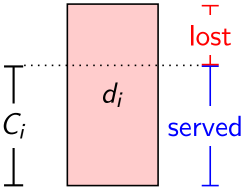

$$\mathbb{P}(\text{successful transmission}) = \min\left(1, \frac{C_i}{d_i}\right)$$

$$\Rightarrow \text{Expected number of transmissions} = \frac{1}{\mathbb{P}(\text{success})} = \max\left(1, \frac{d_i}{C_i}\right)$$

Only “regulation”: pay for what you send

The price p_i at each provider i is per packet *sent* Marbach’02
$\Rightarrow$ If several transmissions are needed, the user pays several times

$$\bar{p}_i := \textit{perceived} \text{ price at } i = \mathbb{E}[\text{price per packet}] = p_i \max\left(1, \frac{d_i}{C_i}\right)$$

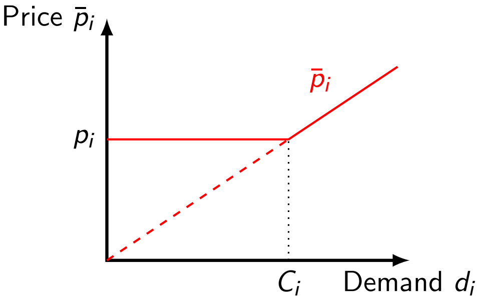

Model for user choices: Wardrop equilibrium

- Users choose the provider(s) i with lowest $\bar{p}_i = p_i \max\left(1, \frac{d_i}{C_i}\right)$

$\Rightarrow$ For a given coverage zone Z, all providers with customers from that zone end up with the same perceived price $\bar{p}_i = \bar{p}_z$ Wardrop'52

Model for user choices: Wardrop equilibrium

- Users choose the provider(s) i with lowest $\bar{p}_i = p_i \max\left(1, \frac{d_i}{C_i}\right)$

$\Rightarrow$ For a given coverage zone Z, all providers with customers from that zone end up with the same perceived price $\bar{p}_i = \bar{p}_z$ Wardrop'52

- The total amount of data that users want to successfully transmit in a zone z depends on that price:

$$\sum_i d_{i,z} \min(1, C_i/d_i) = \alpha_z D(\bar{p}_z),$$

$$i.e. \quad \bar{p}_z = \underbrace{v}_{\text{marg. val. function}}\left(\frac{\sum_i d_{i,z} \min(1, C_i/d_i)}{\alpha_z}\right)$$

with D the total demand function, α_z the population proportion in zone z, and $d_{i,z}$ the demand in zone z for provider i.

Higher level: price competition game

- Providers set their price p_i *anticipating users reaction*
  $\Rightarrow$ Providers are Stackelberg leaders
- We can assume management costs of the form $\underbrace{\ell_i(d_i)}_{\text{nondecreasing, convex}}$

  Provider i's objective: $R_i := p_i d_i - \ell_i(d_i)$.

Competition model

- Simplified topology: common coverage area
- N competing providers declaring price and capacity ($\mathcal{I} := \{1, \ldots, N\}$)

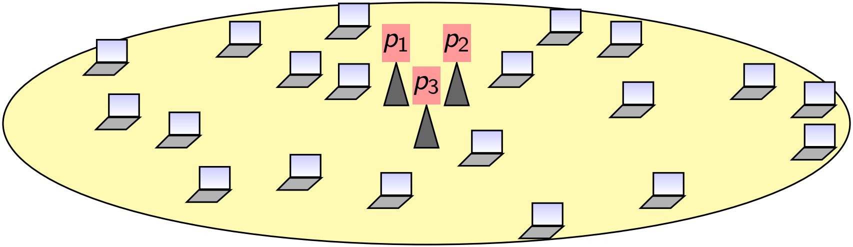

User equilibrium

- Users choose the provider(s) i with lowest $\bar{p}_i = p_i \max\left(1, \frac{d_i}{C_i}\right)$

$\Rightarrow$ All providers with customers end up with the same perceived price $\bar{p}_i = \bar{p}$ Wardrop'52

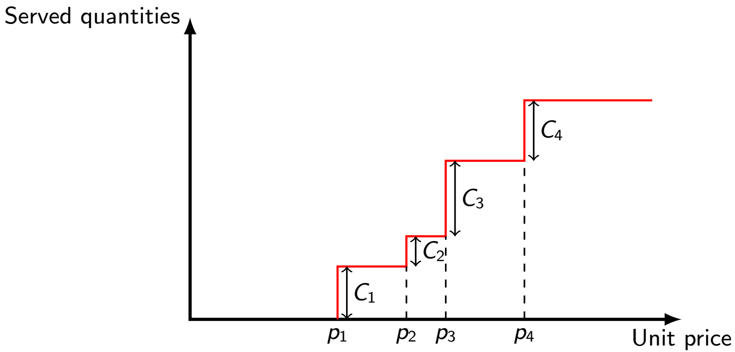

User equilibrium

- Users choose the provider(s) i with lowest $\bar{p}_i = p_i \max\left(1, \frac{d_i}{C_i}\right)$

$\Rightarrow$ All providers with customers end up with the same perceived price $\bar{p}_i = \bar{p}$ Wardrop'52

- The total demand level depends on that price:

$$\bar{p} = \underbrace{v}_{\text{marg. val. function}}\left(\sum \min(C_i, d_i)\right)$$

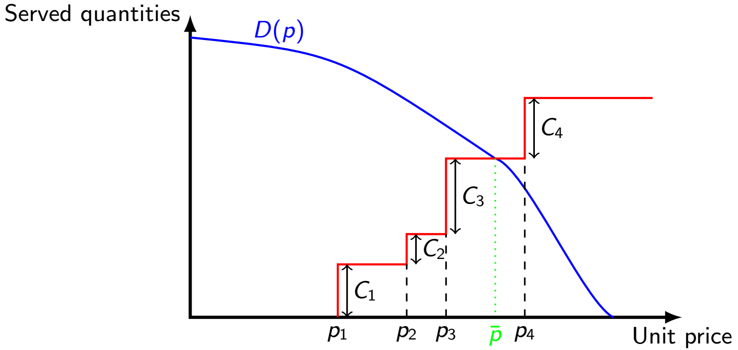

User equilibrium: formal description

$$\begin{aligned}
\bar{p}_i &= p_i \max\left(1, \frac{d_i}{C_i}\right) \\
\bar{p}_i > \min_j \bar{p}_j &\Rightarrow d_i = 0 \\
\sum_i \underbrace{d_{i,z} \min(1, C_i/d_i)}_{\text{effectively received at } i} &= D(\min_j \bar{p}_j).
\end{aligned}$$

Proposition

There exist a (possibly not unique) user (Wardrop) equilibrium demand configuration. The common perceived unit price $\bar{p}$ of providers i with $d_i > 0$ is unique and equals

$$\bar{p} = \min\{p : D(p) \leq \sum_i f_i(p)\},$$

where $f_i(p) = C_i 1_{\{p \geq p_i\}}$, *with* 1_X *indicator function.*

Non-uniqueness happens only when several providers have price $p_i = \bar{p}$: users can choose indifferently those providers.

Price competition, main result

Proposition

Under sufficient condition A, there exists a **unique Nash equilibrium** *on price war among providers, given by*

$$\forall i \in \mathcal{I}, \quad \begin{cases} p_i &= v\left(\sum_{j\in\mathcal{I}} C_j\right) \\ d_i &= C_i. \end{cases}$$

- **Sufficient condition A:** each ℓ_i is Lipschitz with constant κ_i, and $\forall y \geq p^* := v\left(\sum_{j\in\mathcal{I}} C_j\right)$, the demand function D is sufficiently elastic:

$$\frac{-yD'(y)}{D(y)} \geq \frac{1}{1-\kappa/y}, \quad (1)$$

  where $\kappa := \max_{i\in\mathcal{I}} \kappa_i$.
- Without cost functions, it just means a demand elasticity larger than -1.

Price competition, main result

Proposition

Under sufficient condition A, there exists a **unique Nash equilibrium** *on price war among providers, given by*

$$\forall i \in \mathcal{I}, \quad \begin{cases} p_i &= v\left(\sum_{j \in \mathcal{I}} C_j\right) \\ d_i &= C_i. \end{cases}$$

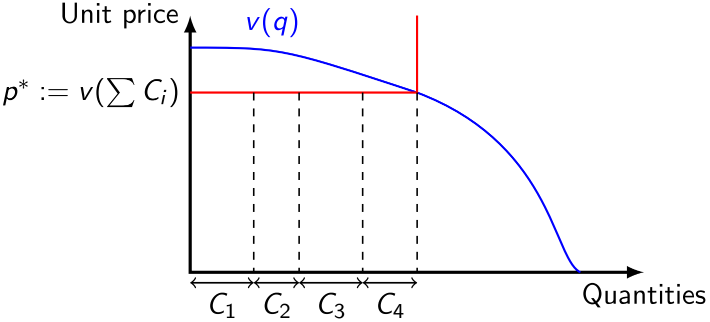

Social Welfare considerations

- A performance measure of the outcome $(d_1, ..., d_I)$ of the game $=$ overall value of the system

$$\text{Social Welfare} := \int_{u=0}^{\sum_{i\in\mathcal{I}} d_i} \frac{\sum_{i\in\mathcal{I}} \min(d_i, C_i)}{\sum_{i\in\mathcal{I}} d_i} v(u)du - \sum_i \ell_i(d_i).$$

First term: total valuation for the service experienced. Comes from actual (per traffic unit) utility of a user having (per traffic unit) willingness-to-pay v is its willingness-to-pay times the probability to be served, i.e.,

$$\frac{\sum_{i\in\mathcal{I}} \min(d_i, C_i)}{\sum_{i\in\mathcal{I}} d_i} v.$$

- **Remark:** the Social Welfare maximization problem leads to the same outcome $d_i = C_i \quad \forall i$ as the price war.
- **Consequence:** The Nash equilibrium corresponds to the socially optimal situation: the Price of Anarchy is 1!.

Game on declared capacities: a third level

We now consider a 3-stage game:

1. Providers $i \in \mathcal{I}$ declare their capacity C_i
2. Providers fix their selling price p_i
3. Users select their providers

Opposite effects of lowering one's capacity:

- the unit selling price at equilibrium increases and the managing cost decreases because the quantity sold decreases
- whereas on the other hand less quantity sold means less revenue.

Proposition

Under the same conditions about **demand elasticity**, *no provider can increase its revenue by artificially lowering its capacity.*

Competition model

Assumptions

- Two competing providers declaring price and capacity
- One coverage area included in the other

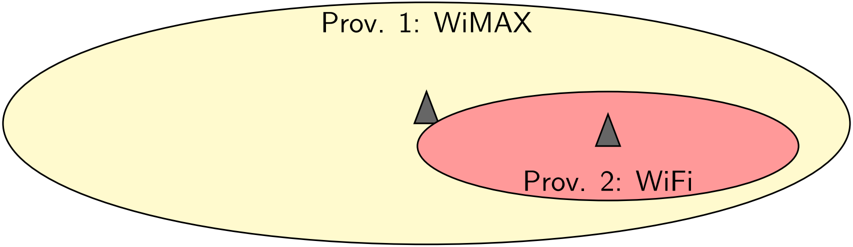

Competition model

Assumptions

- Two competing providers declaring price and capacity
- One coverage area included in the other

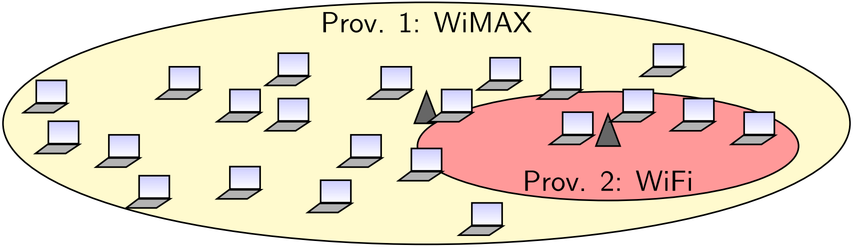

User equilibrium: illustration

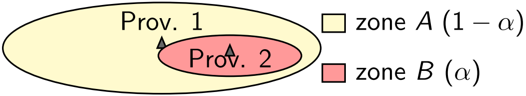

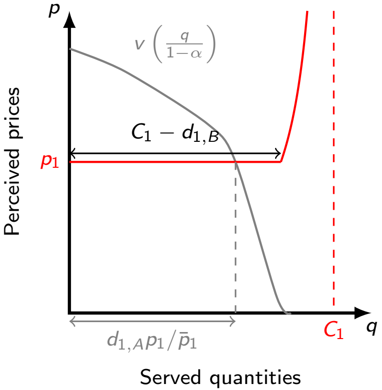

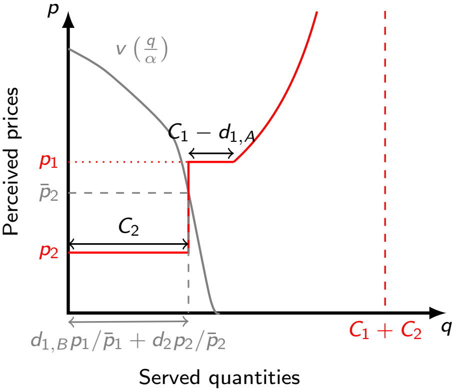

User equilibrium: mathematical formulation

At user equilibrium, according to Wardrop principle

$$\bar{p}_1 = p_1 \max\left(1, \frac{d_{1,A}+d_{1,B}}{C_1}\right)$$
$$\bar{p}_2 = p_2 \max\left(1, \frac{d_2}{C_2}\right)$$

$$d_{1,A}\min\left(1, \frac{C_1}{d_{1,A}+d_{1,B}}\right) = (1-\alpha)D(\bar{p}_1)$$
$$d_{1,B}\min\left(1, \frac{C_1}{d_{1,A}+d_{1,B}}\right) + d_2\min(1, C_2/d_2) = \alpha D(\min(\bar{p}_1, \bar{p}_2))$$

$$\bar{p}_1 > \bar{p}_2 \Rightarrow d_{1,B} = 0$$
$$\bar{p}_1 < \bar{p}_2 \Rightarrow d_2 = 0.$$

User equilibrium: existence and uniqueness

Proposition

For all price profile, there exists at least a user (Wardrop) equilibrium. Moreover, the corresponding perceived prices of each provider are unique.

NB: demand repartition among providers is not necessarily unique.

Higher level: price competition game

- Provider i's objective: $R_i := p_i d_i - \ell_i(d_i)$.

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$, *then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

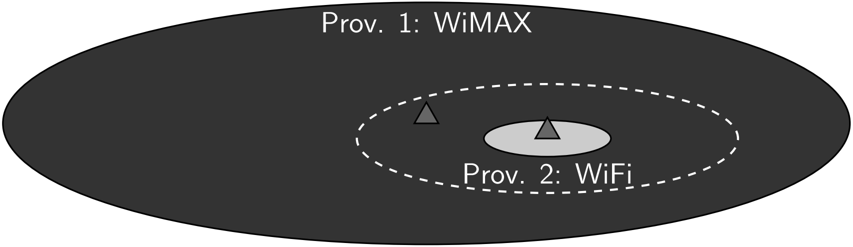

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$*, then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

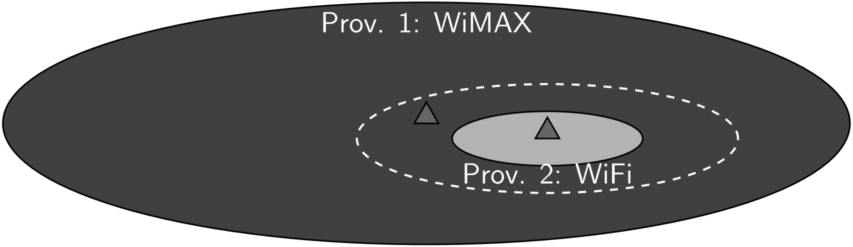

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$, *then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

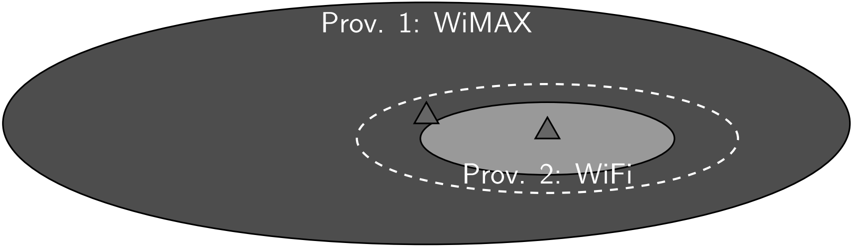

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$, *then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

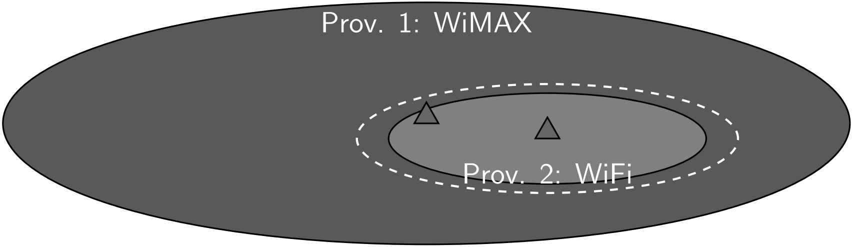

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$, *then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

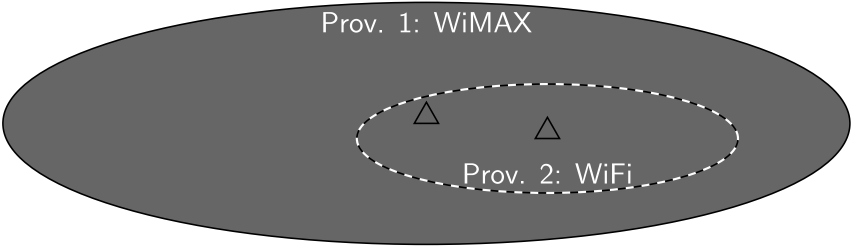

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$, *then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

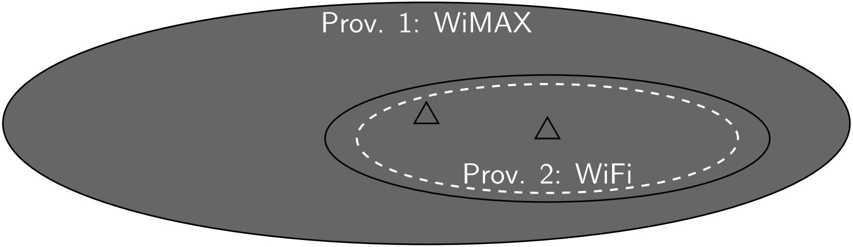

Proposition

If $-\frac{D'(p)p}{D(p)} > 1, \quad \forall p$ *(elastic demand), then there exists a unique Nash equilibrium* (p_1^*, p_2^*) *in the price war between providers.*

- *If* $\alpha \leq \frac{C_2}{C_1+C_2}$, *then* $p_1^* = v\left(\frac{C_1}{1-\alpha}\right) \geq p_2^* = v\left(\frac{C_2}{\alpha}\right)$. *The common zone is left to provider* 2 *by provider* 1.
- *If* $\alpha > \frac{C_2}{C_1+C_2}$ *then* $p_1^* = p_2^* = p^* = v(C_1 + C_2)$. *The common zone is shared by the providers.*

(Darker=more expensive)

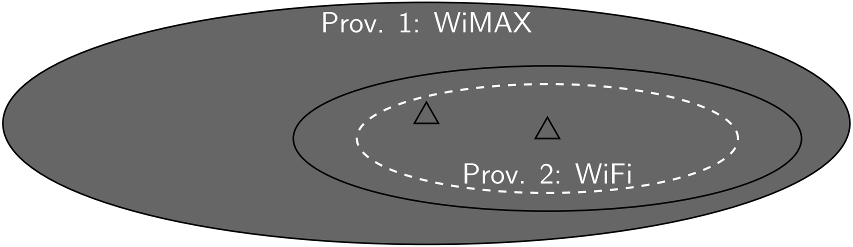

Outline

Concluding remarks

- Game Theory has gained a lot of attention in the networking community (see the number of related publications in major conferences such as IEEE Infocom).
- It allows to model and study the behavior of selfish users in competition for resources.
- We can then play on parameters of the model to drive the equilibrium to a better point.
- Applications in all areas of networking.
- Pricing is a typical (and quite natural) way to yield proper incentives.