Rien ne va plus?

Game Theory and the Internet

Stefan Schmid

Thanks to Yvonne Anne Pignolet (ETH Zurich)
for collaborating on some slides!

TECHNISCHE
UNIVERSITÄT
MÜNCHEN

Wroclaw Information
Technology Initiative (2008)

Game Theory for Machines?!

- Case study „peer-to-peer computing“

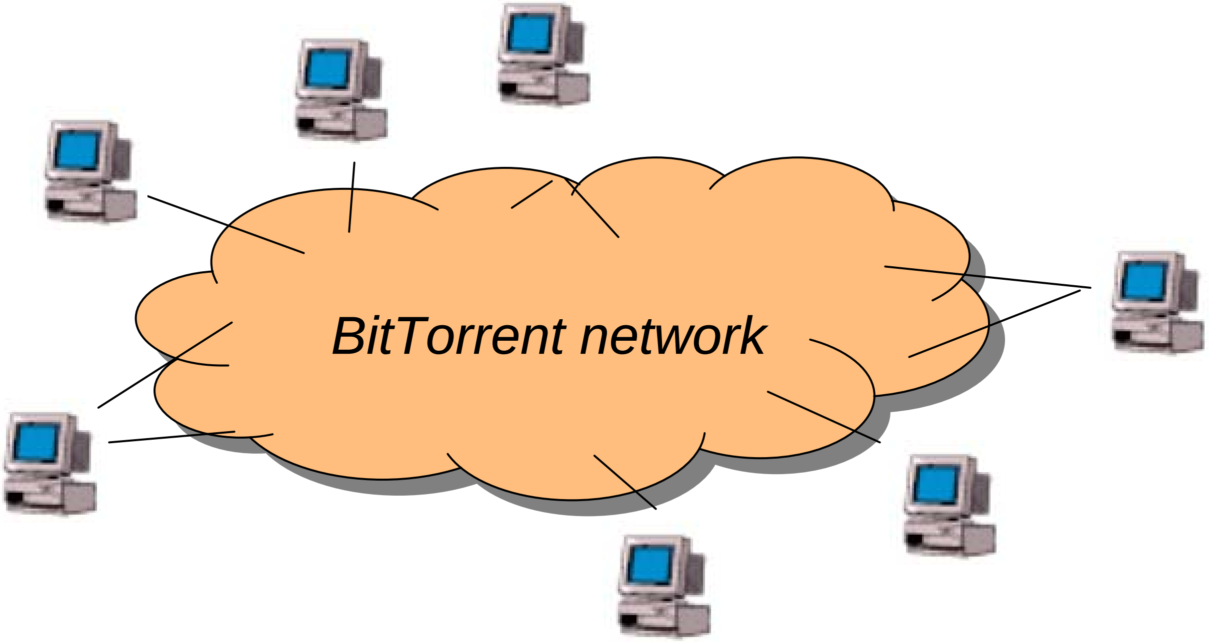

Game Theory for Machines?!

- Can machines be selfish?

Yes!

- Machine and software is under the control of the user!
- Users can implement own client (e.g., BitThief),
- may remove all files from their shared folder,
- can cap their upload channel,
- may only be online when they download something themselves,
- etc.!

Game Theory for Machines?!

- Model for a peer-to-peer network?

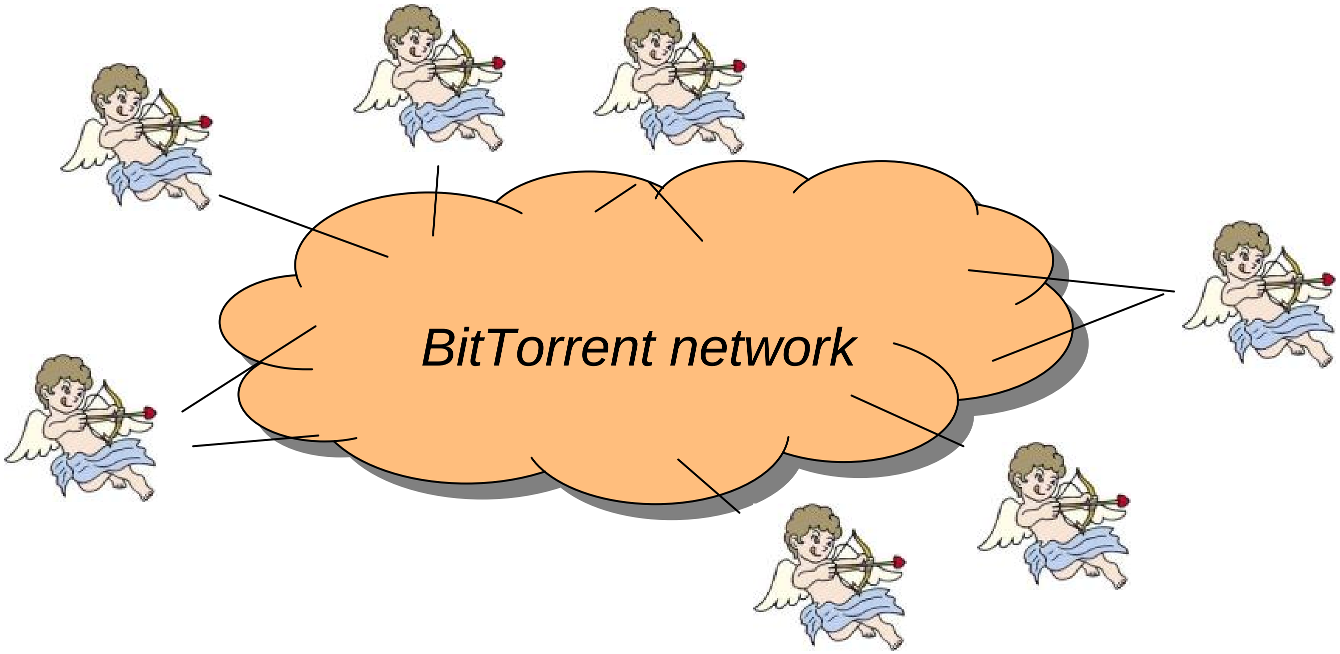

Game Theory for Machines?!

- Model for a peer-to-peer network?

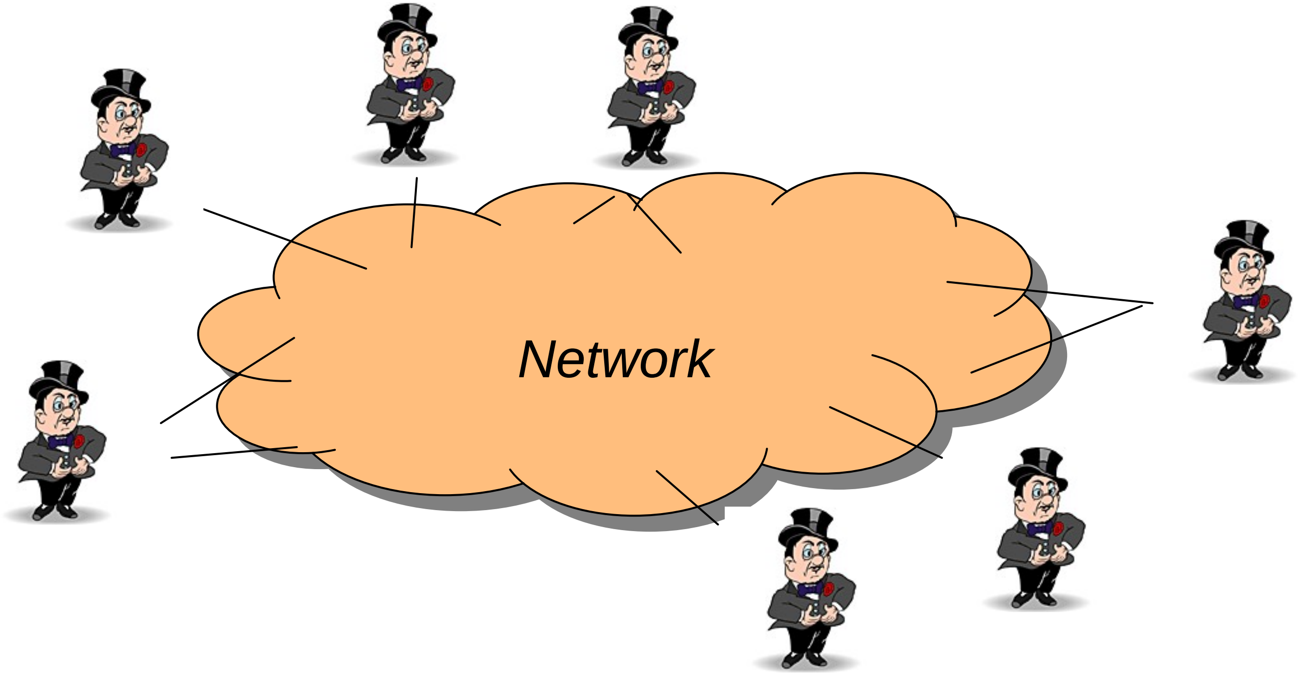

Game Theory for Machines?!

- Model for a peer-to-peer network?

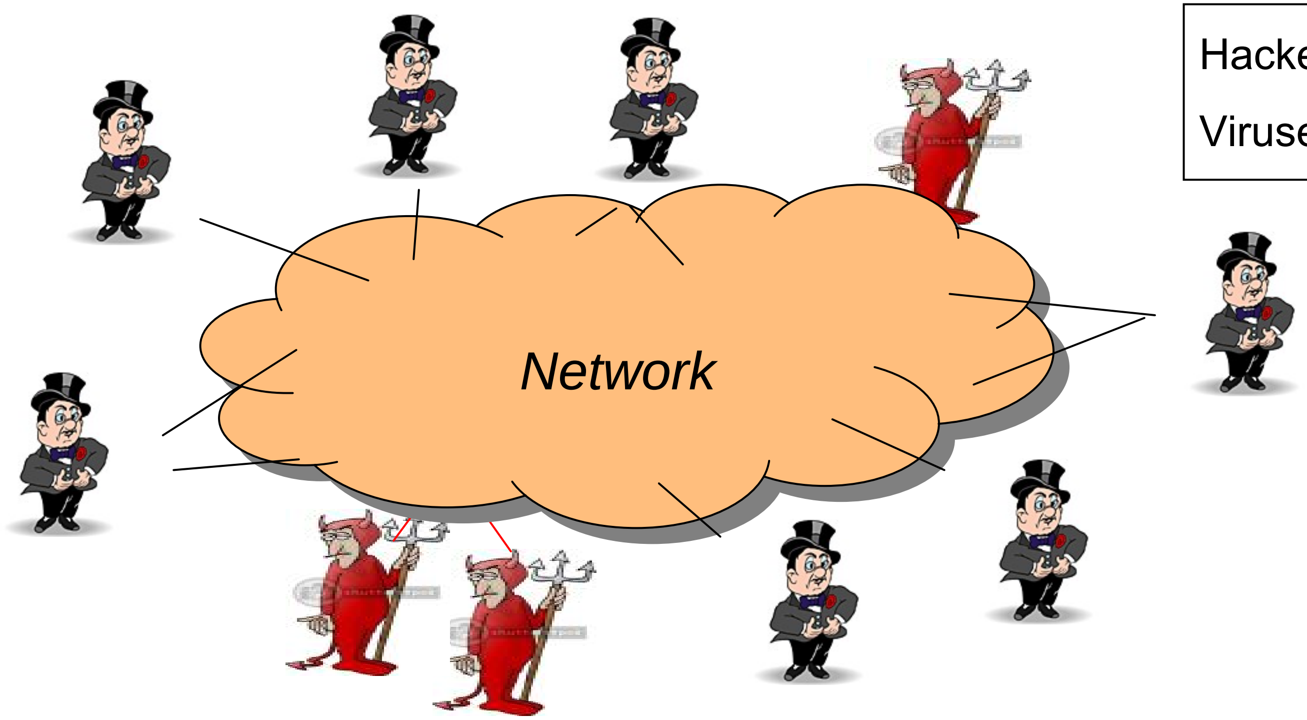

Hackers, Polluters,

Viruses, DOS attacks

Influential Talk by Ch. Papadimitriou (STOC 2001)

- Last 50 years:
  Theoretical CS studied von Neumann machines (plus software)
  with logic and combinatorics

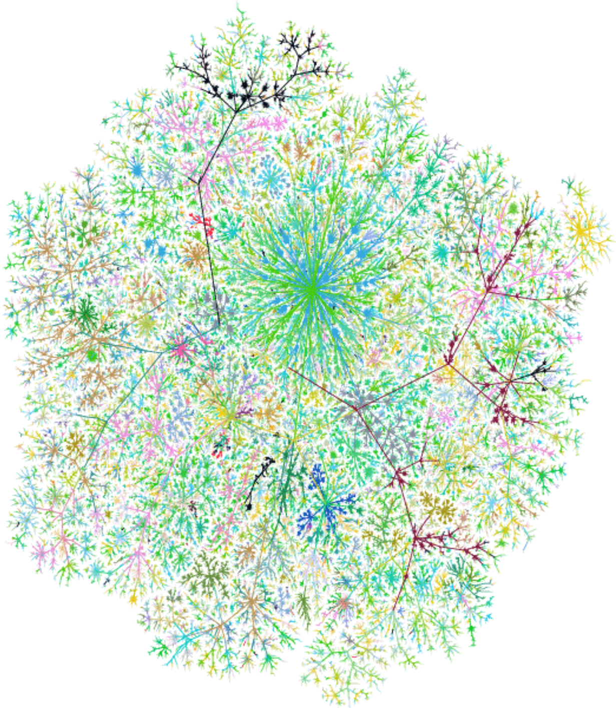

- Internet has become most complex computational artifact of our time
  - characteristics: size and growth,
    spontaneous emergence,
    availability, universality, ...
  - but most importantly:
    socio-economic complexity

Influential Talk by Ch. Papadimitriou (STOC 2001)

„The Internet is unique among all computer systems in that it is built, operated, and used by a multitude of diverse economic interests, in varying relationships of collaboration and competition with each other.“

„This suggests that the mathematical tools and insights most appropriate for understanding the Internet may come from a fusion of algorithmic ideas with concepts and techniques from Mathematical Economics and Game Theory.“

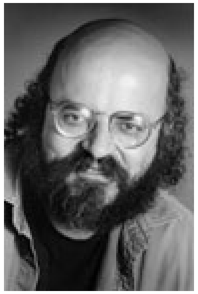

Game Theory

- Main objective: Understand rational behavior

- Model:
  - n players
  - each player i has choice of one strategy in S_i
  - utility given by all the players‘ strategies, i.e., player i has utility u_i: $S_1 \times ... \times S_n \in \mathbf{R}$

Notion of Equilibrium

- Concept of rationality: equilibria

- Most importantly: Nash equilibria

A combination of strategies $x_1 \in S_1 \times \ldots \times x_n \in S_n$ is called a Nash equilibrium if no player can be better off by unilaterally (= given strategies of other players) changing her strategy.

The Most Popular Game?

- Prisoner‘s Dilemma

- Two alleged bank robbers Dave and Henry
  - both are interrogated individually
  - options: confess bank robbery or not
  - if A confesses and B does not, then A only needs to go to prison for one year, and B five
  - if nobody confesses, they can still charged for a minor crime (2 years each)
  - if both confess, they get 3 years each

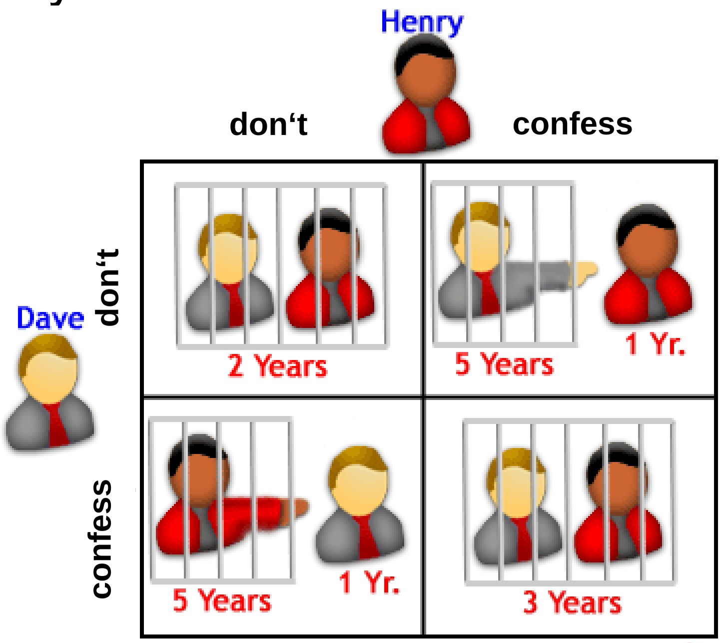

- Utilities / disutilities can be described by a payoff matrix

The Most Popular Game?

- Dilemma:
  - given the other player‘s choice, it‘s always better to confess (1 yr instead of 2 yrs, and 3 yrs instead of 5 yrs)
  - *(c,c)* constitutes a Nash equilibrium
  - in fact, *c* is even a dominant strategy

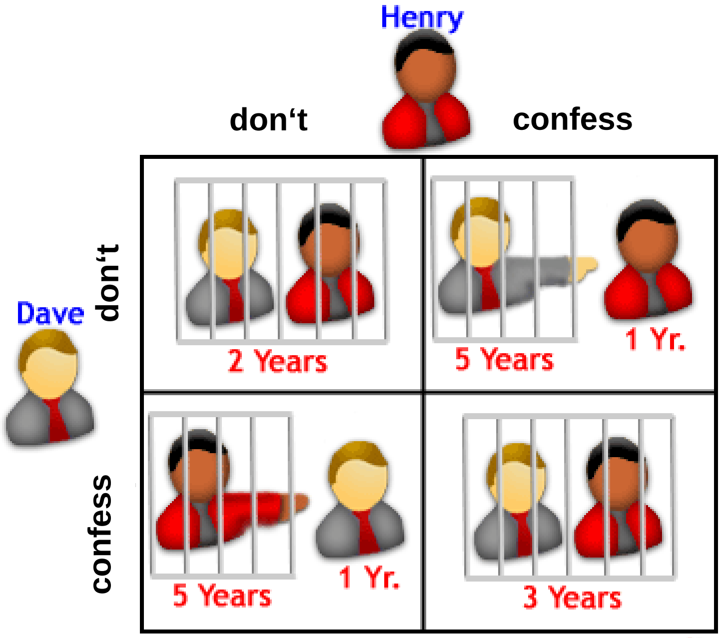

- Thus, in the Nash equilibrium, both players go 3 years to prison (= 6 years in total, = social cost of Nash equilibrium)

- However, if they were smart, they could have got away with 4 years in total (= social optimum)

Comment on Nash Equilibria

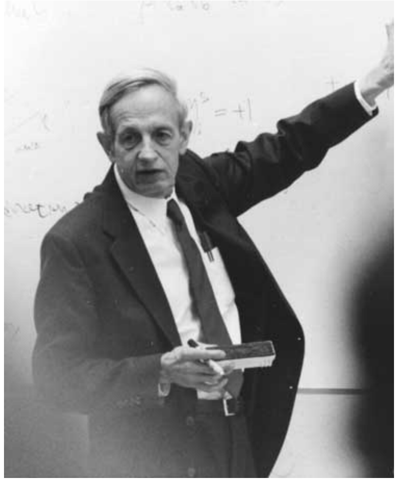

- Games can have more than one Nash equilibrium
  - unclear which to consider
- One distinguishes between pure (deterministic) and mixed equilibria
  - pure equilibria do not always exist...
  - ... but mixed do! *[Nash 1952]*
- Finding (pure and mixed) equilibria can be very time-consuming
  - good model for real economies?
  - **„Together with factoring, the complexity of finding a Nash equilibrium is in my opinion the most important open question on the boundary of *P* today.“**

Alternative Concepts: Pareto Efficiency

- Pareto improvement: Transition to outcome where >=1 player increases payoff, no payoff decrease for anyone
- Pareto optimum (PO): no Pareto improvement possible

Example:

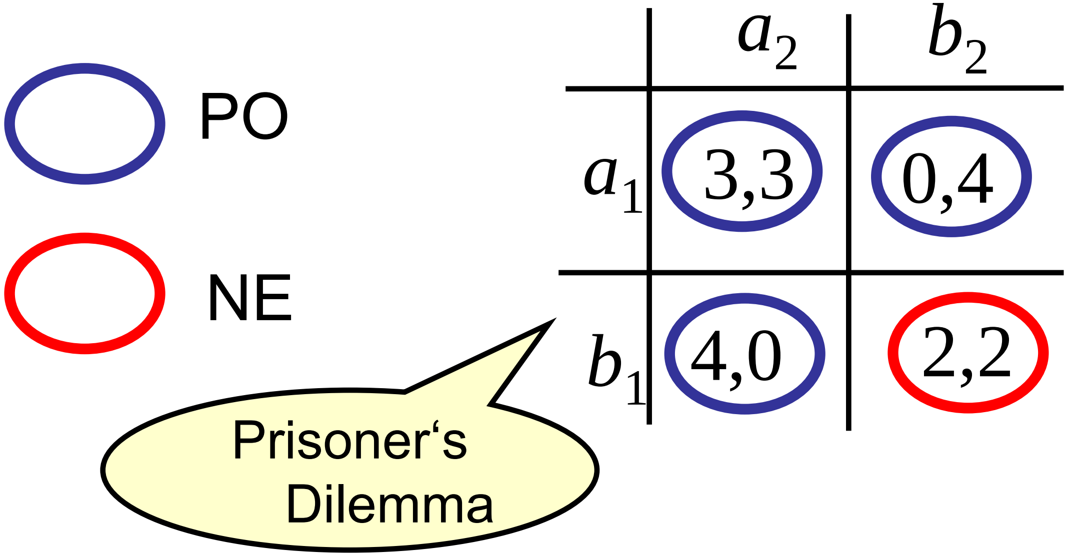

	a_2	b_2
a_1	3,3	0,4
b_1	4,0	2,2

Criticism: not necessarily optimal for society! (unjust and inefficient outcomes)

Alternative Concepts: Correlated Equilibrium

Example:

	C_2	D_2
C_1	6,6	2,7
D_1	7,2	0,0

○ Pure NE

Mixed NE?

Worse than C,C !!

Player 2 probability to choose „Dare“: $p_2 = 1/3$

E[payoff player_1] $< 6 \cdot (1-p_1) \cdot 2/3 + 2 \cdot (1-p_1) \cdot 1/3 + 7 \cdot p_1 \cdot 2/3$

=> maximal if player 1 dares with $p_1 = 1/3$ too => equilibrium

Mixed NE: $p_1 = p_2 = 1/3$ E[payoff]: 4.6

Correlation could help a lot!

Alternative Concepts: Correlated Equilibrium

Example:

Chicken

	C_2	D_2
C_1	6,6	2,7
D_1	7,2	0,0

○ Pure NE

Mixed NE:
$p_1 = p_2 = 1/3$
E[payoff]: 4.6

Third party or shared information necessary!

- Third party draws one of (C,C [illegible] C), (C,D) at random
- Inform players about thei[illegible]gned strategy (but not entire card)
- Player received D: d[illegible]OT deviate supposing other player plays card strat[illegible], gives maximal payoff = 7
- Player recei[illegible] other player plays C or D with probability ½, expec[illegible]ility larger for C, so does NOT deviate, E[payoff]=4

=> utility (6+6)+(2+7)+(7+2) = 30 => / 6 = 5

Correlated Equilibrium E[payoff]=5

Alternative Concepts: Correlated Equilibrium

- More general than NE
- Robert Aumann (1974)
- Each player acts according to observation of same public signal

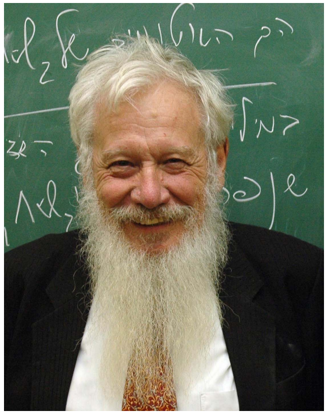

Internet Equilibria (1)

- Concrete examples where Internet actors are selfish?

- Selfish participants:
  - browsers
  - routers
  - servers
  - etc.

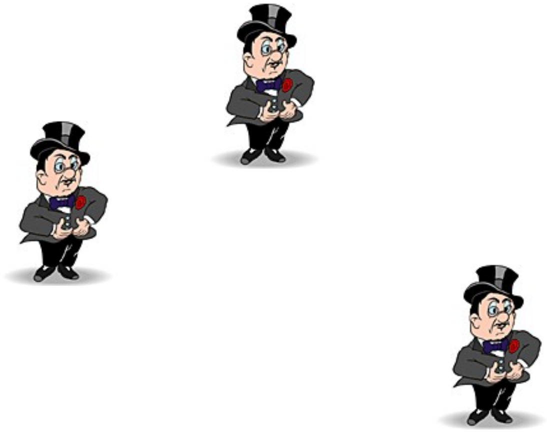

- Example ISPs: With which other providers should an ISP cooperate / „peer“?
  - Network creation games

Internet Equilibria (2)

- E.g., routing: The Internet is operated by thousands of autonomous systems which collaborate to deliver end-to-end flows (e.g., BGP protocol)
  - How is the payoff / income of this service distributed among the AS?
  - Routing and congestion games

- E.g.: Rate control algorithms
  - Why should a user throttle his transmission speed?
  - But TCP seems to work! Of which game is TCP the Nash equilibrium?

Price of Anarchy and Price of Stability

- Is strategic behavior harmful?
- Traditional economic theory by Adam Smith

“It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity but to their self-love, and never talk to them of our own necessities but of their advantages.”

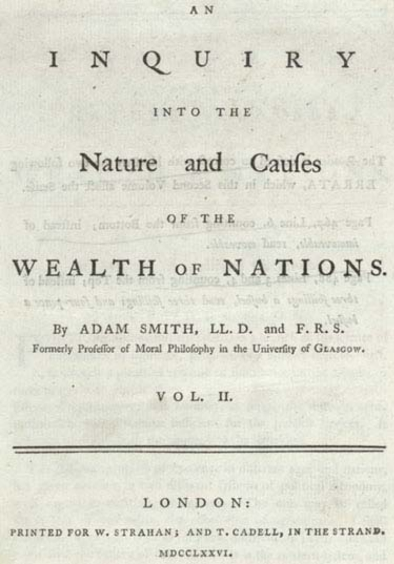

AN
INQUIRY
INTO THE
Nature and Cauſes
OF THE
WEALTH OF NATIONS.

By ADAM SMITH, LL. D. and F. R. S.
Formerly Profeſſor of Moral Philoſophy in the Univerſity of GLASGOW.

VOL. II.

LONDON:
PRINTED FOR W. STRAHAN; AND T. CADELL, IN THE STRAND.
MDCCLXXVI.

- However, it has been shown that Nash equilibria are often not equivalent to the social optimum. Rather, there are situations where selfishness comes at a certain cost.

Price of Anarchy and Price of Stability

- Impact on socio-economic systems such as the Internet?
- The Price of Anarchy: ratio of the total (social) cost of the worst Nash equilibrium divided by the socially optimal cost
- The Price of Stability: ratio of the total (social) cost of the best Nash equilibrium divided by the socially optimal cost

„If the competitive analysis reveals the price of not knowing the future, and approximability captures the price of not having exponential resources, this analysis seeks the price of uncoordinated individual utility-maximizing decisions.“

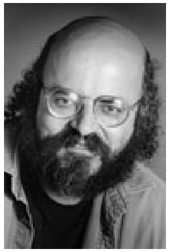

This Talk

- What is the Price of Anarchy in today‘s Internet?
  => Game theoretical analyses

- Going beyond selfish players: impact of malicious and social players

- What can be done about it?
  => Theory of mechanism design

- Limitations of game theory?

Example: Congestion Sensitive Load Balancing

- Delay as a function of load

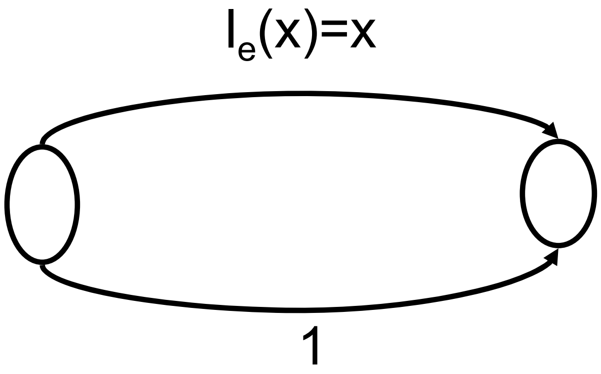

- What is the Price of Anarchy?

Example: Congestion Sensitive Load Balancing

- Total traffic 1, infinitely many players
- Nash equilibrium:

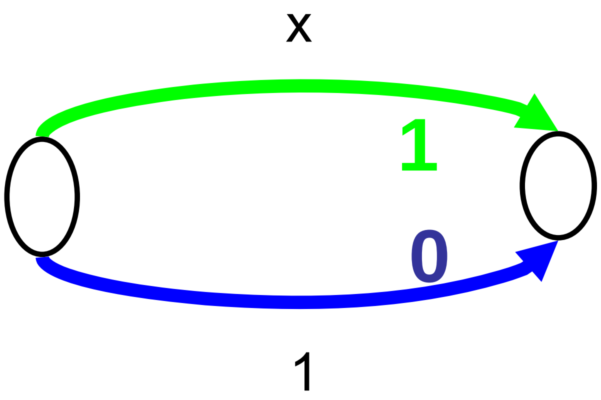

- Because players would change if

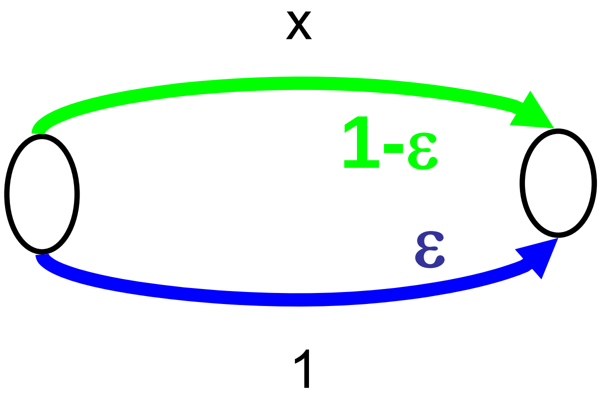

Example: Congestion Sensitive Load Balancing

- Nash equilibrium:
  latency 1 per player

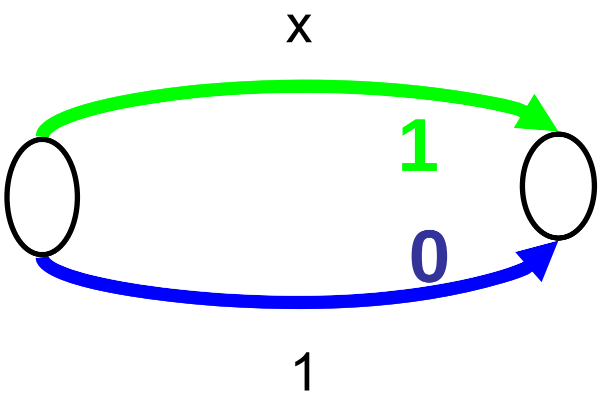

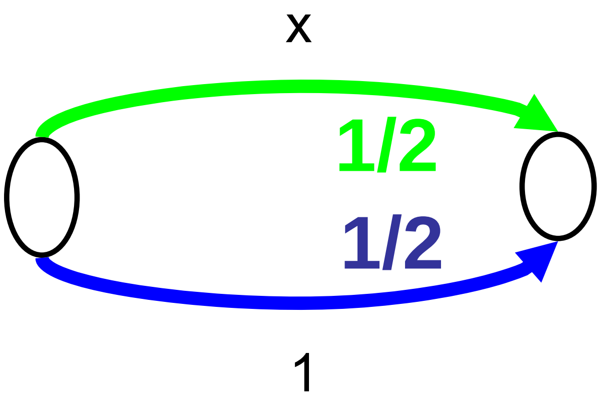

- Social optimum:
  half of the players have latency 1,
  half of the players have latency 1/2

- Price of Anarchy: 1/(3/4) = 4/3

Example: Congestion Sensitive Load Balancing

This Price of Anarchy is an upper bound for all networks with all kind of linear functions. ***[Roughgarden & Tardos, 2000]***

Good to Know: Breass‘s Paradox

- Game theoretical analyses can reveal interesting phenomena

- Travel times can increase with additional links!
  - Assume 4000 users
  - Edge delays as in the picture

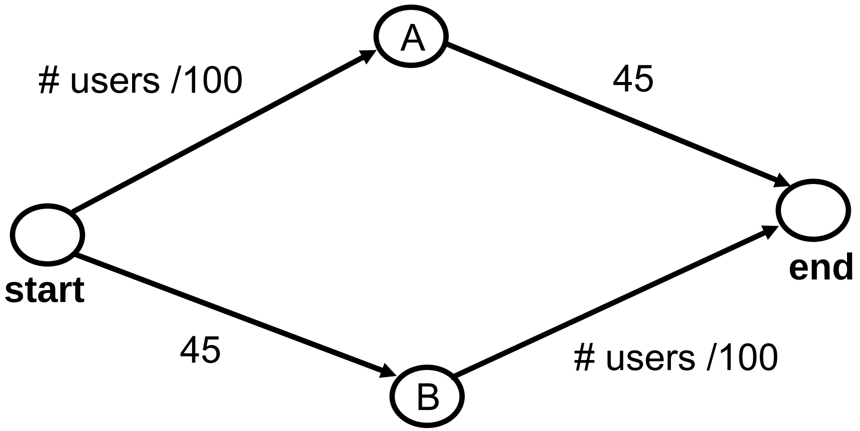

- Nash equilibrium: Half of the drivers drive via A, half via B!
  - Travel time: 2000/100 + 45 = <u>65 minutes</u>

Good to Know: Breass‘s Paradox

- Introduce a free road:

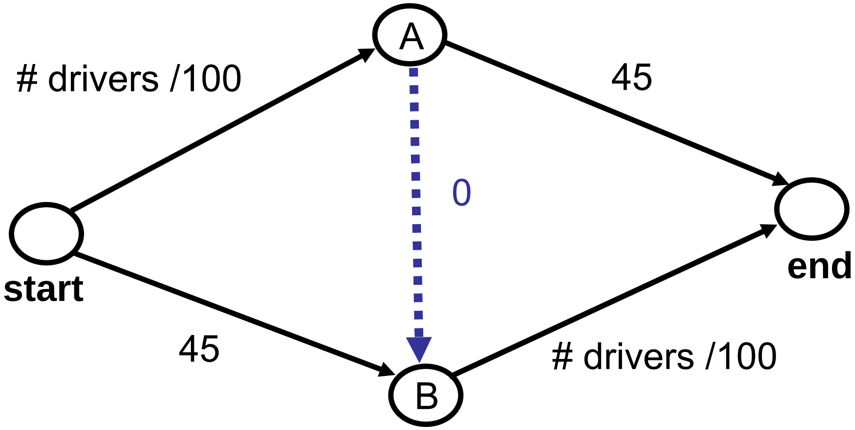

- Nash equilibrium: All drivers will drive
  start -> A -> B -> end
  - Travel time: 4000/100 + 0 + 4000/100 = <u>80 minutes</u> (> 65 minutes!)

On the Price of Anarchy of Unstructured Peer-to-Peer Topologies

Model – The “Locality Game”

- *n* peers {π_0, …, π_{n-1}} distributed in a metric space
  - defines distances (→ latencies) between peers
  - triangle inequality holds
  - examples: Euclidean space, doubling or growth-bounded metrics, 1D line,…
- Each peer can choose to which other peer(s) it connects
- Yields a directed graph…

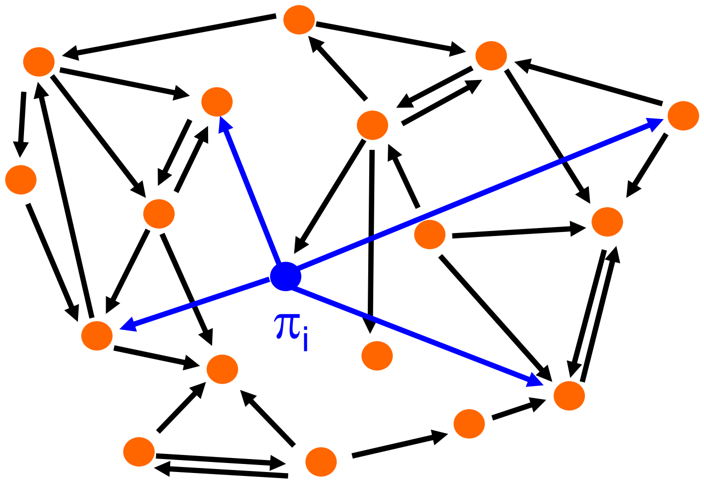

Model – The “Locality Game”

- Goal of a selfish peer:

  - Only little memory used
  - Small maintenance overhead

(1) Maintain a small number of neighbors only (out-degree)

(2) Small stretches to all other peers in the system

Fast lookups!

– Shortest path using links in G…

– … divided by shortest direct distance

LOCALITY!

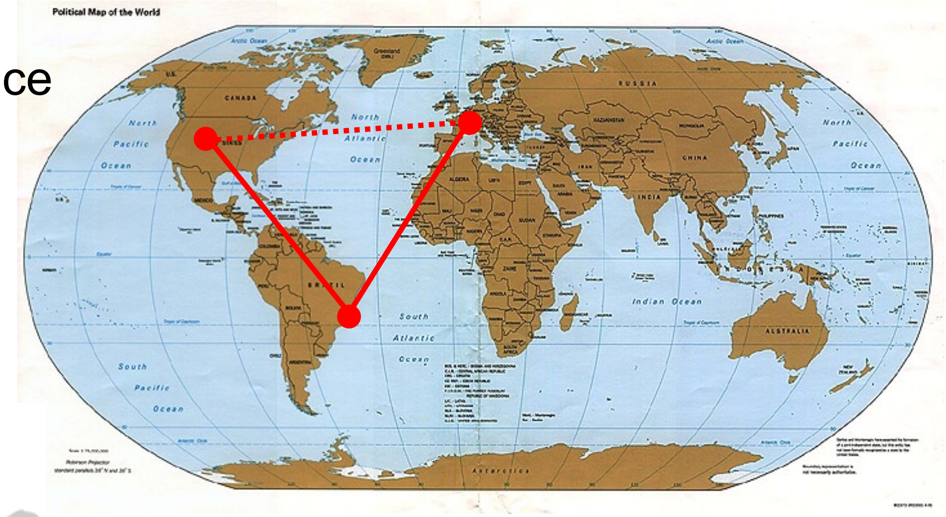

Classic P2P trade-off!

Model – The "Locality Game"

- Cost of a peer π_i:
  - Number of neighbors (out-degree) times a parameter α
  - plus stretches to all other peers
  - α captures the trade-off between link and stretch cost

$$cost_i = \alpha \cdot outdeg_i + \sum_{i \neq j} stretch_G(\pi_i, \pi_j)$$

- Goal of a peer: Minimize its cost!

- α is cost per link
- >0, otherwise solution is a complete graph

Locality Game: Upper Bound

$$cost_i = \alpha \cdot outdeg_i + \sum_{i \neq j} stretch_G(\pi_i, \pi_j)$$

- For connectivity, at least *n* links are necessary
  → OPT $\geq \alpha$ n
- Each peer has at least stretch 1 to all other peers
  → OPT $\geq$ n · (n-1) · 1 = $\Omega(n^2)$

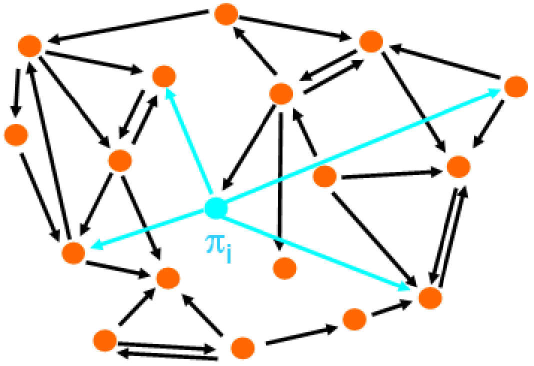

OPT $\in \Omega(\alpha$ n + n$^2)$

- Now: Upper Bound for NE? In any Nash equilibrium, no stretch exceeds α+1: total stretch cost at most O(α n^2)
  → otherwise it's worth connecting to the corresponding peer (stretch becomes 1, edge costs α)

- Total link cost also at most O(α n^2)

Really…?

NASH $\in$ O(αn^2)

Can be bad for large α

Price of Anarchy $\in$ O(min{α,n})

Locality Game: Price of Anarchy (Lower Bound)

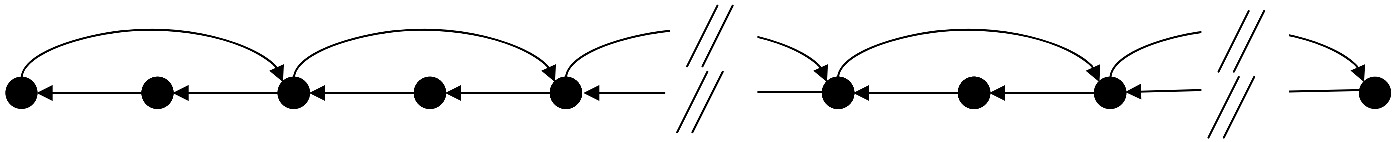

Peer:	π_1	π_2	π_3	π_4	π_5	…	π_{i-1}	π_i	π_{i+1}	…	π_n
Position:	$\frac{1}{2}$	α	$\frac{1}{2}\alpha^2$	α^3	$\frac{1}{2}\alpha^4$	…	$\frac{1}{2}\alpha^{i-2}$	α^{i-1}	$\frac{1}{2}\alpha^i$	…	$\frac{1}{2}\alpha^{n-1}$

To prove:

(1) “is a selfish topology” = instance forms a Nash equilibrium

(2) “has large costs compared to OPT”

= the social cost of this Nash equilibrium is $\Theta(\alpha n^2)$

Note: Social optimum is at most $O(\alpha n + n^2)$:

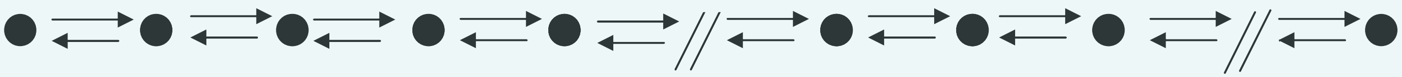

$O(n)$ links of cost α, and all stretches = 1

Lower Bound: Topology is Nash Equilibrium

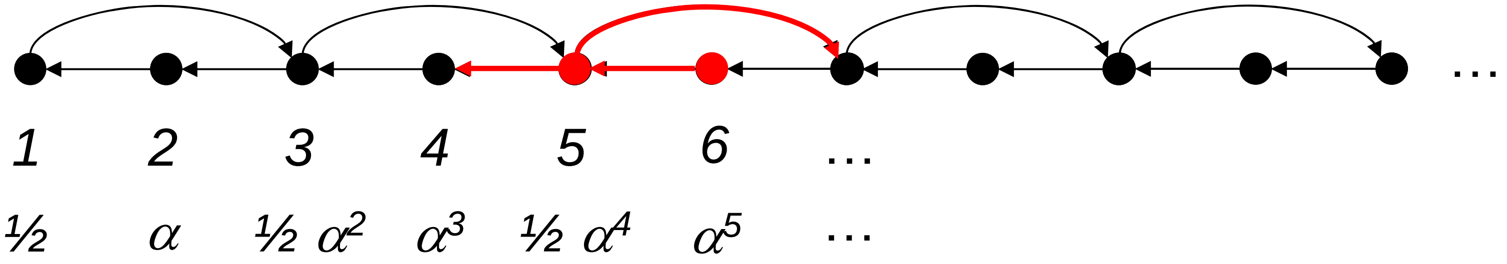

- Proof Sketch: Nash?
  - Even peers:
    - For connectivity, at least one link to a peer on the left is needed (cannot change neighbors without increasing costs!)
    - With this link, all peers on the left can be reached with an optimal stretch 1
    - Links to the right cannot reduce the stretch costs to other peers by more than α
  - Odd peers:
    - For connectivity, at least one link to a peer on the left is needed
    - With this link, all peers on the left can be reached with an optimal stretch 1
    - Moreover, it can be shown that all alternative or additional links to the right entail larger costs

Lower Bound: Topology has Large Costs

- Idea why social cost are $\Theta(\alpha\ n^2)$: $\Theta(n^2)$ stretches of size $\Theta(\alpha)$

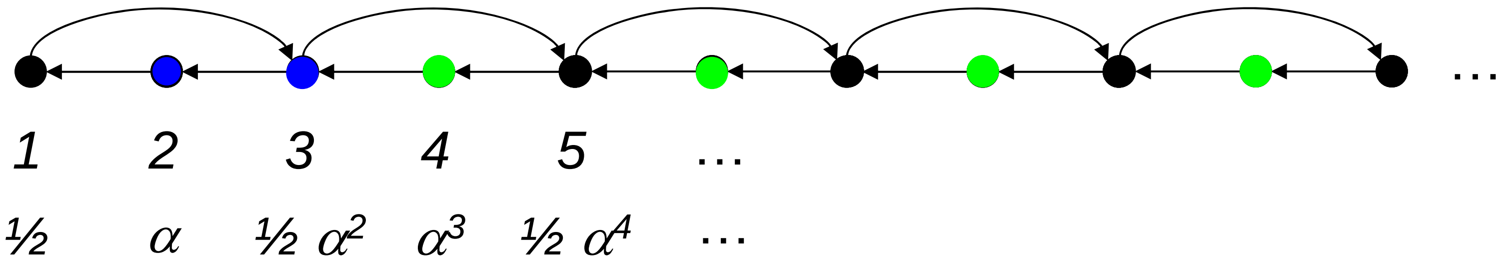

- The stretches from all odd peers *i* to a even peers *j>i* have stretch > $\alpha/2$

- And also the stretches between even peer *i* and even peer *j>i* *are >* $\alpha/2$

What about Stability…?

- Consider the following simple toy-example
- Let α=0.6 (for illustration only!)

- 5 peers in Euclidean plane as shown below (other distances implicit)
- What topology do they form…?

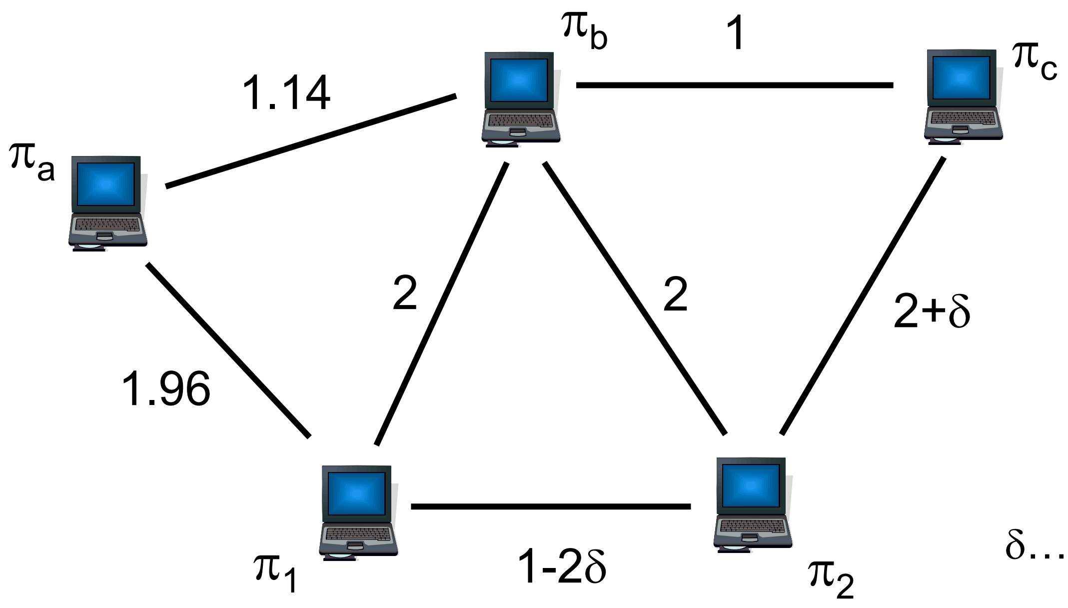

δ…arbitrary small number

What about Stability…?

- Example sequence:
  - Bidirectional links shown must exist in any NE, and peers at the bottom must have directed links to the upper peers somehow: considered now! (ignoring other links)

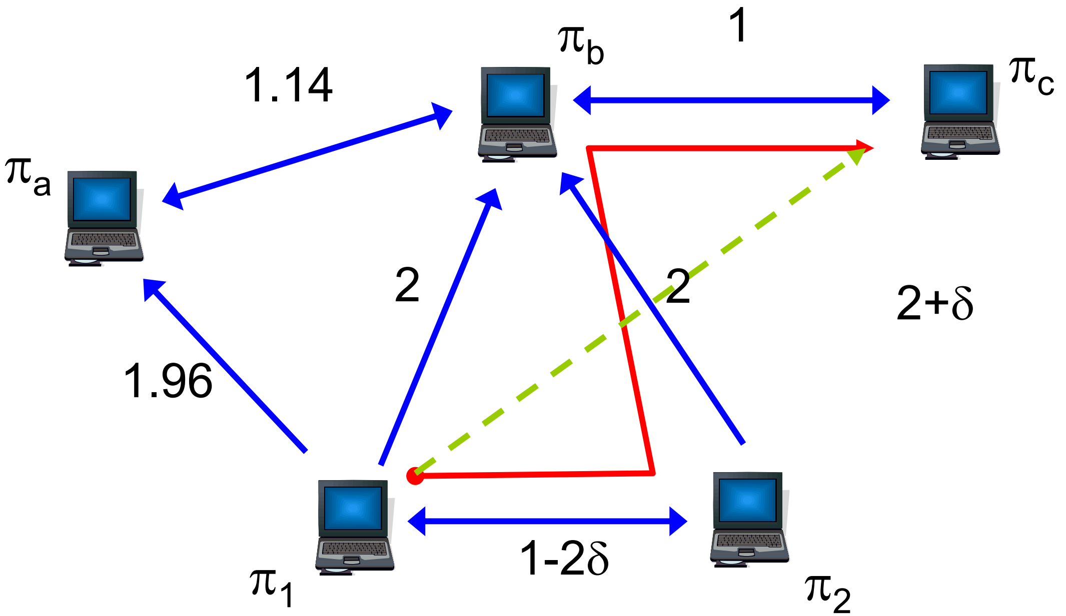

$$\underbrace{\frac{1-2\delta+2+1}{d(\pi_1,\pi_c)}}_{\text{stretch}(\pi_1,\pi_c)}+\underbrace{\frac{1-2\delta+2}{2}}_{\text{stretch}(\pi_1,\pi_b)}>\alpha+1+\underbrace{\frac{3}{d(\pi_1,\pi_c)}}_{\text{stretch}(\pi_1,\pi_c)}$$

What about Stability…?

- Example sequence:

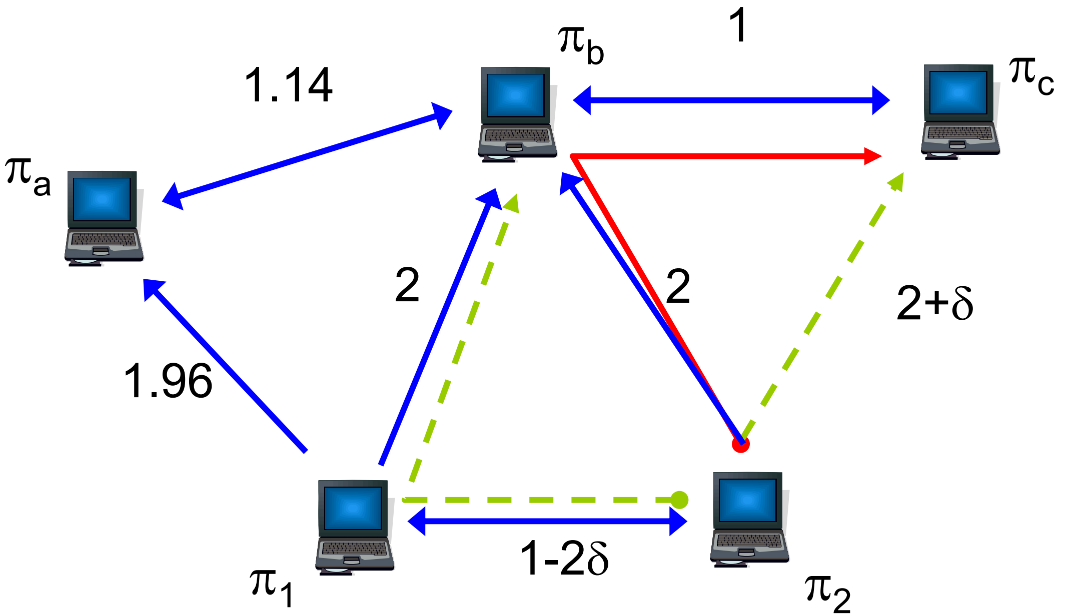

$$1 + \underbrace{\frac{1+2}{2+\delta}}_{\text{stretch}(\pi_2,\pi_c)} > \underbrace{\frac{1-2\delta+2}{2}}_{\text{stretch}(\pi_2,\pi_b)} + 1$$

What about Stability…?

- Example sequence:

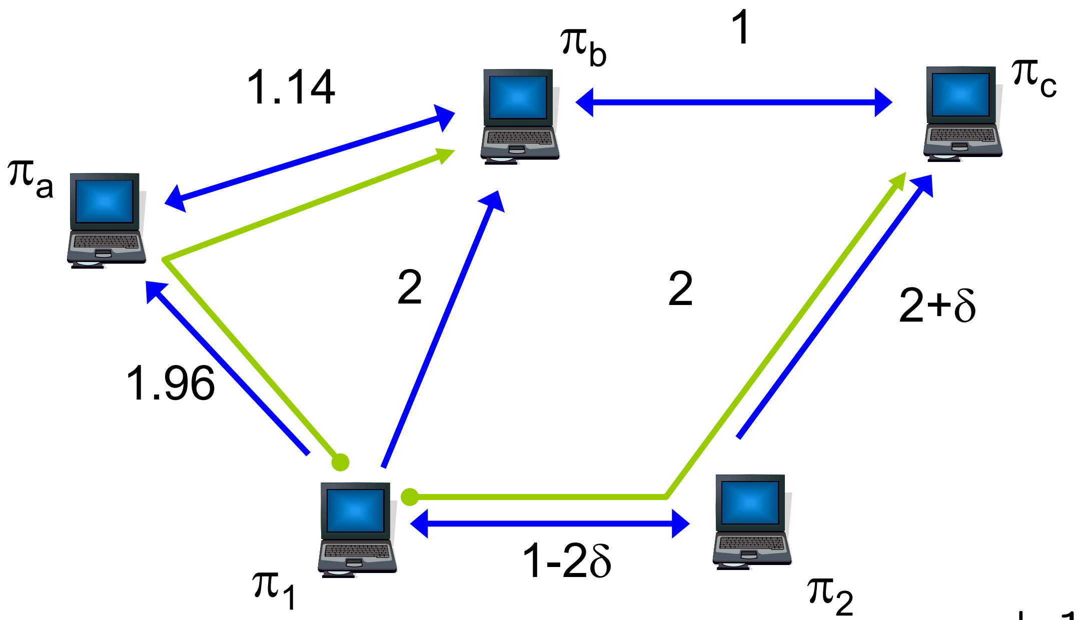

$$\alpha + 1 > \underbrace{\frac{1.96 + 1.14}{2}}_{\text{stretch}(\pi_1,\pi_b)}$$

What about Stability…?

- Example sequence:

Again initial situation
→ Changes repeat forever!

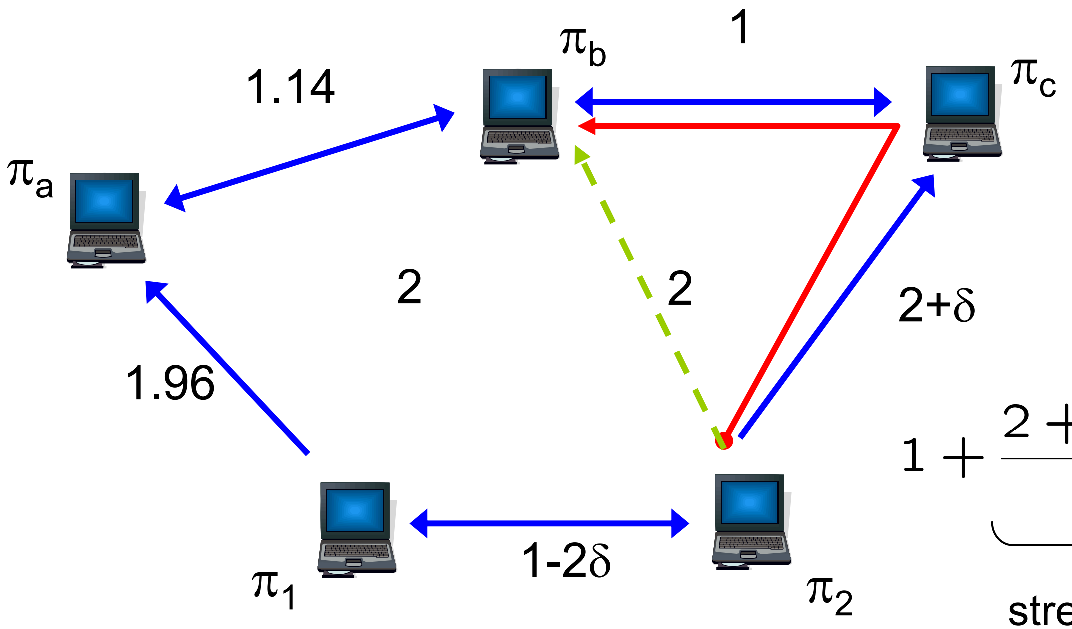

$$\underbrace{1+\frac{2+\delta+1}{2}}_{\text{stretch}(\pi_2,\pi_b)} > \underbrace{1+\frac{3}{2+\delta}}_{\text{stretch}(\pi_2,\pi_c)}$$

Generally, it can be shown that for all α , there are networks, that do not have a Nash equilibrium → that may not stabilize!

Locality Game: Stability for general α?

- So far, only a result for α=0.6
- With a trick, we can generalize it to all magnitudes of α
- Idea, replace one peer by a cluster of peers
- Each cluster has k peers → The network is instable for α=0.6k
- Trick: between clusters, at most one link is formed (larger α -> larger groups); this link then changes continuously as in the case of k=1.

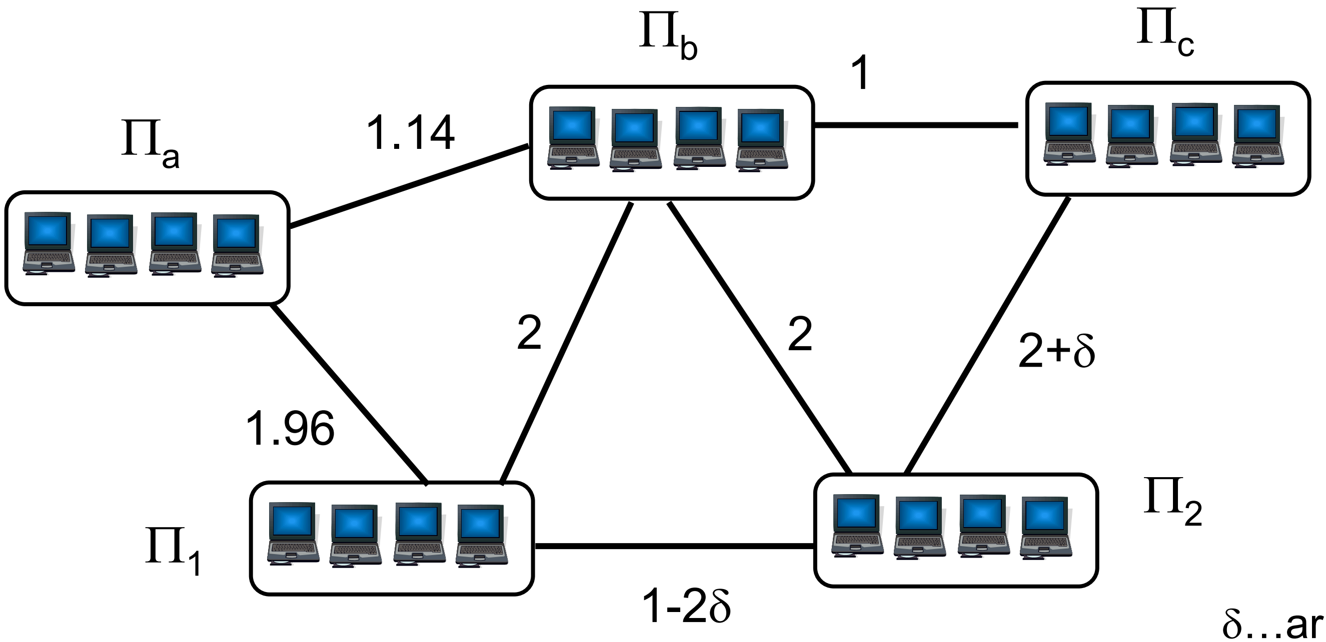

δ…arbitrary small number

Locality Game: Complexity of Nash Equilibrium

Proof Idea: reduction from SAT in CNF form (each clause has 2 or 3 literals)

- Polynomial time reduction: SAT formula -> distribution of nodes in metric space
- If each clause is satisfiable -> there exists a Nash equilibrium
- Otherwise, it does not.
- As reduction is fast, determining the complexity must also be NP-hard, like SAT!
- Remark: Special SAT, each variable in at most 3 clauses, still NP hard.

Locality Game: Complexity of Nash Equilibrium

- Arrange nodes as below
  - For each clause, our old instable network! (cliques -> for all magnitudes of α!)
  - Distances not shown are given by shortest path metric
  - Not Euclidean metric anymore, but triangle inequality etc. ok!
  - Two clusters at bottom, three clusters per clause, plus a cluster for each literal (positive and negative variable)
  - Clause cluster node on the right has short distance to those literal clusters appearing in the clause!

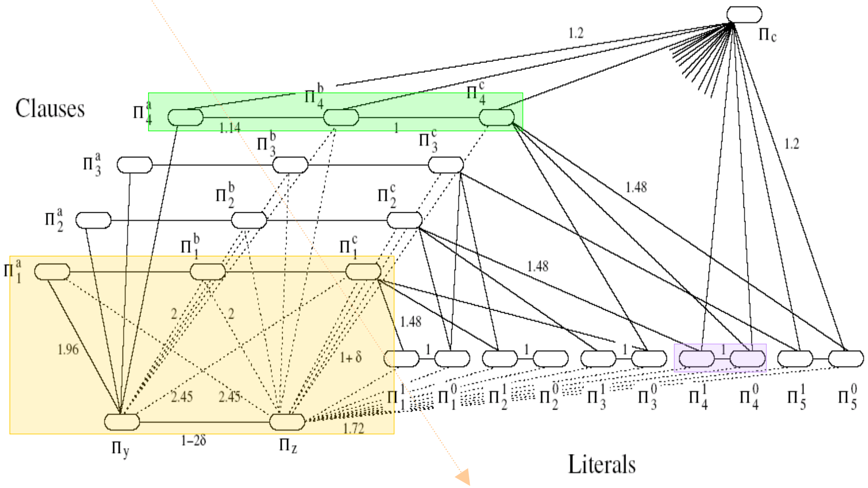

Game Theory with Malicious Players

Some Definitions from Game Theory

- Goal of a selfish player: minimize her own cost
- Social Cost is the sum of costs of selfish players
- Social Optimum (OPT)
  - Minimal social cost of a given problem instance
  - “solution formed by collaborating players”!
- Nash equilibrium
  - “Result” of selfish behavior
  - State in which no selfish player can reduce its costs by changing her strategy, given the strategies of the other players
- Measure impact of selfishness: Price of Anarchy
  - Captures the impact of selfishness by comparison with optimal solution
  - Formally: social costs of worst Nash equilibrium divided by optimal social cost

$$PoA := \frac{\text{worst Nash equilibrium}}{\text{social optimum}}$$

Large PoA ->
Selfish player are harmful!

“Byzantine* Game Theory”

- Game framework for malicious players
- Consider a system (network) with n players
- Among these players, s are selfish
- System contains b=n-s malicious players
- Malicious players want to *maximize* social cost!
- Define Byzantine Nash Equilibrium:

Social Cost:
Sum of costs of
selfish players:

$$Cost_{tot} = \sum_{i \in Selfish} cost_i(a)$$

A situation in which no selfish player can improve its perceived costs by changing its strategy!

Of course, whether a selfish player is happy with its situation depends on what she knows about the malicious players!

Do they know that there are malicious players? If yes, it will take this into account for computing its expected utility! Moreover, a player can react differently to knowledge (e.g. risk averse).

* „malicious“ is better… but we stick to paper notation in this talk.

Actual Costs vs. Perceived Costs

- Depending on selfish players' knowledge, actual costs (-> social costs) and perceived costs (-> Nash eq.) may differ!
- Actual Costs: $cost_i(a)$ ← Players do not know !

  → The cost of selfish player i in strategy profile a
- Perceived Costs: $\widehat{cost}_i(a)$ ← Byz. Nash Equilibrium

  → The cost that player i expects to have in strategy profile a, given preferences and his knowledge about malicious players!

← Many models conceivable →

Risk-averse…

Risk-seeking…

Neutral…

Nothing…,

Number of malicious players…

Distribution of malicious players…

Strategy of malicious players…

How to Measure the Impact of Malicious Players?

- Game theory with selfish players only studies the Price of Anarchy:

$$PoA := \frac{\text{worst Nash equilibrium}}{\text{social optimum}}$$

- We define Price of Byzantine Anarchy:

$$PoB(b) := \frac{\text{worst Byz. NE with } b \text{ malicious players}}{\text{social optimum}}$$

- Finally, we define the Price of Malice!

$$PoM(b) := \frac{\text{worst NE with } b \text{ malicious players}}{\text{worst NE}}$$

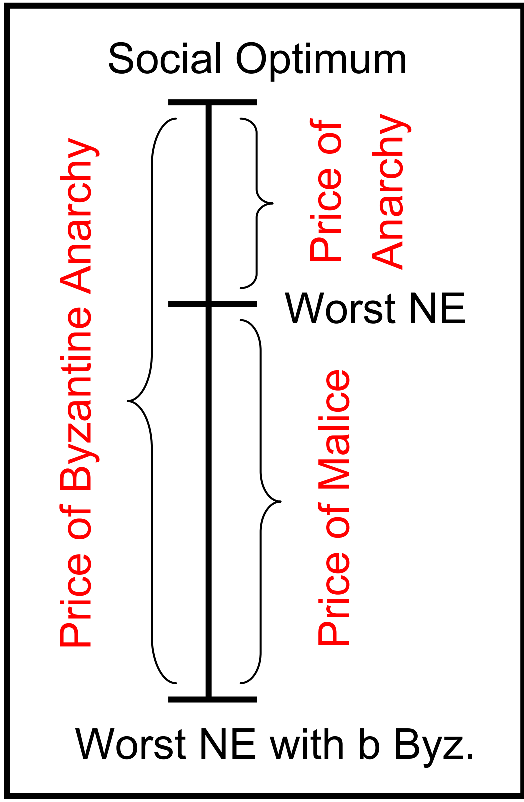

The Price of Malice captures the degradation of a system consisting of selfish agents due to malicious participants!

Remark on “Byzantine Game Theory”

- Are malicious players different from selfish players...? Also egoists?!

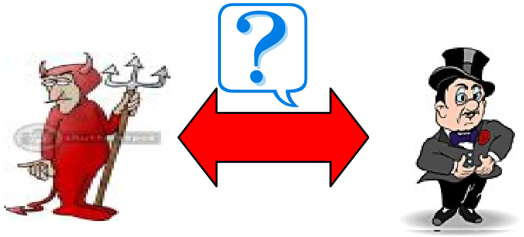

- Theoretically, malicious players are also selfish...
.... just with a different utility function!

Everyone is selfish!

→ Difference: Malicious players‘ utility function depends inversely on the total social welfare! („irrational“: utility depends on more than one player‘s utility)

→ When studying a specific game/scenario, it makes sense to distinguish between selfish and malicious players.

Sample Analysis: Virus Inoculation Game

- Given n nodes placed in a grid network
- Each peer or node can choose whether to install anti-virus software
- Nodes who install the software are secure (costs 1) ●
- Virus spreads from one randomly selected node in the network
- All nodes in the same insecure connected component are infected (being infected costs L, L>1)

→ Every node selfishly wants to minimize its expected cost!

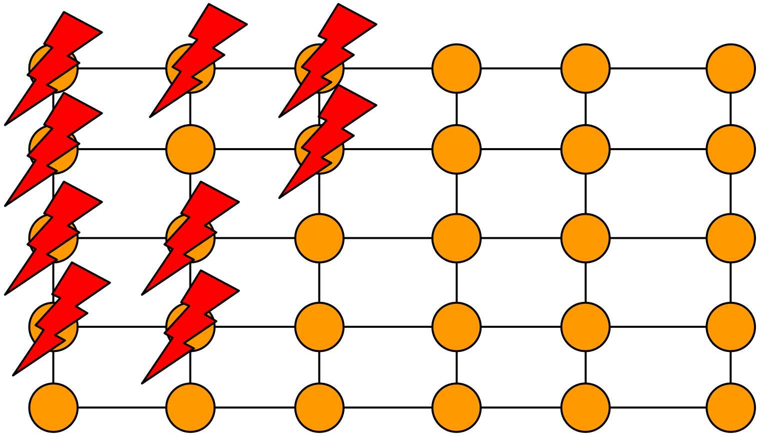

<u>Related Work:</u>

The VIG was first studied by Aspnes et al. [SODA'05]

- General Graphs
- No malicious players

Virus Game (1)

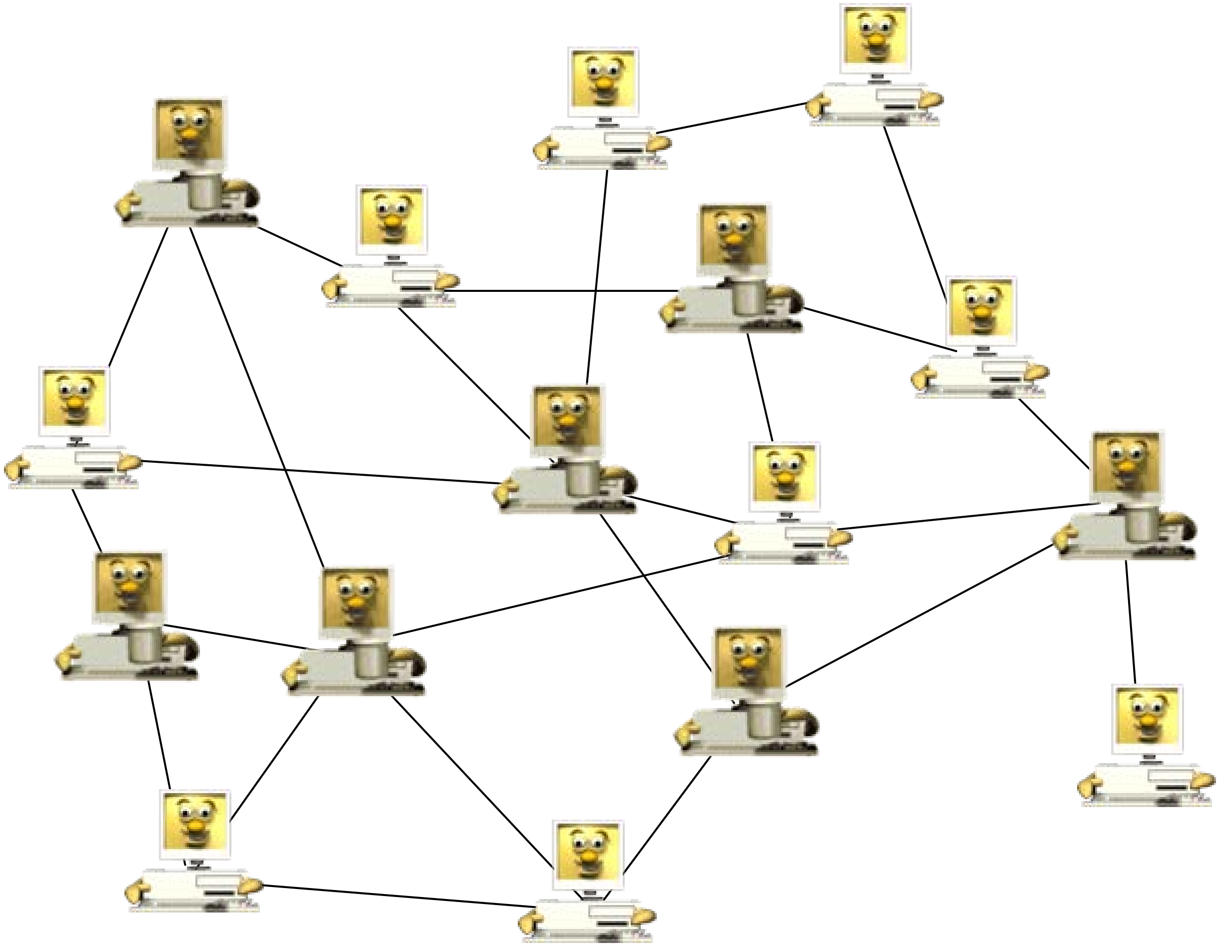

Virus Game (1)

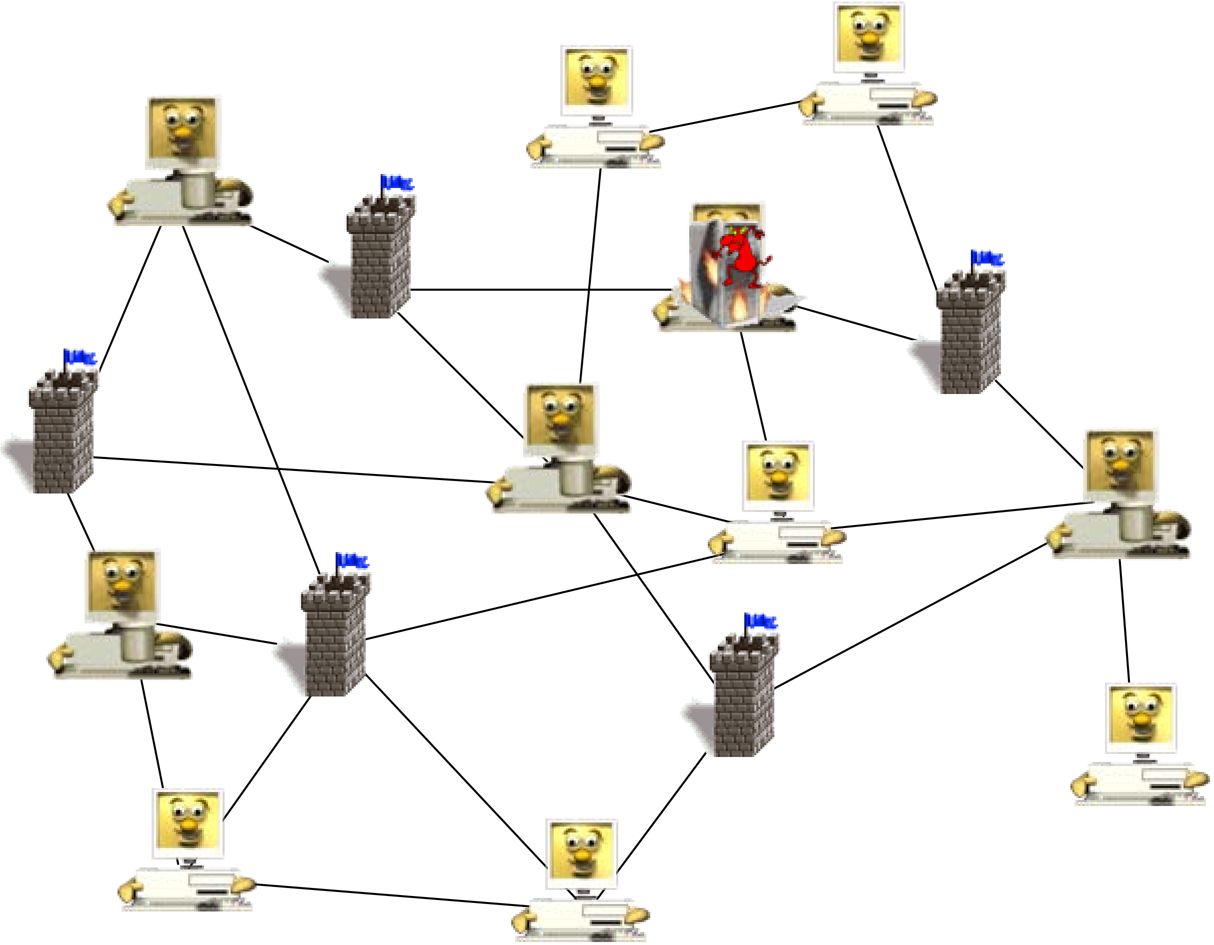

Virus Game (1)

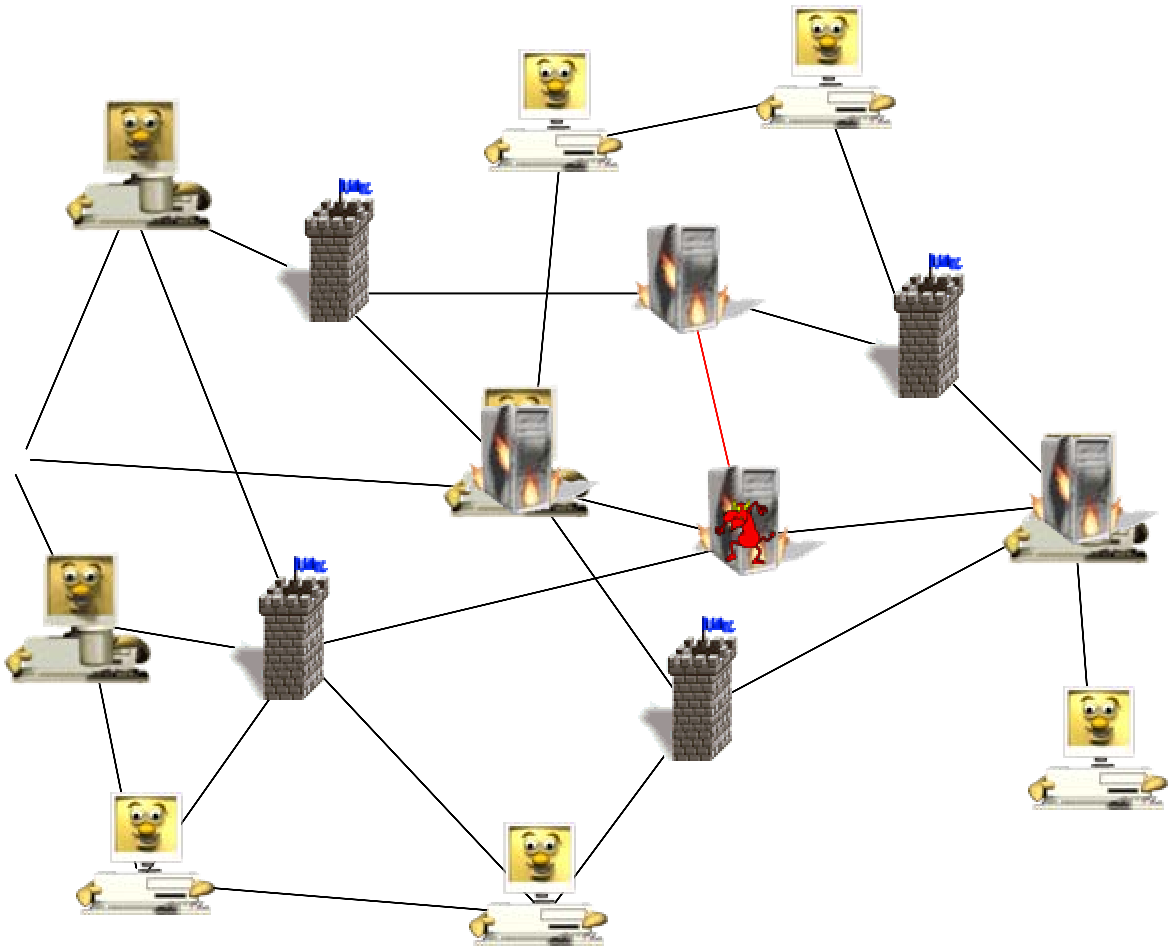

Virus Game (1)

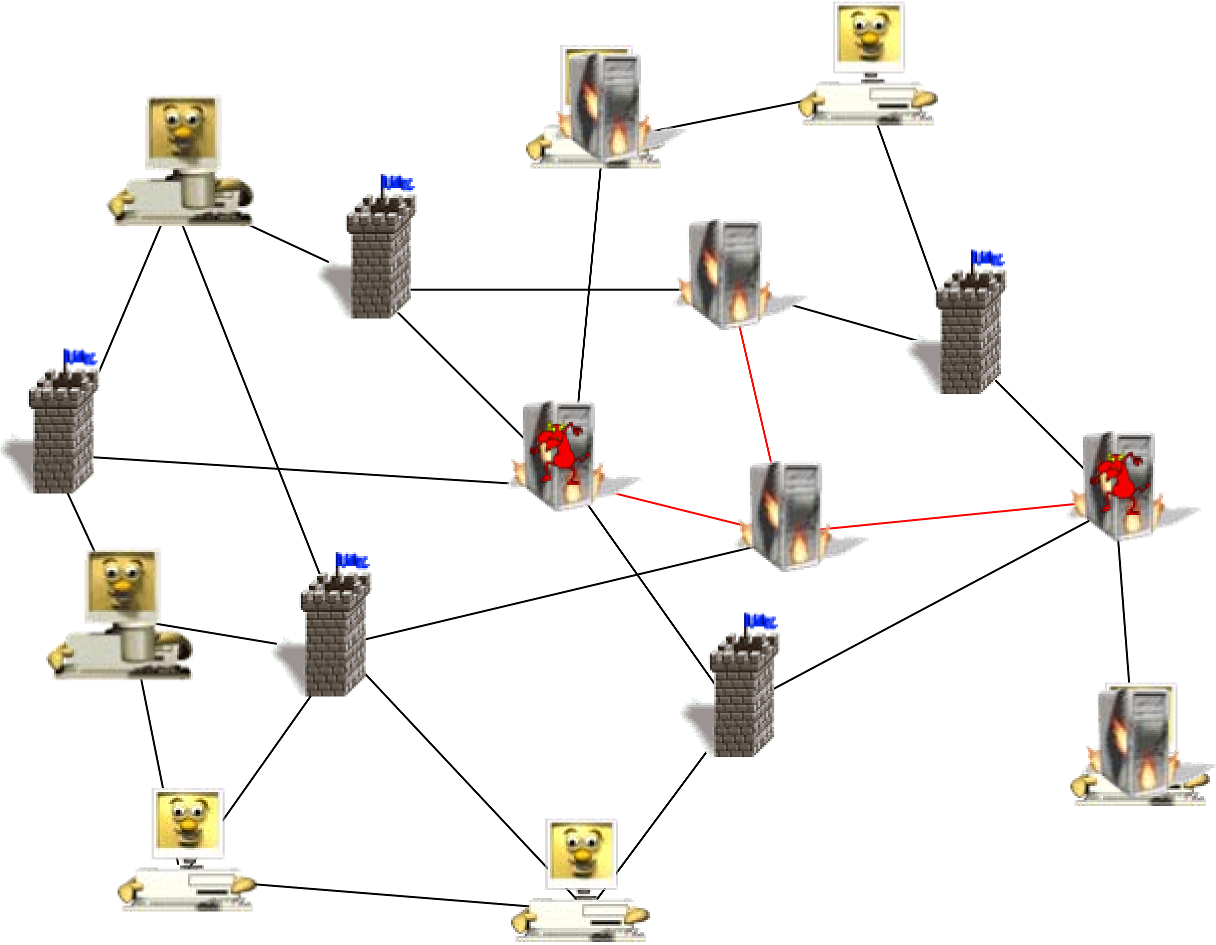

Virus Game (1)

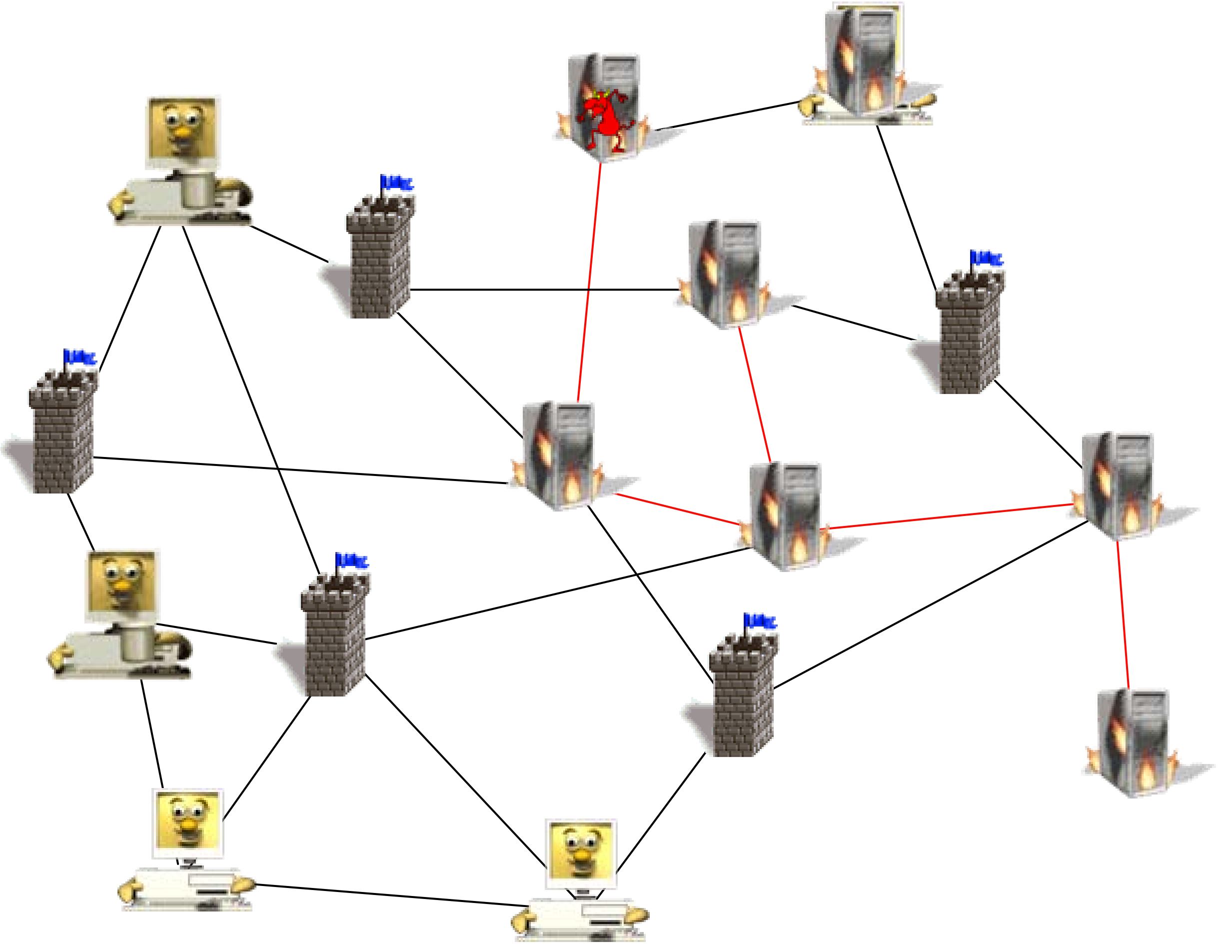

Virus Inoculation Game: Selfish Players Only

- What is the impact of selfishness in the virus inoculation game?
- What is the Price of Anarchy?
- Intuition:

Expected infection cost of

nodes in an insecure component A: quadratic in |A|

|A|/n * |A| * L = $|A|^2$ L/n

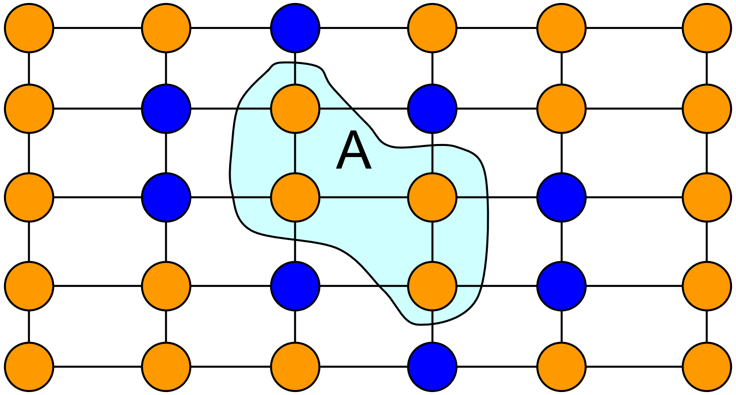

Total infection cost: $Cost_{inf} = \frac{L}{n} \sum_i k_i^2$

Total inoculation cost: $Cost_{inoc} = \gamma$

k_i: insecure nodes in the i-th component

γ: number of secure (inoculated) nodes

Optimal Social Cost

$$Cost_{OPT} = \Theta\left(n^{2/3} L^{1/3}\right)$$

Price of Anarchy:

$$PoA = \Theta\left(\sqrt[3]{\frac{n}{L}}\right)$$

Simple …
in NE, size <n/L+1
otherwise inoculate!

Adding Malicious Players…

- What is the impact of malicious agents in this selfish system?
- Let us add b malicious players to the grid!
- Every malicious player tries to maximize social cost!

→ Every malicious player pretends to inoculate, but does not!

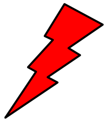

(worst-case: malicious player cannot be trusted and may say s.th. but do s.th. else…)

- What is the Price of Malice…?

→ Depends on what nodes *know* and how they *perceive threat*!

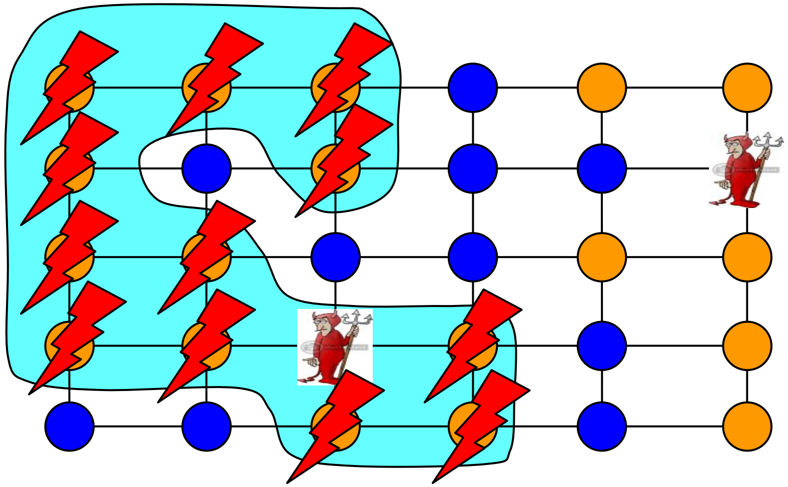

Distinguish between:

- Oblivious model
- Non-oblivious model
  - Risk-averse

Price of Malice – Oblivious case

- Nodes do not know about the existence of malicious agents (oblivious model)!
- They assume everyone is selfish and rational
- How much can the social cost deteriorate…?

- Simple upper bound:

- At most every selfish node can inoculate itself → $Cost_{inoc} \leq s$

- Recall: total infection cost is given by (see earlier: component i is hit with probability k_i/n, and we count only costs of the l_i selfish nodes therein)

$$Cost_{inf} = \frac{L}{n} \sum_i k_i \cdot l_i$$

Size of attack component i (*including* Byz.)

#selfish nodes in component i

Price of Malice – Oblivious case

- Total infection cost is given by: $Cost_{inf} = \frac{L}{n} \sum_i k_i \cdot l_i$
- It can be shown: for all components without any malicious node → $Cost_{inf}^{\overline{Byz}} \in O(s)$ (similar to analysis of PoA!)

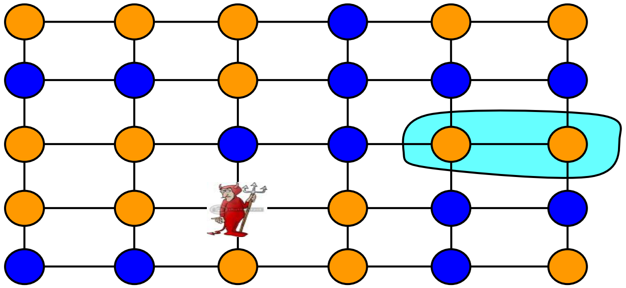

- On the other hand: a component i with $b_i>0$ malicious nodes: $\sum_i b_i = b$
- In any non-Byz NE, the size of an attack component is at most n/L, so

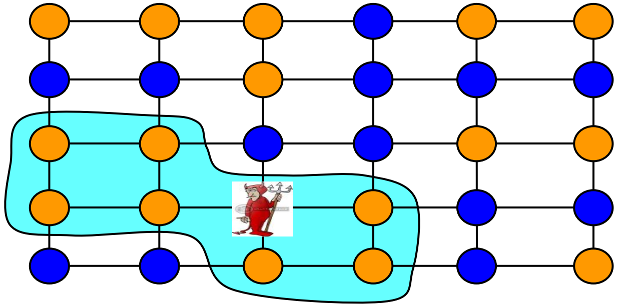

$$\begin{aligned} k_i &\leq (b_i+1) \cdot \frac{n}{L} + b_i \\ l_i &\leq (b_i+1) \cdot \frac{n}{L}. \end{aligned}$$

it can be shown → $Cost_{inf}^{Byz} \in O\left(\frac{b^2 n}{L}\right)$

Price of Malice – Oblivious case

- Adding inoculation and infection costs gives an upper bound on social costs:

$$O\left(s+\frac{b^2 n}{L}\right)$$

- Hence, the Price of Byzantine Anarchy is at most

for b<L/2

(for other case see paper)

$$PoB(b) \in \frac{O\left(s+\frac{b^2 n}{L}\right)}{\Theta(s^{2/3}L^{1/3})} \in O\left(\left(\frac{n}{L}\right)^{1/3}\cdot\left(1+\frac{b^2}{L}+\frac{b^3}{sL}\right)\right)$$

- The Price of Malice is at most

Because PoA is $\Theta\left(\left(\frac{n}{L}\right)^{1/3}\right)$

$$\boxed{PoM(b) \in O\left(1+\frac{b^2}{L}+\frac{b^3}{sL}\right)}$$

if L<n

Oblivious Case Lower Bound: Example Achieving It…

- In fact, these bounds are tight! I.e., there is instance with such high costs.

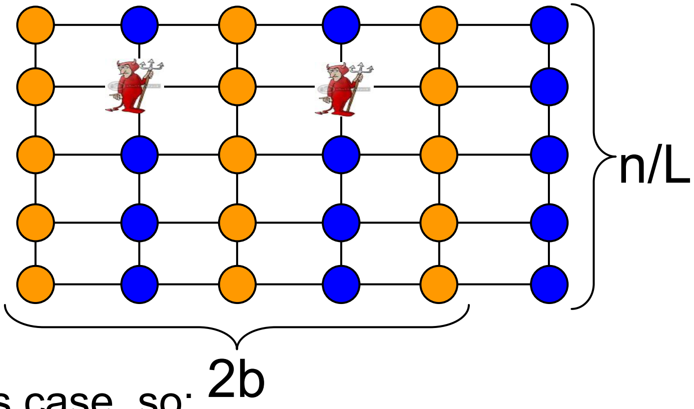

→ bad example: components with large surface (many inoculated nodes for given component size

=> bad NE! All malicious players together,

=> and one large attack component, large BNE)

→ this scenario where every second column is is fully inoculated is a Byz Nash Eq. in the oblivious case, so:

$$Cost_{inoc} = s/2 - b$$

→ What about infection costs? With prob. ((b+1)n/L+b)/n, infection starts at an insecure or a malicious node of an attack component of size (b+1)n/L

→ With prob. (n/2-(b+1)n/L)/n, a component of size n/L is hit

Combining all these costs yields $\Omega\left(s + \frac{b^2 n}{L}\right)$

Price of Malice – Oblivious case

- So, if nodes do not know about the existence of malicious agents!
- They assume everyone is selfish and rational
- Price of Byzantine Anarchy is:

This was Price of Anarchy…

$$PoB(b) = \Theta\left(\left(\frac{s}{L}\right)^{1/3} \cdot \left(1 + \frac{b^2}{L} + \frac{b^3}{sL}\right)\right)$$

- Price of Malice is:

$$PoM(b) = \Theta\left(1 + \frac{b^2}{L} + \frac{b^3}{sL}\right)$$

- Price of Malice grows more than linearly in b
- Price of Malice is always $\geq$ 1

This is clear, is it…?!

→ malicious players cannot improve social welfare!

Price of Malice – Non-oblivious Case

- Selfish nodes know the number of malicious agents b (non-oblivious)
- Assumption: they are risk-averse
- The situation can be totally different…
- …and more complicated!
- For intuition: consider the following scenario…: more nodes inoculated!

Each player wants to minimize its maximum possible cost (assuming worst case distribution)

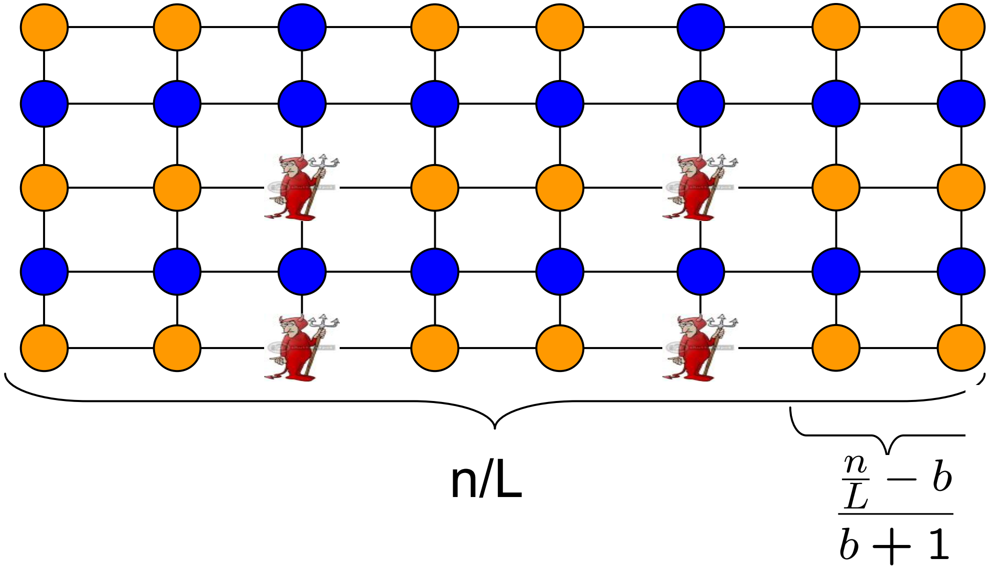

This constitutes a Byzantine Nash equilibrium!

Any b nodes can be removed while attack component size is at most n/L!

(n/L = size where selfish node is indifferent between inoculating or not in absence of malicious players)

Price of Malice – Lower Bound for Non-oblivious Case

- What is the social cost of this Byzantine Nash equilibrium…?

(all b malicious nodes in one row, every second column fully inoculated, attack size =< n/L)

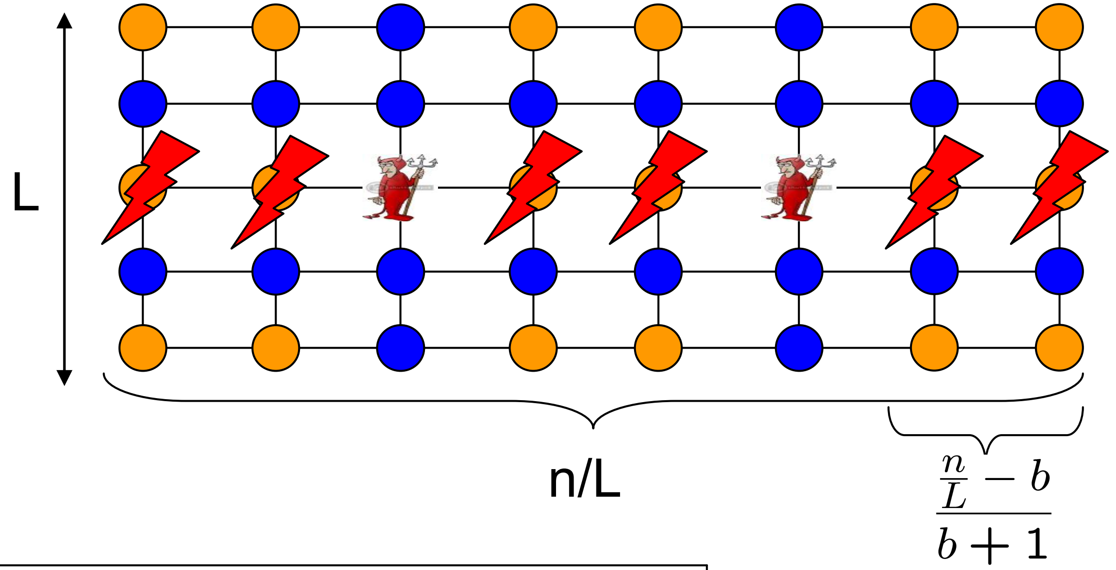

Infection cost of selfish nodes in infected row…

n/L-b selfish nodes

(b > n/L -> all s nodes inoculate)

It can be shown that expected infection cost for this row is:

$$Cost_{inf}^{0} = \frac{n}{L} - b$$

Infection cost of selfish nodes in other rows…

$$Cost_{inf}^{no} = \mu \cdot \frac{\frac{n}{L}-b}{b+1} \cdot \frac{L}{n}$$

number of insecure nodes in other rows

Total Cost:

$$Cost \geq \frac{s}{2} + \frac{bL}{4}$$

Total inoculation cost:

$$Cost_{inoc} = \frac{s}{2} + \frac{bL}{2} - b$$

Price of Malice – Non-oblivious Case: Lower Bound Results

- Nodes know the number of malicious agents b
- Assumption: Non-oblivious, risk-averse
- Price of Byzantine Anarchy is:

$$\boxed{PoB(b) \geq \frac{1}{8}\left(\left(\frac{n}{L}\right)^{1/3} + b\left(\frac{L}{n}\right)^{2/3}\right)}$$

- Price of Malice is:

$$\boxed{PoM(b) \geq \frac{\sqrt{\pi}}{48}\left(1 + \frac{bL}{n}\right)}$$

- Price of Malice grows at least linearly in b
- Price of Malice may become less than 1…!!!

→ Existence of malicious players can improve social welfare!

(malicious players cannot do better as we do not trust them in our model, i.e., not to inoculate still is the best thing for them to do!)

The Windfall of Malice: the “Fear Factor”

- In the non-oblivious case, the presence (or at least believe) of malicious players may improve social welfare!
- Selfish players are more willing to cooperate in the view of danger!
- Improved cooperation outweighs effect of malicious attack!
- In certain selfish systems:

Everybody is better off in case there are malicious players!

- Define the Fear-Factor Ψ

$$\Psi(b) := \frac{1}{PoM(b)}$$

$$\Downarrow$$

Ψ describes the achievable performance gain when introducing b Byzantine players to the system!

In virus game:

$$\Psi \leq \frac{48}{\sqrt{\pi}}$$

Price of Malice – Interpretations & Implications

- What is the implication in practical networking…?
- If Price of Anarchy is high

  → System designer must cope with selfishness (incentives, taxes)
- If Price of Malice is high

  → System must be protected against malicious behavior! (e.g., login, etc.)

orthogonal

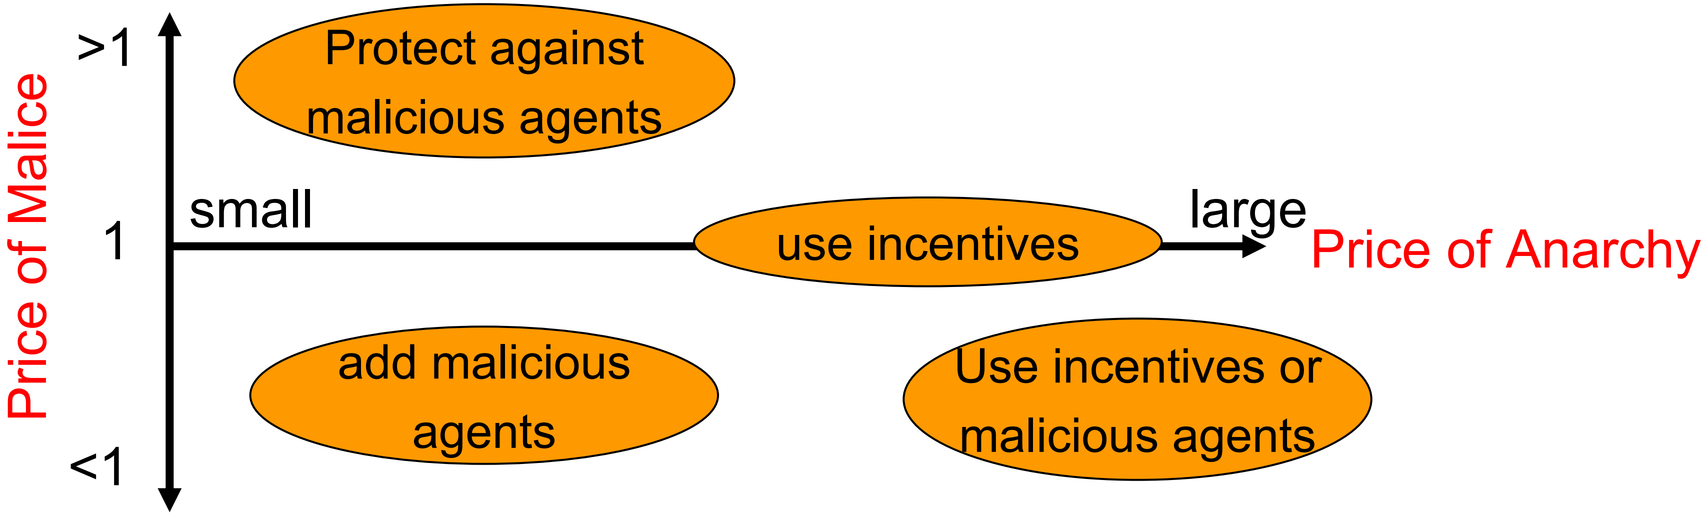

Reasoning about the Fear Factor

- What is the implication in practical networking…?
- Fear-Factor can improve network performance of selfish systems!

  (if Price of Malice < 1)

- Are there other selfish systems with Ψ>1 ?
- If yes… make use of malicious participants!!!
- Possible applications in P2P systems, multi-cast streaming, …

  → Increase cooperation by threatening malicious behavior!
- In our analysis: we theoretically upper bounded fear factor in virus game!

  → That is, fear-factor is fundamentally bounded by a constant

  (independent of b or n)

Game Theory with Social Players?

Impact of Social Players?

- In the following, we want to study social peers
- Motivation: Social networks
  - E.g., Skype contact lists

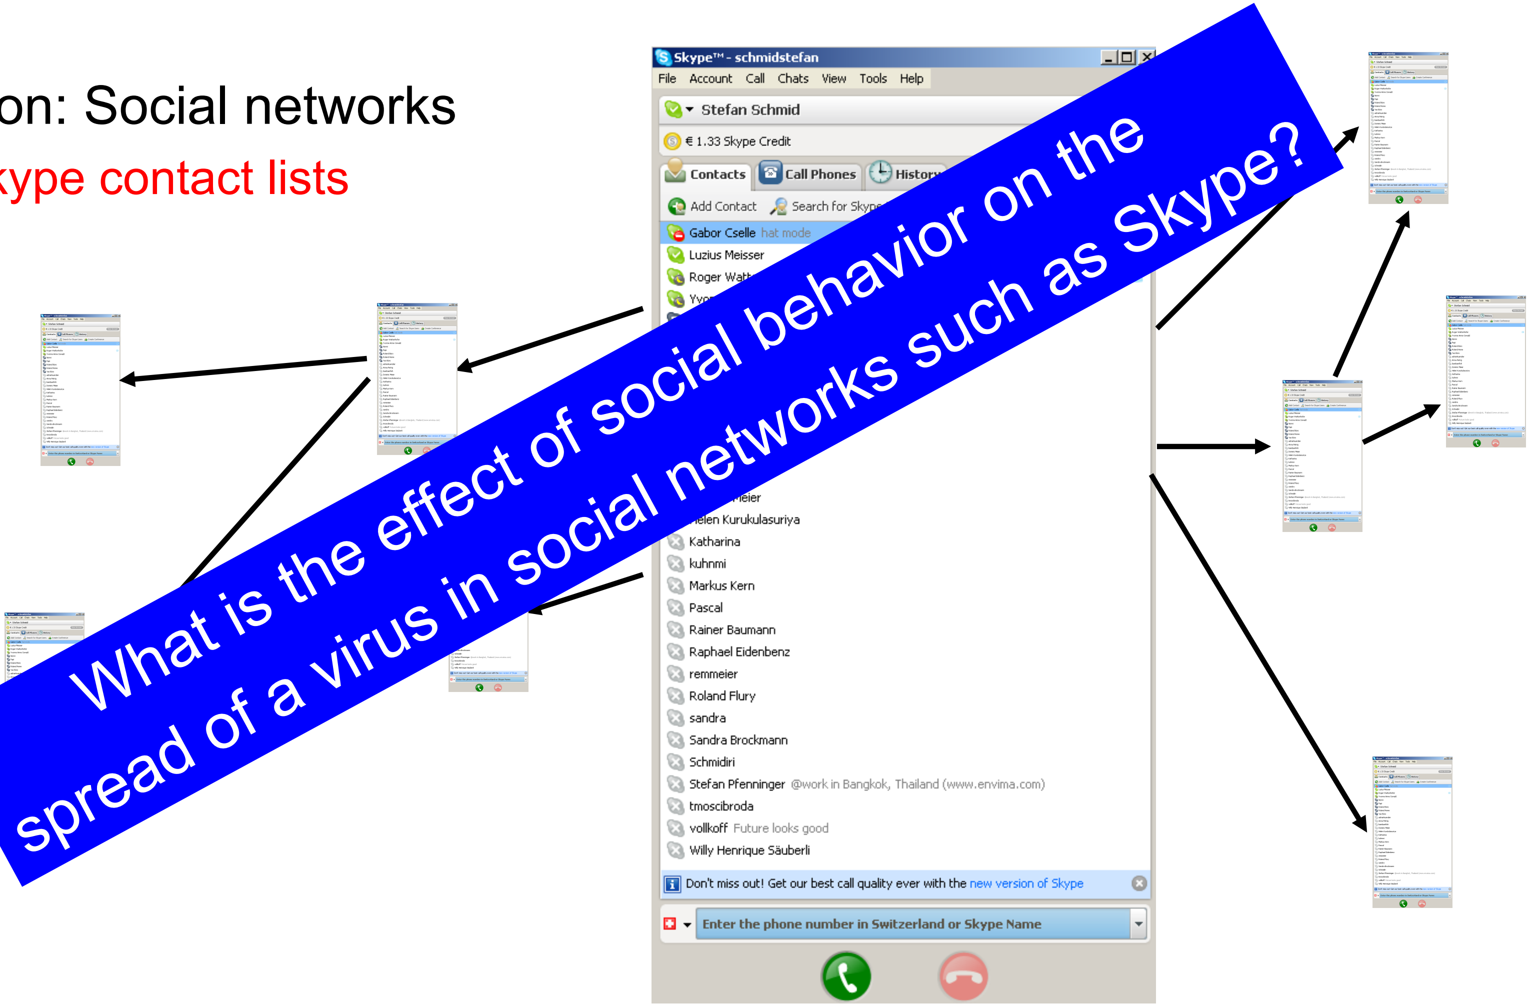

Gabor Cselle hat mode
Luzius Meisser
Roger Watt
Yv
eier
elen Kurukulasuriya
Katharina
kuhnmi
Markus Kern
Pascal
Rainer Baumann
Raphael Eidenbenz
remmeier
Roland Flury
sandra
Sandra Brockmann
Schmidiri
Stefan Pfenninger @work in Bangkok, Thailand (www.envima.com)
tmoscibroda
vollkoff Future looks good
Willy Henrique Säuberli

Don't miss out! Get our best call quality ever with the new version of Skype

Enter the phone number in Switzerland or Skype Name

What is the effect of social behavior on the spread of a virus in social networks such as Skype?

A Sample Game

- Sample game: virus inoculation

- The game
  - Network of n peers (or players)
  - Decide whether to inoculate or not
  - Inoculation costs C
  - If a peer is infected, it will cost L>C

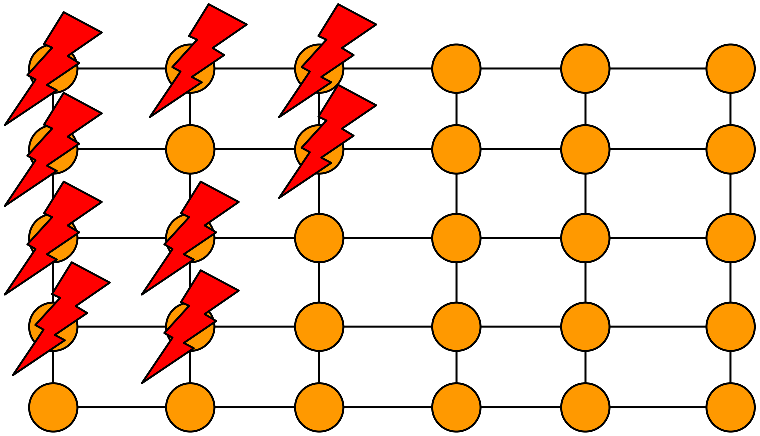

- At runtime: virus breaks out at a random player, and (recursively) infects all insecure adjacent players

Modelling Peers...

- Peers are selfish, maximize utility

- However, utility takes into account friends‘ utility
  - „local game theory“

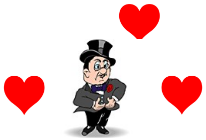

- Utility / cost function of a player
  - Actual individual cost:

k_i = attack component size

$$c_a(i,\vec{a}) = a_i \cdot C + (1-a_i)L \cdot \frac{k_i}{n}$$

a_i = inoculated?

  - Perceived individual cost:

$$c_p(i,\vec{a}) = c_a(i,\vec{a}) + F \cdot \sum_{p_j \in \Gamma(p_i)} c_a(j,\vec{a})$$

F = friendship factor,
extent to which players care about friends

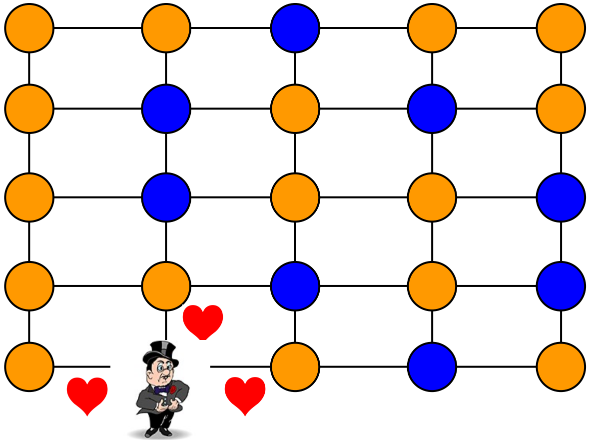

Social Costs and Equilibria

- In order to quantify effects of social behavior...

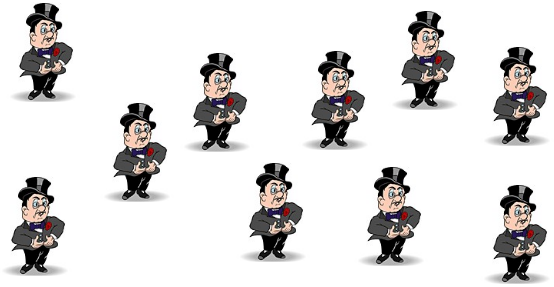

- Social costs
  - Sum over all players' actual costs

- Nash equilibria
  - Strategy profile where each player cannot improve her welfare...
  - ... given the strategies of the other players
  - Nash equilibrium (NE): scenario where all players are selfish
  - Friendship Nash equilibrium (FNE): social scenario
  - FNE defined with respect to perceived costs!
  - Typical assumption: selfish players end up in such an equilibrium (if it exists)

Evaluation

- What is the impact of social behavior?

- Windfall of friendship
  - Compare (social cost of) worst NE where every player is selfish (perceived costs = actual costs)...

  - ... to worst FNE where players take friends‘ actual costs into account with a factor F (players are „social“)

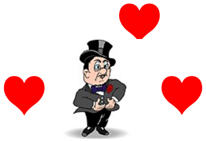

Windfall of Friendship

- Formally, the windfall of friendship (WoF) is defined as

$$\Upsilon(F, I) = \frac{\max_{NE} C_{NE}(I)}{\max_{FNE} C_{FNE}(F, I)}$$

instance I describes graph, C and L

- WoF >> 1 => system benefits from social aspect
  - Social welfare increased

- WoF < 1 => social aspect harmful
  - Social welfare reduced

Characterization of NE

- In regular (and pure) NE, it holds that...

- Insecure player is in attack component A of size at most Cn/L
  - otherwise, infection cost
    > (Cn/L)/n * L = C

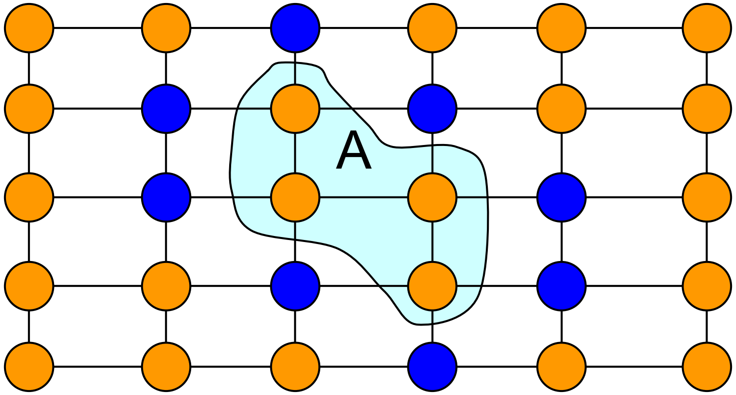

- Secure player: if she became insecure, she would be in attack component of size at least Cn/L
  - same argument: otherwise it‘s worthwhile to change strategies

Characterization of Friendship Nash Equilibria

- In friendship Nash equilibria, the situation is more complex
- E.g., problem is asymmetric
  - One insecure player in attack component may be happy...
  - ... while other player in same component is not
  - Reason: second player may have more insecure neighbors

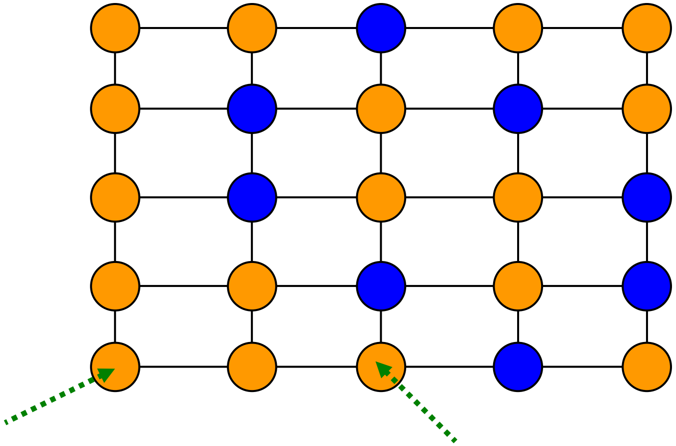

not happy, two insecure neighbors
(with same actual costs)

happy, only one insecure neighbor
(with same actual costs)

Bounds for the Windfall

THEOREM 4.2. *For all instances of the virus inoculation game and* $0 \leq F \leq 1$, *it holds that*

$$1 \leq \Upsilon(F, I) \leq PoA(I).$$

- It is always beneficial when players are social!

- The windfall can never be larger than the price of anarchy

  - Price of anarchy = ratio of worst Nash equilibrium cost divided by social optimum cost

- Actually, there are problem instances (with large F) which indeed have a windfall of this magnitude („tight bounds“, e.g., star network)

Example for Star Graph

- In regular NE, there is always a (worst) equilibrium where center is insecure, i.e., we have n/L insecure nodes and n-n/L secure nodes (for C=1):

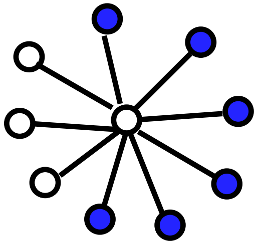

Social cost = (n/L)/n * n/L * L + (n-n/L) ~ n

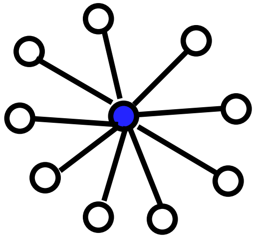

- In friendship Nash equilibrium, there are situations where center *must* inoculate, yielding optimal social costs of (for C=1):

Social cost = „social optimum“
= 1 + (n-1)/n * L ~ L

WoF as large as maximal price of anarchy in arbitrary graphs (i.e., n for constant L).

A Proof Idea for Lower Bound

- WoF $\geq$ 1 because...:

- Consider arbitrary FNE (for any F)

- From this FNE, we can construct (by a best response strategy) a regular NE with at least as large social costs
  - Component size can only increase: peers become insecure, but not secure
  - Due to symmetry, a player who joins the attack component (i.e., becomes insecure) will not trigger others to become secure
  - It is easy to see that this yields larger social costs

- In a sense, this result matches our intuitive expectations...

Monotonicity

But the windfall does not increase monotonously:
WoF can decline when players care more about their friends!

- Example again in simple star graph...

Monotonicity: Counterexample

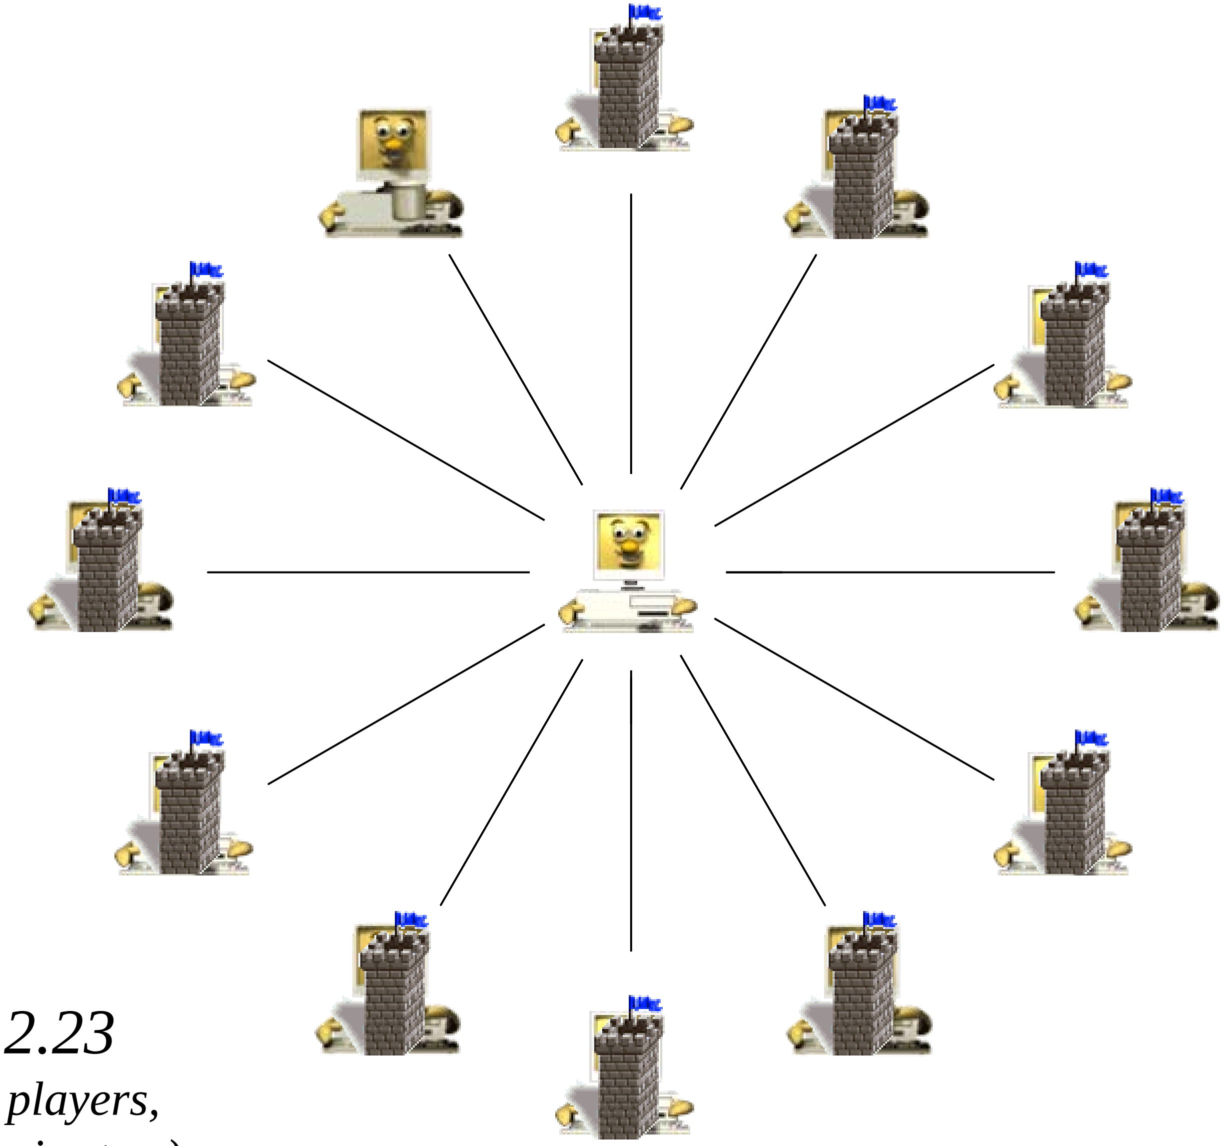

$n = 13$

$C = 1$

$L = 4$

$F = 0.9$

total cost = 12.23

(many inoculated players,
attack component size two)

Monotonicity: Counterexample

n = 13

C = 1

L = 4

F = 0.1

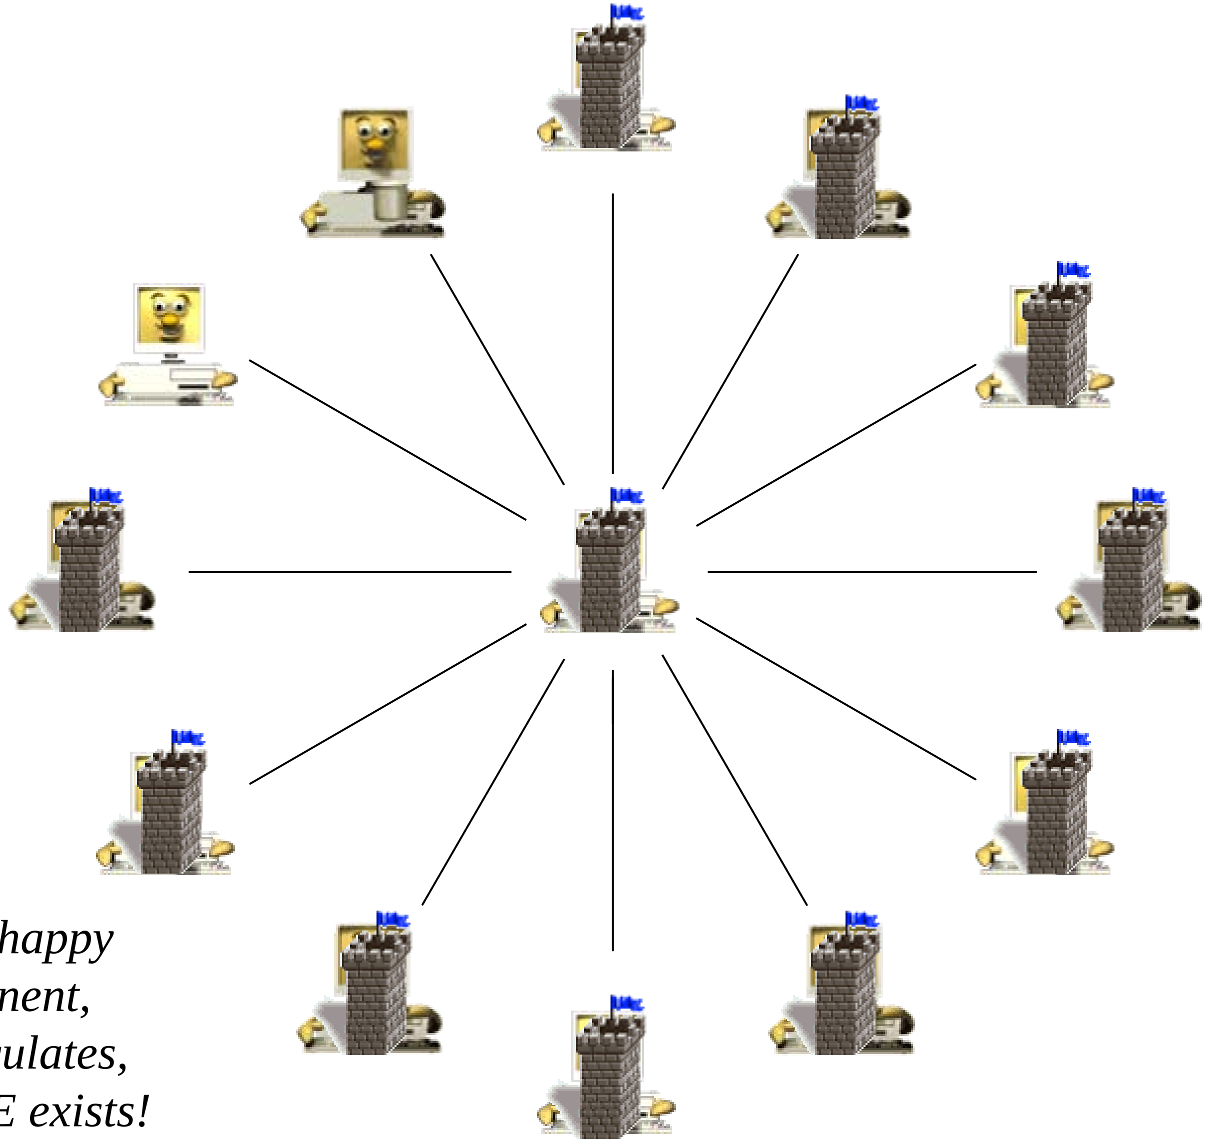

Boundary players happy with larger component, center always inoculates, thus: only this FNE exists! total cost = 4.69

Overcoming the Tragedy of the Commons:

Mechanism Design

Incentives in Peer-to-Peer Computing (1)

- Peer-to-peer systems rely on contributions

- Real peer-to-peer networks
  - Evidence of strategic behavior
  - Excellent live laboratory!

- Peer-to-peer interactions: strangers that will never meet again

- Plus: hidden actions, creation of multiple identities, etc.

Incentives in Peer-to-Peer Computing (2)

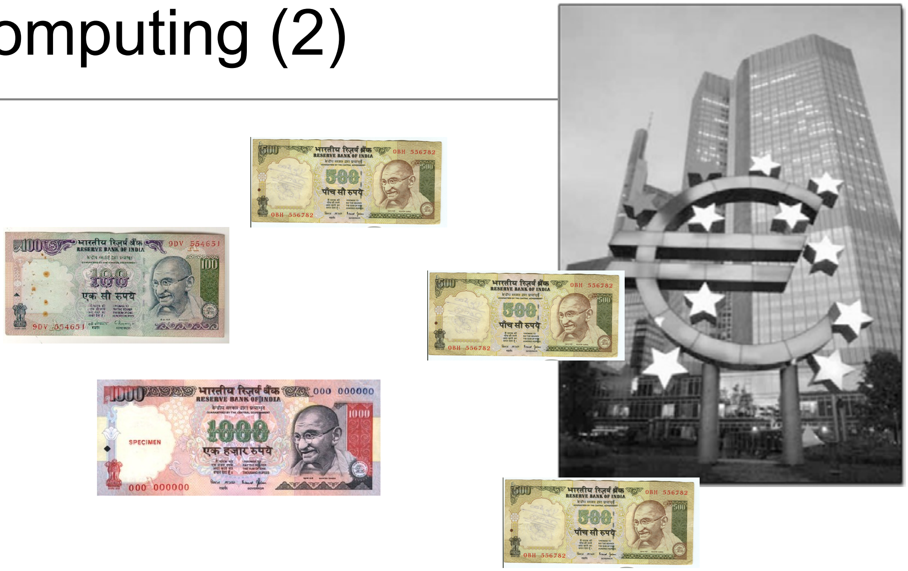

- Solutions?

- Naïv solution: Kazaa
  - client monitors contributions
  - Kazaa lite client hardwires to max

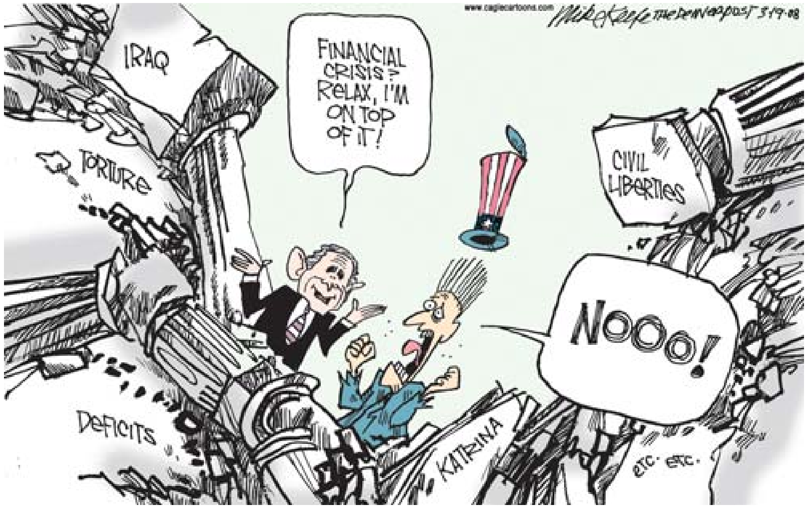

- Idea: do it like in the real economy!
  - virtual money systems
  - e.g., Karma (peer-to-peer currency)
  - complex issue... (fight inflation/deflation? bubbles?! integrity?)

Incentives in Peer-to-Peer Computing (3)

- Sometimes, barter systems are a good solution!

- Bram Cohen heralded a paradigm shift by showing that cooperation is possible on a single file!

- BitTorrent uses a tit-for-tat mechanism on file blocks
  - Bootstrap problem solved with optimistic unchoking

Incentives in Peer-to-Peer Computing (4)

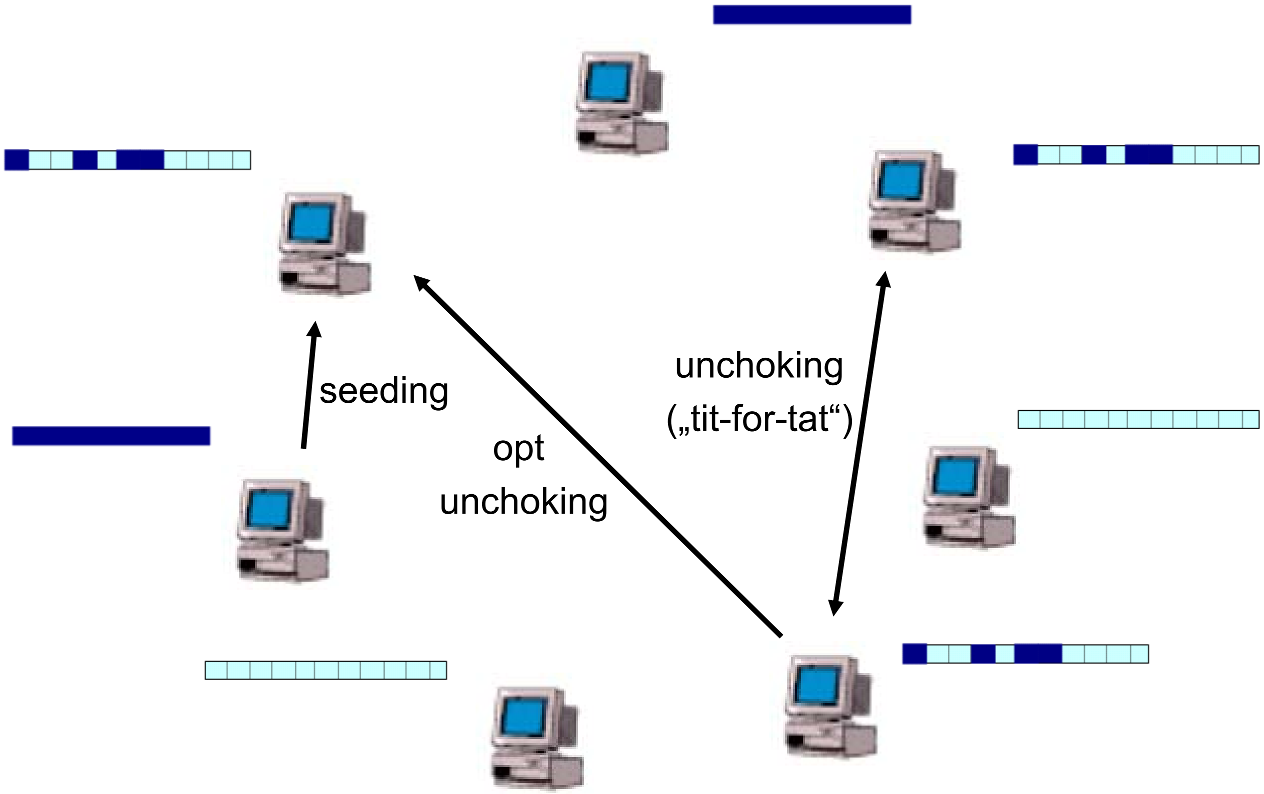

Mechanism Design

Design rules of a system leading to good results with selfish participants

- good results?
  fairness, happiness, social welfare: objective function
- selfishness?
  rationality, bounded rationality, maliciousness, altruism

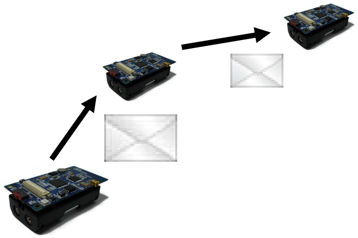

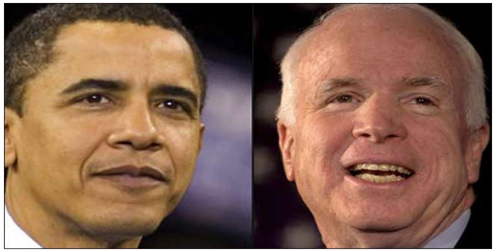

Mechanism Design

Design rules of a system leading to good results with selfish participants

- good results?
  fairness, happiness, social welfare, objective function
- selfishness?
  rationality, bounded rationality, maliciousness, altruism

Mechanism Design sometimes called „inverse“ Game Theory:

GT
game is given,
we have to figure out
how to act

MD
we know how we would
like the agents to act,
have to figure out game

Example: (Single-item) Auctions

Sealed-bid auction: every bidder submits bid in a sealed envelope

- First-price sealed-bid auction: highest bid wins, pays amount of own bid (e.g. *Dutch auction* where value for product is decreased iteratively, break ties randomly)
- Second-price sealed-bid auction: highest bid wins, pays amount of second-highest bid (e.g. *English auction* where price rises; sold as soon as second player quits)

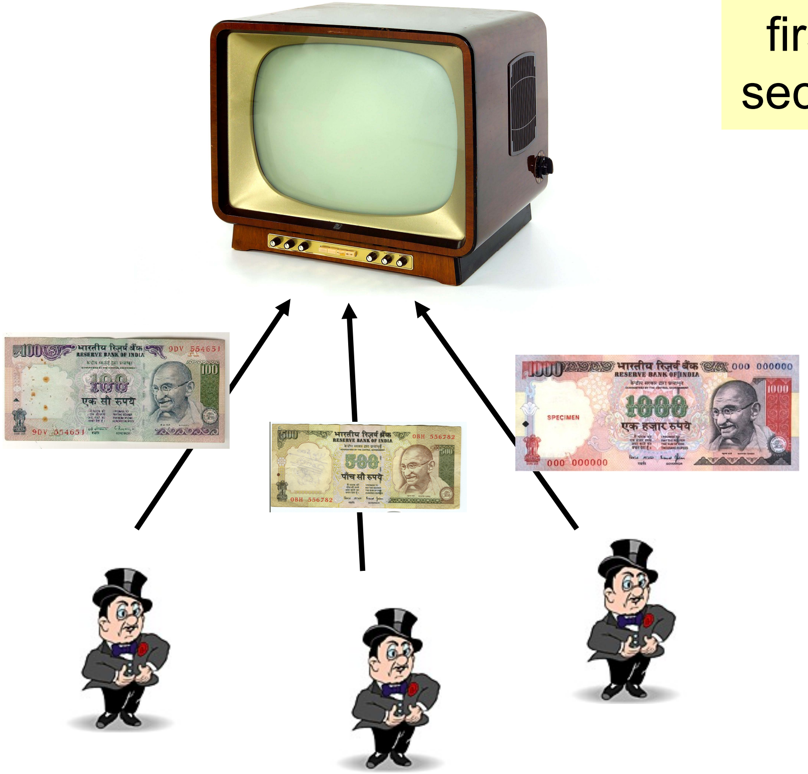

first-price: bid 3 wins, pays 1000
second-price: bid 3 wins, pays 500

Each bid depends on

- –bidder's true valuation for the item (utility = valuation – payment)
- –bidder's beliefs over what others will bid (→ game theory)
- –and... the auction mechanism used

first-price auction: it does not make sense to bid your true valuation

- –Even if you win, your utility will be 0…

second-price auction: it always makes sense to bid your true valuation (get only more, price independent of own bid => Vickrey mechanism)

Combinatorial Auctions

Simultaneously for sale:

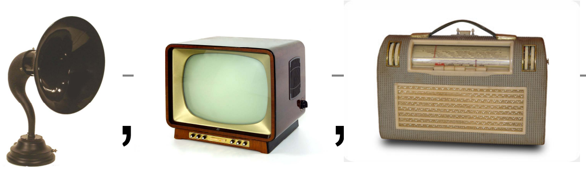

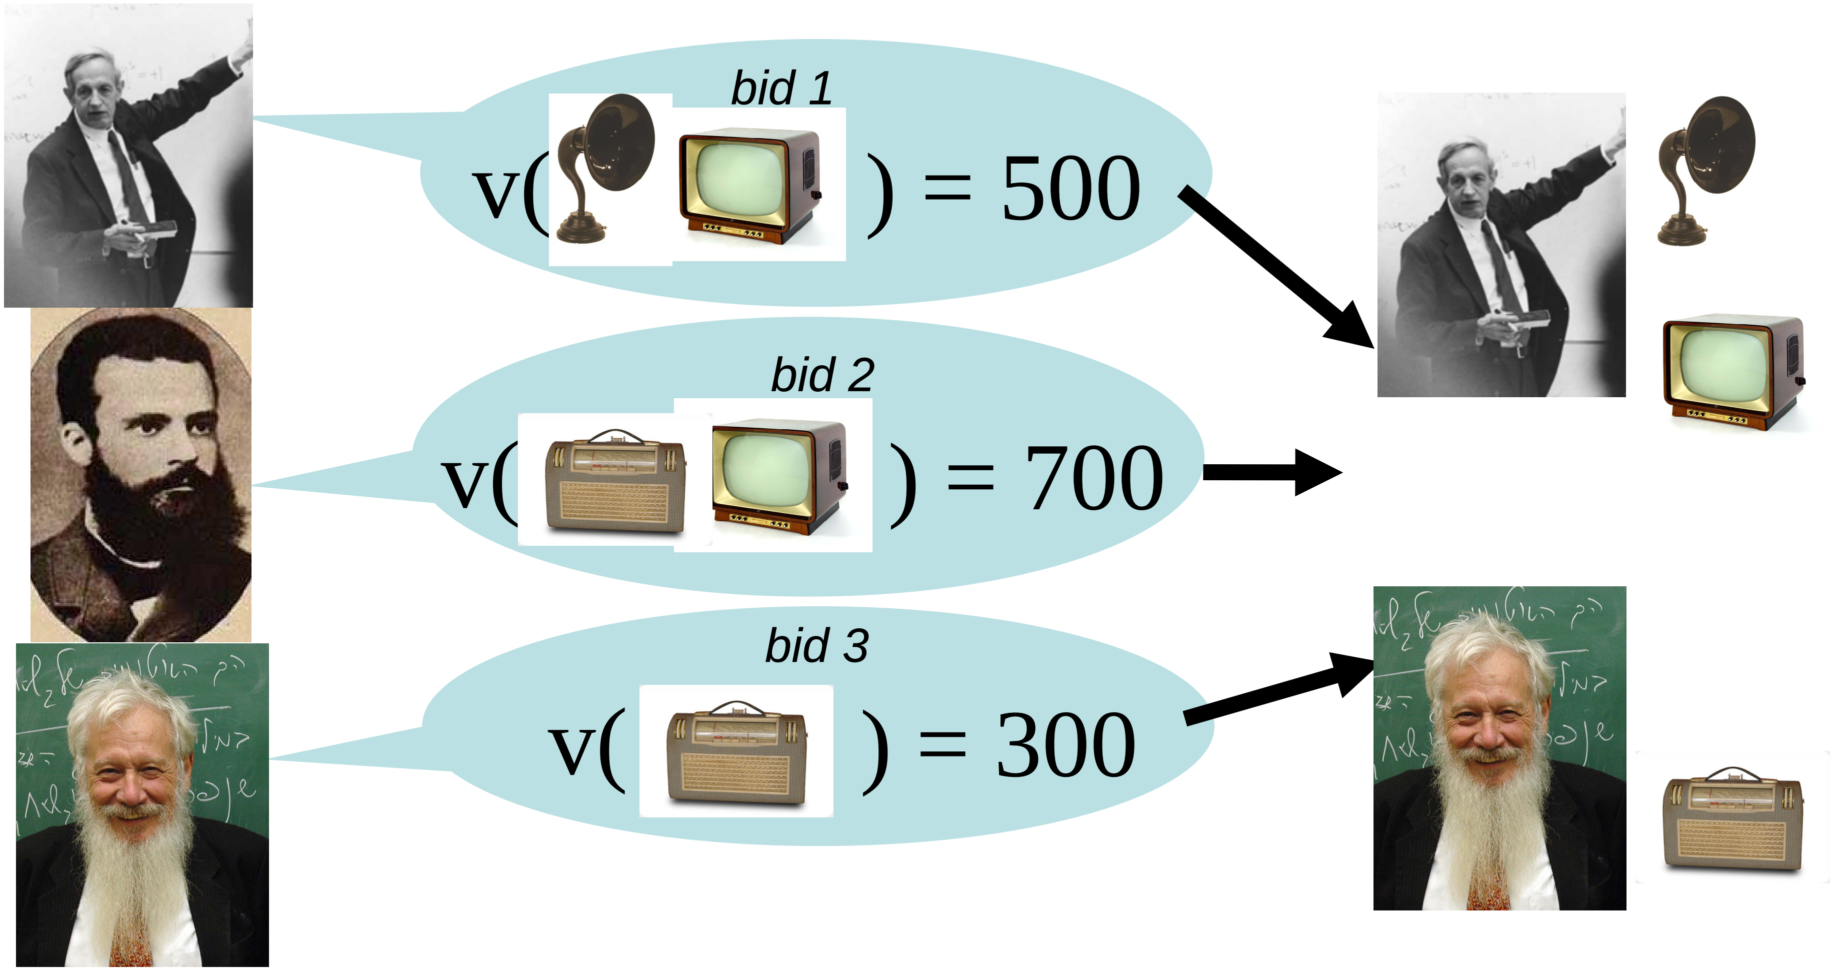

bid 1

v( ) = 500

bid 2

v( ) = 700

bid 3

v( ) = 300

used in truckload transportation, industrial procurement, radio spectrum allocation, …

Combinatorial Auction Problems

- **Winner determination** problem
  - Deciding which bids win is a hard computational problem (in general NP-hard)
- **Preference elicitation** (communication) problem
  - In general, each bidder may have a different value for each bundle
  - But it may be impractical to bid on every bundle (there are exponentially many bundles)
- **Mechanism design** problem
  - How do we get the bidders to behave so that we get good outcomes?

- These problems interact in nontrivial ways
  - E.g. limited computational or communication capacity limits mechanism design options

A Sample Mechanism (1)

Important goals of mechanism

1. Truth revealing / strategy proof
   (each player tells her true private utility to the mechanism)
2. No positive payments
   (no monetary transfer from mechanism to player)
3. Voluntary participation
   (no one must participate in mechanism)
4. Consumer sovereignity
   (if player declares very high utility, she should be serviced)
5. Maximize overall welfare
   (maximize sum of utility of players selected by mechanism minus costs)
6. Budget balance
   (sum of payments of selected players = cost of solution)

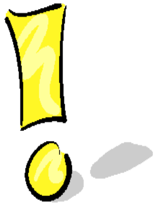

It‘s *impossible* to reach all these goals.

However, it is possible to reach goals 1-5.

A Sample Mechanism (2)

- Example: High-speed train from Wroclaw to Paris
  - costs billions => need to ask tax payers!

A Sample Mechanism (3)

Key idea: whether you get it depends on your declaration, but the price for which you get it does not!

- Example: High-speed train from Wroclaw to Paris
  - costs billions => need to ask tax payers!


© dean chamberland

- Mechanism: we ask all tax payers about their utility of this train
  - If (sum of utilities > cost) => build train
  - otherwise don‘t

A Sample Mechanism (4)

- One of the only mechanisms that works!
- Basic idea:

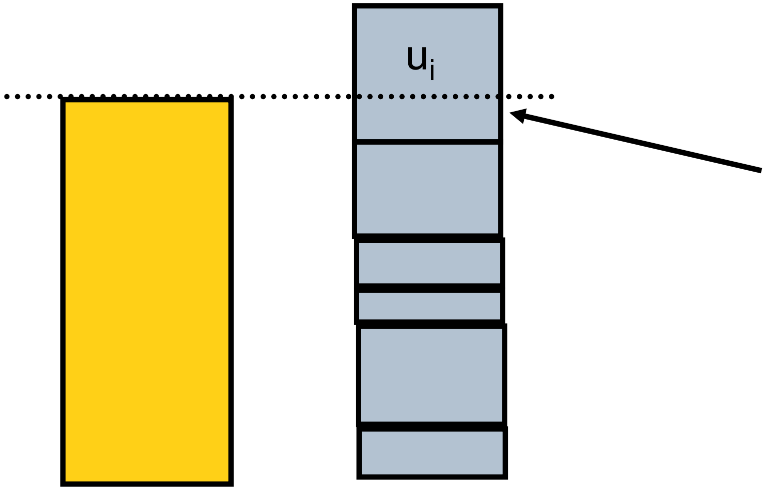

Player i needs to pay the remaining amount of money if all other utilities are summed up.

Cost of train

Utilities of players

A Sample Mechanism (5)

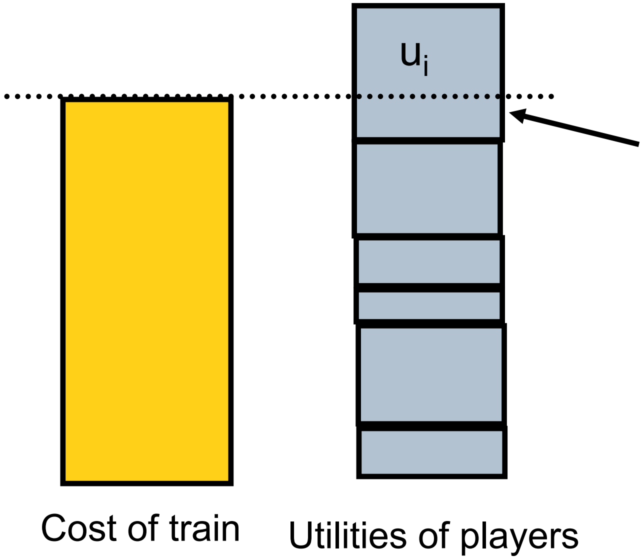

Player i needs to pay the remaining amount of money if all other utilities are summed up.

- Truth-revealing?

  - If without me it‘s already built => it‘s free for me!

  - If not, project depends on my contribution to the overall utility (I never have to pay more than my utility!)

  - It‘s bad to say less than my true utility, project may not be built!

  - It‘s bad to say more than my true utility, may have to pay more than utility!

A Sample Mechanism (6)

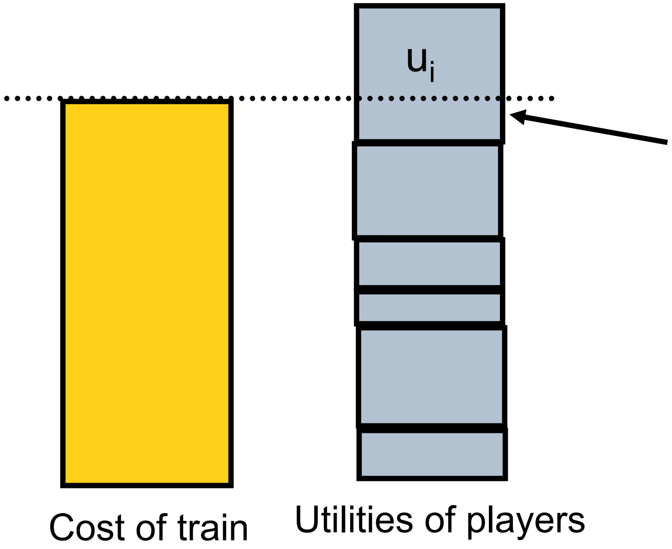

Player i needs to pay the remaining amount of money if all other utilities are summed up.

- However, of course, mechanism is not budget balanced!

  - project may be almost free for everybody!

Mechanism Design without Payments?

Extended Prisoners’ Dilemma (1)

- Mechanism design by creditability
  - creditabile designer can sometimes implement profiles for free
  - interesting, e.g., in wireless networks where monetary transfers are problematic!

- A bimatrix game with two bank robbers
  - A bank robbery (unsure, *video tape*) and a minor crime (sure, *DNA*)
  - Players are interrogated independently

Robber 2

		silent	testify	confess
Robber 1	silent	**3 3**	**0 4**	**0 0**
	testify	**4 0**	**1 1**	**0 0**
	confess	**0 0**	**0 0**	**0 0**

Extended Prisoners' Dilemma (2)

- A bimatrix game with two bank robbers

Robber 2

		silent	testify	confess
Robber 1	silent	3 3	0 4	0 0
	testify	4 0	1 1	0 0
	confess	0 0	0 0	0 0

Payoff = number of *saved* years in prison

Silent = Deny bank robbery

Testify = Betray other player (provide evidence of other player's bankrobbery)

Confess = Confess bank robbery (prove that they acted together)

Extended Prisoners' Dilemma (3)

- Concept of non-dominated strategies

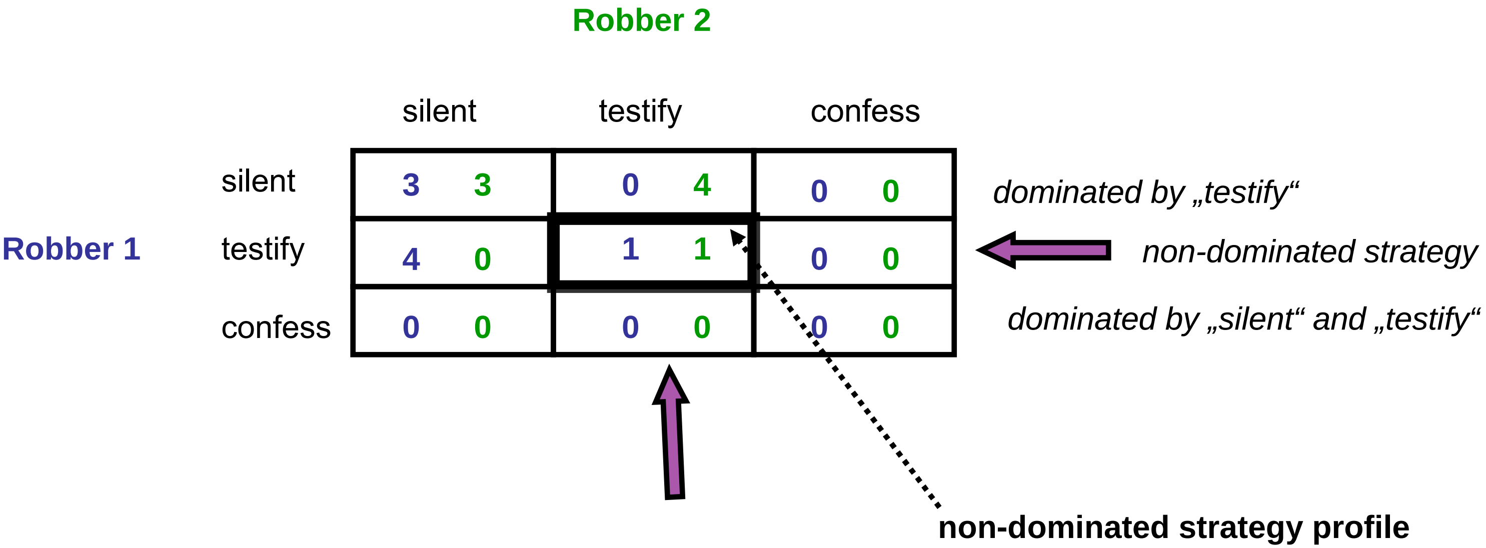

Robber 2

Robber 1

	silent	testify	confess
silent	**3 3**	**0 4**	**0 0**
testify	**4 0**	**1 1**	**0 0**
confess	**0 0**	**0 0**	**0 0**

dominated by „testify"

non-dominated strategy

dominated by „silent" and „testify"

non-dominated strategy profile

- Non-dominated strategy may not be unique!

In this talk, we use weakest assumption that players choose any non-dominated strategy. (here: both will testify)

Mechanism Design by Al Capone (1)

- Hence: both players testify = go 3 years to prison each.

Robber 2

		silent	testify	confess
	silent	3 3	0 4	0 0
Robber 1	testify	4 0	1 1	0 0
	confess	0 0	0 0	0 0

- Not good for gangsters‘ boss Al Capone!
  - Reason: Employees in prison!
  - Goal: Influence their decisions
  - Means: Promising certain payments for certain outcomes!

Mechanism Design by Al Capone (2)

	s	t	c
s	3 3	0 4	0 0
t	4 0	1 1	0 0
c	0 0	0 0	0 0

Original game **G**...

New non-dominated strategy profile!
Al Capone has to pay money worth 2 years in prison, but saves 4 years for his employees!
Net gain: 2 years!

+ ... plus Al Capone‘s monetary promises **V** ...

	s	t	c
s	1 1	2 0	
t	0 2		
c			

= ... yields new game **G(V)**!

	s	t	c
s	4 4	2 4	0 0
t	4 2	1 1	0 0
c	0 0	0 0	0 0

! Al Capone can save his employees 4 years in prison at low costs!

? Can the police do a similar trick to increase the total number of years the employees spend in prison?

Mechanism Design by the Police

	s	t	c
s	3 3	0 4	0 0
t	4 0	1 1	0 0
c	0 0	0 0	0 0

Original game **G**...

New non-dominated strategy profile! Both robbers will confess and go to jail for *four years each*! Police does not have to pay anything at all!

Net gain: 2

+

... plus the police‘ monetary promises **V** ...

	s	t	c
s			0 5
t			0 2
c	5 0	2 0	

=

... yields new game **G(V)**!

	s	t	c
s	3 3	0 4	0 5
t	4 0	1 1	0 2
c	5 0	2 0	0 0

Definition:

Strategy profile implemented by Al Capone has leverage (potential) of two: at the cost of money worth 2 years in prison, the players in the game are better off by 4 years in prison.

Strategy profile implemented by the police has a malicious leverage of two: at no costs, the players are worse off by 2 years.

Exact vs Non-Exact (1)

- Goal of a mechanism designer: implement a certain set of strategy profiles at low costs

  - I.e., make this set of profiles the (newly) non-dominated set of strategies

- Two options: Exact implementation and non-exact implementation

  - Exact implementation: *All* strategy profiles in the target region *O* are non-dominated

  - Non-exact implementation: Only a *subset* of profiles in the target region *O* are non-dominated

Exact vs Non-Exact (2)

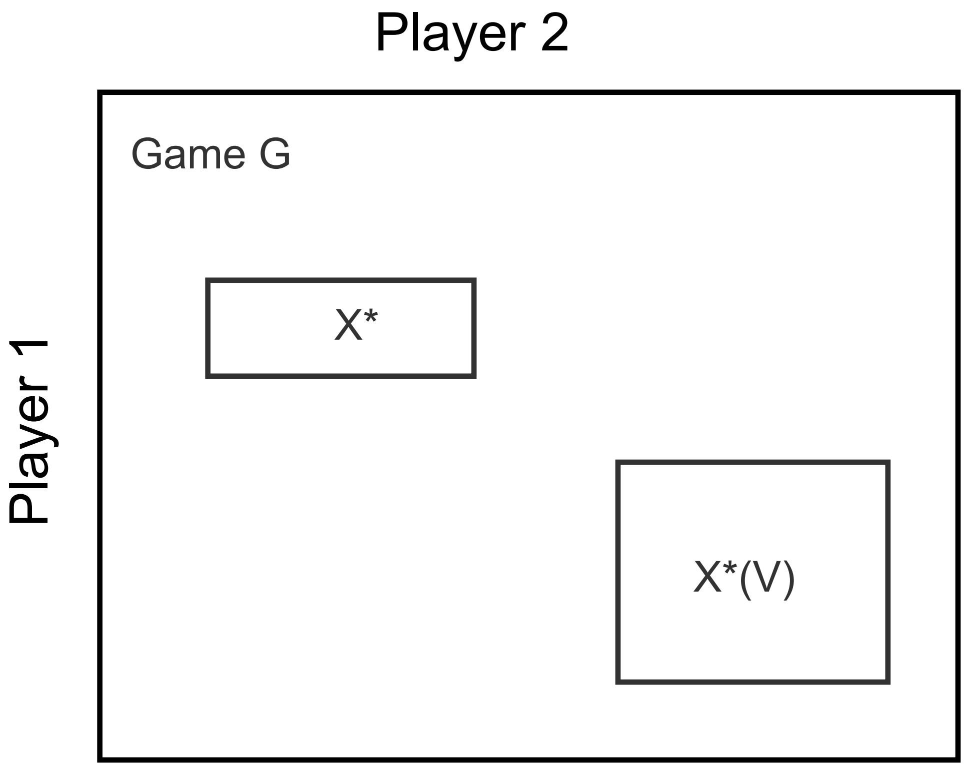

Player 2

Game G

X^*

$X^*(V)$

Player 1

X^* = non-dominated strategies *before* manipulation

$X^*(V)$ = non-dominated strategies *after* manipulation

Exact implementation:
$X^*(V) = O$
Non-exact implementation:
$X^*(V) \subset O$

Non-exact implementations can yield larger gains, as the mechanism designer can choose which subsets to implement!

Worst-Case vs Uniform Cost

What is the cost of implementing a target region *O*?

Two different cost models: worst-case implementation cost and uniform implementation cost

- Worst-case implementation cost: Assumes that players end up in the worst (most expensive) non-dominated strategy profile.
- Uniform implementation costs: The implementation costs is the average of the cost over all non-dominated strategy profiles. (All profiles are equally likely.)

Overview of Results

- **Worst-case leverage**
  - Polynomial time algorithm for computing leverage of singletons
  - Leverage for special games (e.g., zero-sum games)
  - Algorithms for general leverage (super polynomial time)

- **Uniform leverage**
  - Computing minimal implementation cost is *NP*-hard (for both exact and non-exact implementations); it cannot be approximated better than $\Omega(n \cdot \log(|X_i^* \setminus O_i|))$
  - Computing leverage is also *NP*-hard and also hard to approximate.
  - Polynomial time algorithm for singletons and super-polynomial time algorithms for the general case.

Sample Result: NP-hardness (1)

Theorem: Computing exact uniform implementation cost is *NP*-hard.

- Reduction from Set Cover: Given a set cover problem instance, we can *efficiently* construct a game whose minimal exact implementation cost yields a solution to the minimal set cover problem.

- As set cover is *NP*-hard, the uniform implementation cost must also be *NP*-hard to compute.

Sample Result: NP-hardness (2)

- Sample set cover instance:
  universe of elements **$U = \{e_1,e_2,e_3,e_4,e_5\}$**
  universe of sets **$S = \{S_1, S_2, S_3,S_4\}$**
  where **$S_1 = \{e_1,e_4\}$**, **$S_2=\{e_2,e_4\}$**, **$S_3=\{e_2,e_3,e_5\}$**, **$S_4=\{e_1,e_2,e_3\}$**

- Gives game...:

		elements					helper cols	
		e_1	e_2	e_3	e_4	e_5	d	r
elements	e_1	5	0	0	0	0	1	0
	e_2	0	5	0	0	0	1	0
	e_3	0	0	5	0	0	1	0
	e_4	0	0	0	5	0	1	0
	e_5	0	0	0	0	5	1	0
sets	s_1	5	0	0	5	0	0	0
	s_2	0	5	0	5	0	0	0
	s_3	0	5	5	0	5	0	0
	s_4	5	5	5	0	0	0	0

Player 2: payoff 1 everywhere except for column r (payoff 0)

Also works for more than two players!

Sample Result: NP-hardness (3)

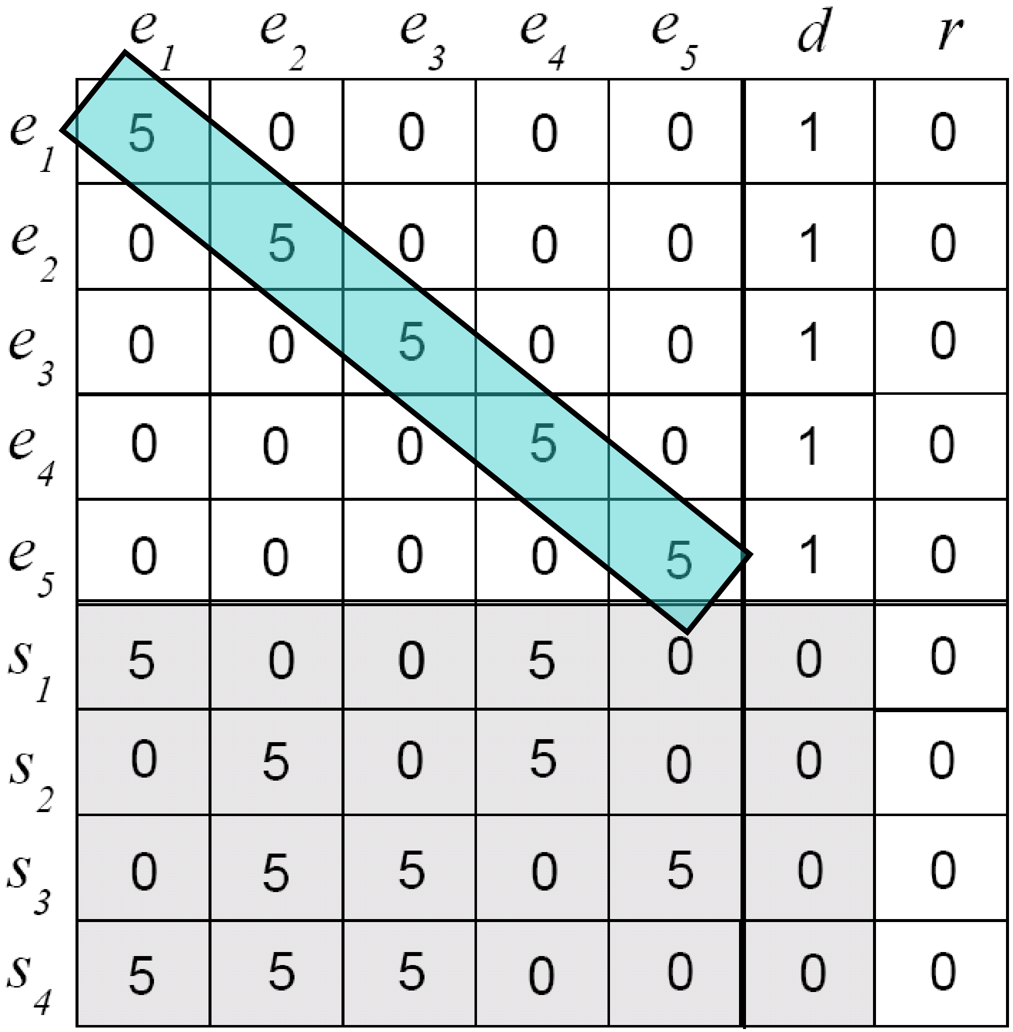

	e_1	e_2	e_3	e_4	e_5	d	r
e_1	5	0	0	0	0	1	0
e_2	0	5	0	0	0	1	0
e_3	0	0	5	0	0	1	0
e_4	0	0	0	5	0	1	0
e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	0	0
s_2	0	5	0	5	0	0	0
s_3	0	5	5	0	5	0	0
s_4	5	5	5	0	0	0	0

All 5s (=number of elements) in diagonal...

Sample Result: NP-hardness (3)

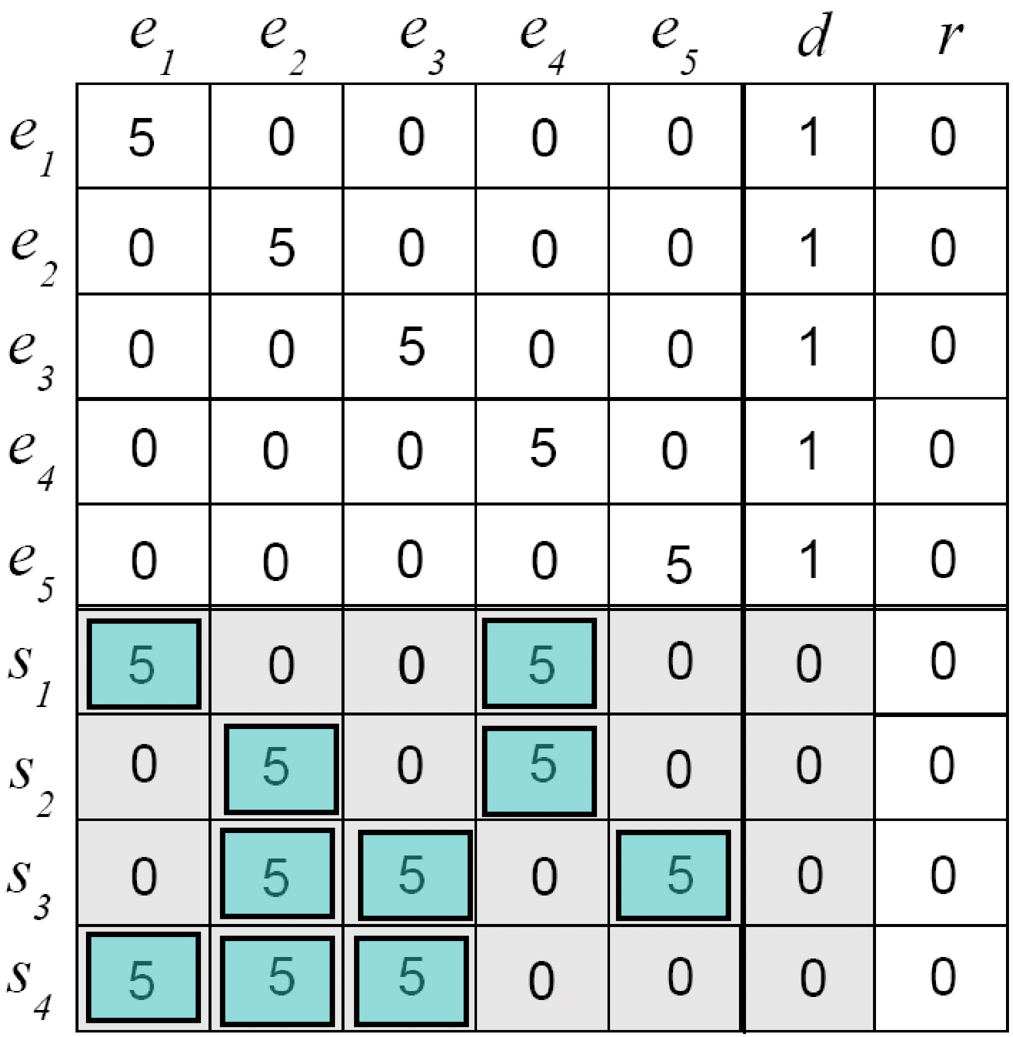

	e_1	e_2	e_3	e_4	e_5	d	r
e_1	5	0	0	0	0	1	0
e_2	0	5	0	0	0	1	0
e_3	0	0	5	0	0	1	0
e_4	0	0	0	5	0	1	0
e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	0	0
s_2	0	5	0	5	0	0	0
s_3	0	5	5	0	5	0	0
s_4	5	5	5	0	0	0	0

Set has a 5 for each element it contains... (e.g., $S_1 = \{e_1,e_4\}$)

Sample Result: NP-hardness (3)

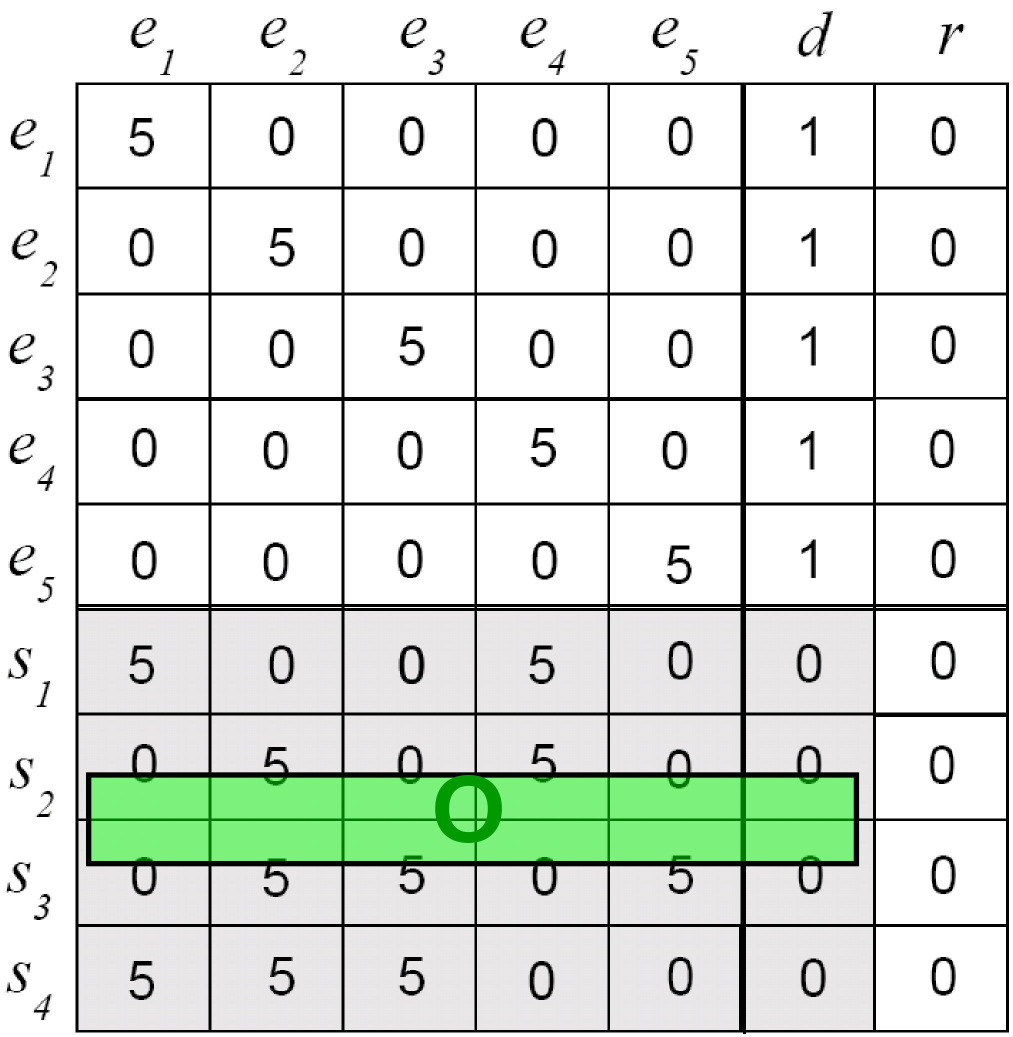

	e_1	e_2	e_3	e_4	e_5	d	r
e_1	5	0	0	0	0	1	0
e_2	0	5	0	0	0	1	0
e_3	0	0	5	0	0	1	0
e_4	0	0	0	5	0	1	0
e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	0	0
s_2	0	5	0	5	0	0	0
s_3	0	5	5	0	5	0	0
s_4	5	5	5	0	0	0	0

Goal: implementing this region O exactly at minimal cost

Sample Result: NP-hardness (3)

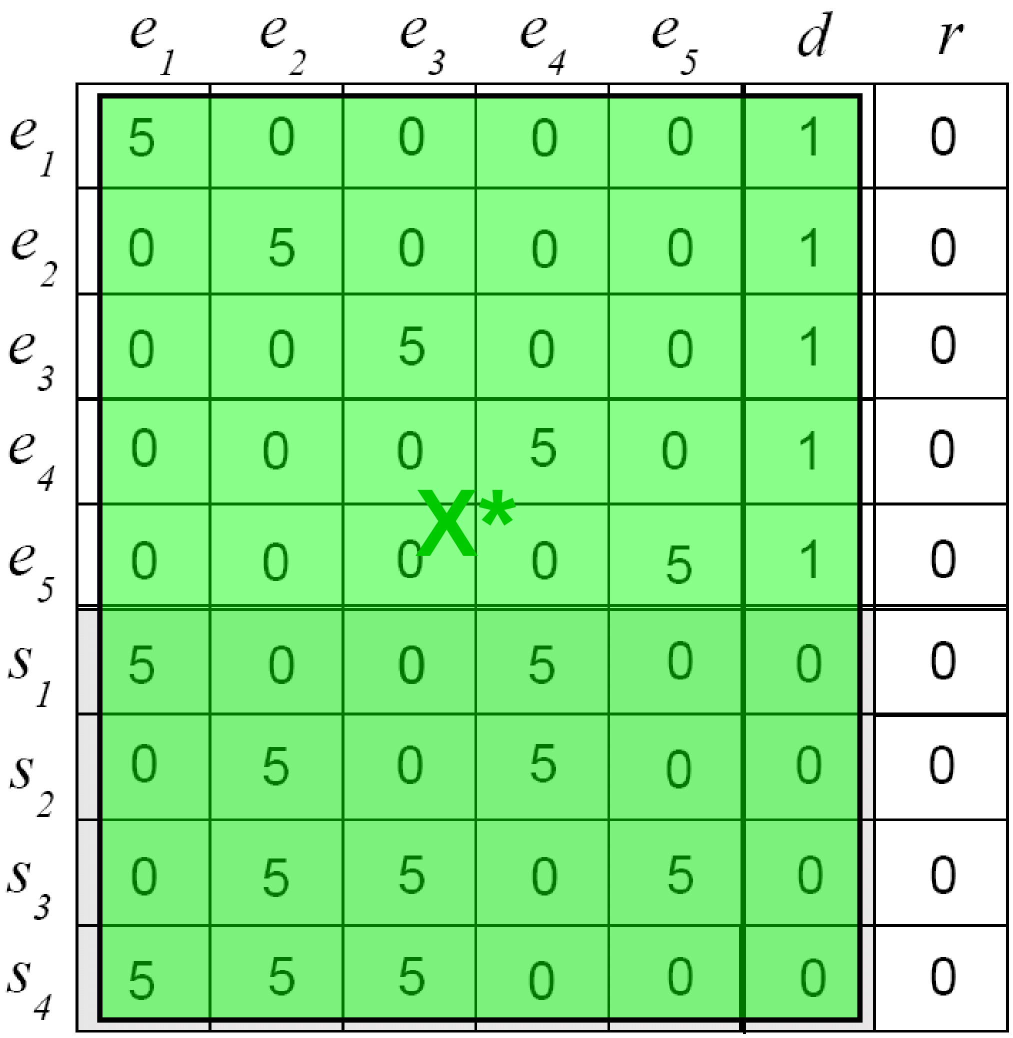

	e_1	e_2	e_3	e_4	e_5	d	r
e_1	5	0	0	0	0	1	0
e_2	0	5	0	0	0	1	0
e_3	0	0	5	0	0	1	0
e_4	0	0	0	5	0	1	0
e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	0	0
s_2	0	5	0	5	0	0	0
s_3	0	5	5	0	5	0	0
s_4	5	5	5	0	0	0	0

Originally, all these strategy profiles are non-dominated...

Sample Result: NP-hardness (3)

	e_1	e_2	e_3	e_4	e_5	d	r
e_1	5	0	0	0	0	1	0
e_2	0	5	0	0	0	1	0
e_3	0	0	5	0	0	1	0
e_4	0	0	0	5	0	1	0
e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	0	0
s_2	0	5	0	5	0	0	0
s_3	0	5	5	0	5	0	0
s_4	5	5	5	0	0	0	0

It can be shown that the minimal cost implementation only makes 1-payments here...

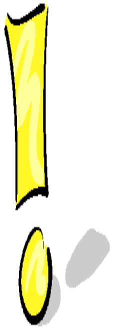

In order to dominate strategies above, we have to select minimal number of sets which covers all elements! (minimal set cover)

Sample Result: NP-hardness (3)

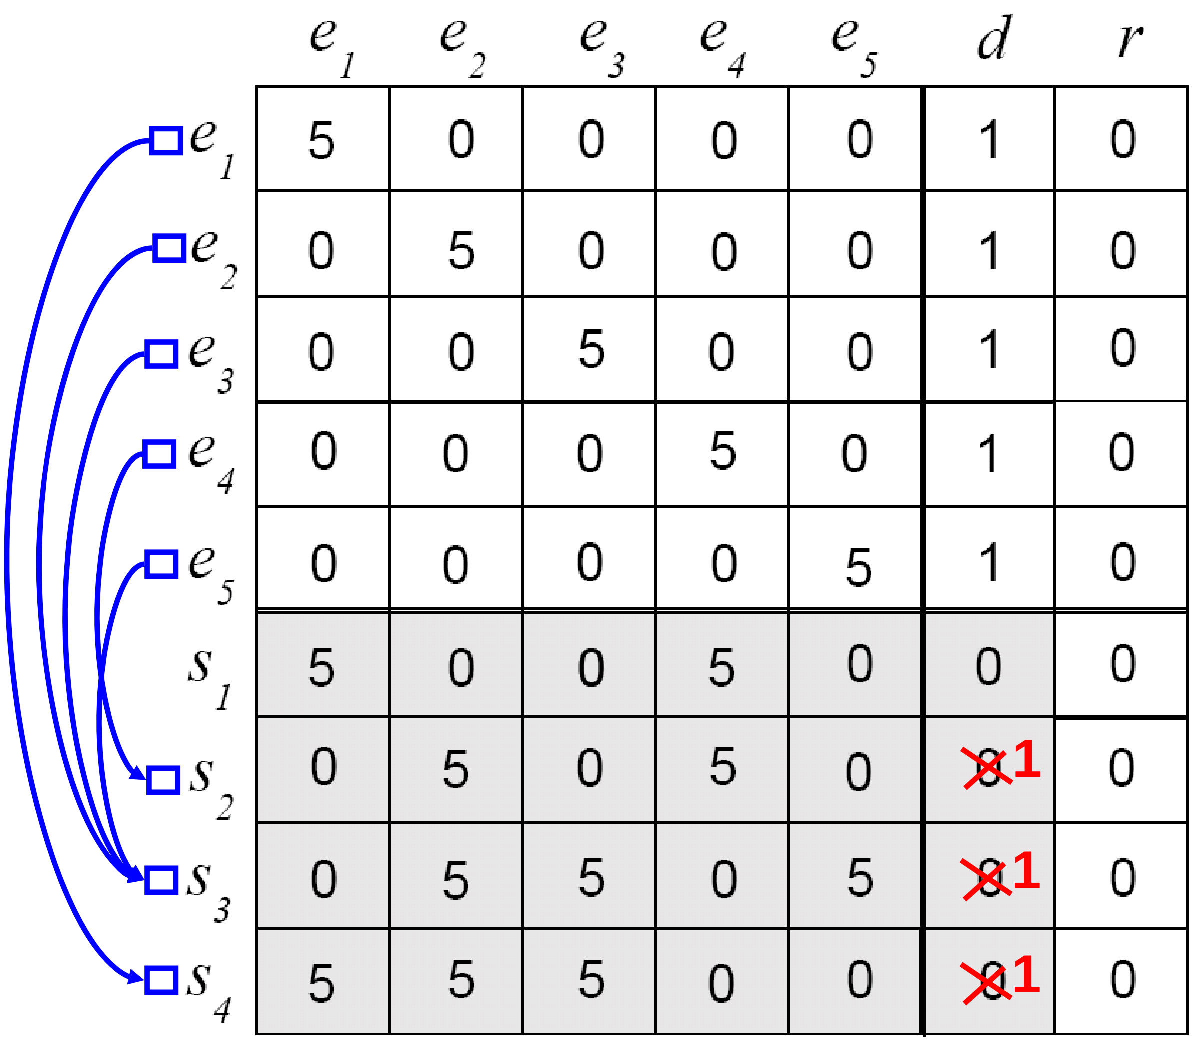

	e_1	e_2	e_3	e_4	e_5	d	r
☐ e_1	5	0	0	0	0	1	0
☐ e_2	0	5	0	0	0	1	0
☐ e_3	0	0	5	0	0	1	0
☐ e_4	0	0	0	5	0	1	0
☐ e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	0	0
☐ s_2	0	5	0	5	0	~~0~~1	0
☐ s_3	0	5	5	0	5	~~0~~1	0
☐ s_4	5	5	5	0	0	~~0~~1	0

A possible solution:
S_2, S_3, S_4
„dominates“ or
„covers“ all
elements above!
Implementation
costs: 3

Sample Result: NP-hardness (3)

	e_1	e_2	e_3	e_4	e_5	d	r
e_1	5	0	0	0	0	1	0
e_2	0	5	0	0	0	1	0
e_3	0	0	5	0	0	1	0
e_4	0	0	0	5	0	1	0
e_5	0	0	0	0	5	1	0
s_1	5	0	0	5	0	~~0~~ 1	0
s_2	0	5	0	5	0	0	0
s_3	0	5	5	0	5	~~0~~ 1	0
s_4	5	5	5	0	0	0	0

A better solution:
cost 2!

Sample Result: NP-hardness (4)

- A similar thing works for non-exact implementations!

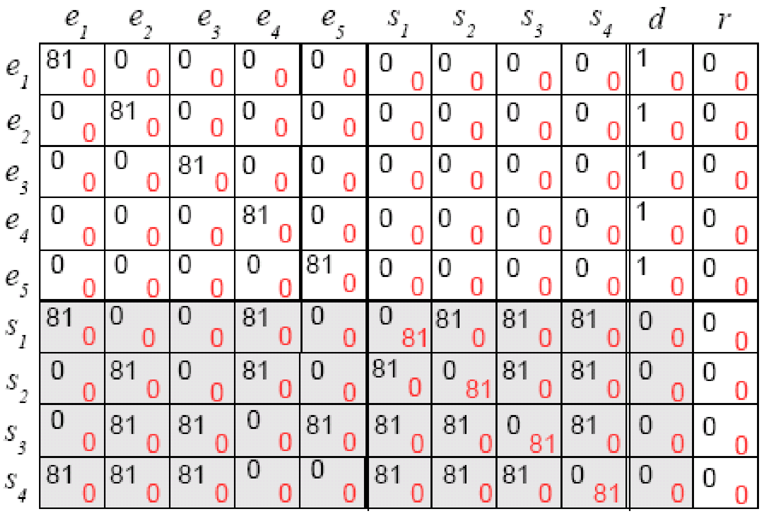

	e_1	e_2	e_3	e_4	e_5	s_1	s_2	s_3	s_4	d	r
e_1	81 0	0 0	0 0	0 0	0 0	0 0	0 0	0 0	0 0	1 0	0 0
e_2	0 0	81 0	0 0	0 0	0 0	0 0	0 0	0 0	0 0	1 0	0 0
e_3	0 0	0 0	81 0	0 0	0 0	0 0	0 0	0 0	0 0	1 0	0 0
e_4	0 0	0 0	0 0	81 0	0 0	0 0	0 0	0 0	0 0	1 0	0 0
e_5	0 0	0 0	0 0	0 0	81 0	0 0	0 0	0 0	0 0	1 0	0 0
s_1	81 0	0 0	0 0	81 0	0 0	0 81	81 0	81 0	81 0	0 0	0 0
s_2	0 0	81 0	0 0	81 0	0 0	81 0	0 81	81 0	81 0	0 0	0 0
s_3	0 0	81 0	81 0	0 0	81 0	81 0	81 0	0 81	81 0	0 0	0 0
s_4	81 0	81 0	81 0	0 0	0 0	81 0	81 0	81 0	0 81	0 0	0 0

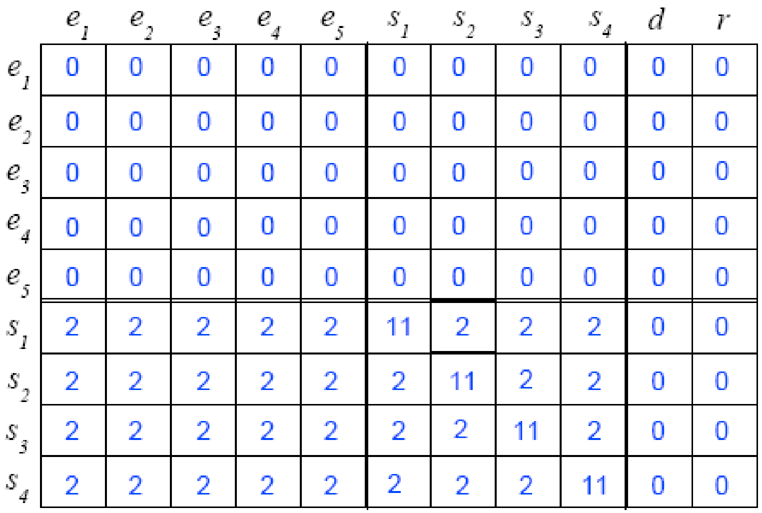

	e_1	e_2	e_3	e_4	e_5	s_1	s_2	s_3	s_4	d	r
e_1	0	0	0	0	0	0	0	0	0	0	0
e_2	0	0	0	0	0	0	0	0	0	0	0
e_3	0	0	0	0	0	0	0	0	0	0	0
e_4	0	0	0	0	0	0	0	0	0	0	0
e_5	0	0	0	0	0	0	0	0	0	0	0
s_1	2	2	2	2	2	11	2	2	2	0	0
s_2	2	2	2	2	2	2	11	2	2	0	0
s_3	2	2	2	2	2	2	2	11	2	0	0
s_4	2	2	2	2	2	2	2	2	11	0	0

- From hardness of costs follows hardness of leverage!

$$lev_{UNI}(O) = \max_{V \in \mathcal{V}^*(O)}\{\varnothing_{z \in X^*(V)}\{U(z) - V(z)\}\} - \varnothing_{z \in X^*} U(x^*) = \max_{V \in \mathcal{V}^*(O)}\{\varnothing_{z \in X^*(V)} U(z) - \varnothing_{z \in X^*(V)} V(z)\} - \varnothing_{x^* \in X^*} U(x^*) = \varnothing_{z \in X^*(V)} U(z) - \min_{V \in \mathcal{V}^*(O)}\{\varnothing_{z \in X^*(V)} V(z)\} - \varnothing_{x^* \in X^*} U(x^*) = \varnothing_{z \in X^*(V)} U(z) - \mathbf{k}^*_{\mathrm{UNI}}(\mathbf{O}) - \varnothing_{x^* \in X^*} U(x^*).$$

$$mlev_{UNI}(O) = \varnothing_{x^* \in X^*} U(x^*) - \min_{V \in \mathcal{V}^*(O)}\{\varnothing_{z \in X^*(V)}\{U(z) + 2V(z)\}\} = \varnothing_{x^* \in X^*} U(x^*) - \varnothing_{z \in X^*(V)} U(z) - 2\min_{V \in \mathcal{V}^*(O)}\{\varnothing_{z \in X^*(V)} V(z)\} = \varnothing_{x^* \in X^*} U(x^*) - \varnothing_{z \in X^*(V)} U(z) - 2\mathbf{k}^*_{\mathrm{UNI}}(\mathbf{O}).$$

Conclusion

Limitations of Game Theory? (1)

- Challenge: heterogeneity and generality
  - The more general the player types, the asymmetric information, etc., the less can be said about the equilibria
  - How lazy are players? How much do they invest to find out what their optimal strategy is at all?
  - Most general: arbitrary or Byzantine behavior, and arbitrary knowledge

- Difficulties in explaining certain phenomena with autonomous agent model
  - For instance wireless networking: why is there a throughput at all?
  - It seems that there is no advantage of not trying to send at any moment of time?

Limitations of Game Theory? (2)

- It is difficult to make predictions or give lower bounds with game theory
  - Game theory is based on behavioral assumptions
  - Some interests which are not taken into account may improve collaboration...

- Your remark?

Wielkie dzieki!

Slides and papers at

http://www14.informatik.tu-muenchen.de/personen/schmiste/

Extra Slides

Potential Games

Types of Games

- Dummy Game: Unilateral deviations imply no change for deviating player, all strategy profiles are pure NE

a,b	c,b
a,d	c,d

- Coordination Game: mutual gain for consistent decision, multiple pure NE, at least one Pareto efficient NE

9,9	0,0
0,0	9,9

drive left, right

9,4	0,0
0,0	4,9

battle of the sexes

9,9	0,7
7,0	7,7

stag hunt

- Exact Potential Games: incentive to change can be expressed in global function P: S → R, such that

$$P(s_i, s_{-i}) - P(s'_i, s_{-i}) = u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})$$

Exact Potential Games

Exact Potential Games: ∃ function P: S → R, such that

$$P(s_i, s_{-i}) - P(s'_i, s_{-i}) = u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})$$

Example:

Is stag hunt an exact potential game?

9,9	0,7
7,0	7,7

$$u_1(S_1,S_2) - u_1(H_1,S_2) = 2 = P(S_1,S_2) - P(H_1,S_2)$$

$$u_1(S_1,H_2) - u_1(H_1,H_2) = -7 = P(S_1,H_2) - P(H_1,H_2)$$

$$u_2(S_1,H_2) - u_2(S_1,S_2) = -2 = P(S_1,H_2) - P(S_1,S_2)$$

$$u_2(H_1,H_2) - u_2(H_1,S_2) = 7 = P(H_1,H_2) - P(H_1,S_2)$$

$$P(s) := \begin{cases} 2 & \text{if } s=(S_1,S_2) \\ 0 & \text{if } s=(H_1,\quad) \\ 0 & \text{if } s=(S_1,H_2) \\ 7 & \text{if } s=(H_1,H_2) \end{cases}$$

✓

YES!

Note: P not unique

Exact Potential Games

Exact Potential Games: ∃ function P: S → R, such that

$$P(s_i, s_{-i}) - P(s'_i, s_{-i}) = u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})$$

Why should we study potential games?

Many networking problems can be formulated as Potential Games!

- Routing Problems
- Power Control in wireless networks
- Sensor coverage problems
- ...

NE are local maxima!
Coincidence?

Potential Games have nice properties!

9,9	0,7
7,0	7,7

$$P(s) := \begin{cases} 2 & \text{if } s=(S_1,S_2) \\ 0 & \text{if } s=(H_1,S_2) \\ 0 & \text{if } s=(S_1,H_2) \\ 7 & \text{if } s=(\mathrm{H}_1,H_2) \end{cases}$$

makes analysis
easier!

Exact Potential Games Properties

Properties:

1. Sum of coordination and dummy games

3,3	0,2
2,0	2,2

Coordination

+

4,3	1,3
4,2	1,2

Dummy

=

7,6	1,5
6,2	3,4

Exact Potential Games Properties

Properties:

1. Sum of coordination and dummy games

2. Pure NE: local maximum of potential function

<u>Proof</u>:

- Let s be the profile s maximizing P
- Suppose it is not a NE, we can improve by deviating to new profile s’, where $P(s') - P(s) = u_i(s') - u_i(s) > 0$
  - Thus, P(s’) < P(s) , contradicting that s maximizes P.

  => set of pure NE is set of local max of P

3,3	0,2
2,0	2,2

Coordination

+

4,3	1,3
4,2	1,2

Dummy

=

7,6	1,5
6,2	3,4

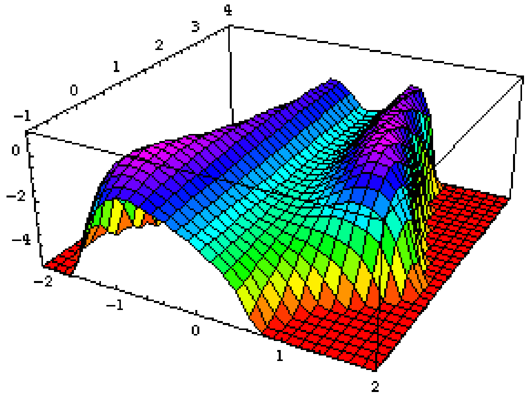

Exact Potential Games Properties

Properties:

1. Sum of coordination and dummy games
2. Pure NE: local maximum of potential function
3. Best-response dynamics converge to NE

Proof:

no cycles in best response graph

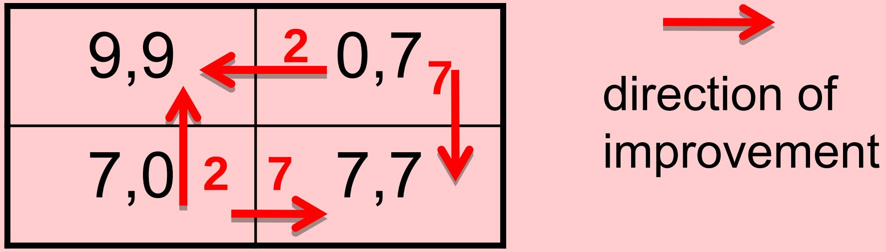

9,9	0,7
7,0	7,7

direction of improvement

If there was a cycle, no potential function possible

3,3	0,2
2,0	2,2

Coordination

+

4,3	1,3
4,2	1,2

Dummy

=

7,6	1,5
6,2	3,4

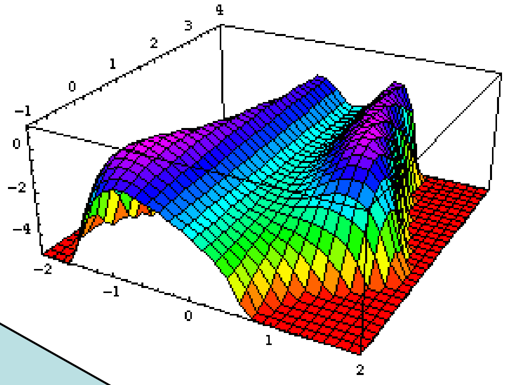

- no communication required
- Convergence Speed?

$$\frac{\max P(s_i, s_{-i})}{\min(P(s_i, s_{-i}) - P(s'_i, s_{-i}))}$$

Exact Potential Games Properties

Properties:

1. Sum of coordination and dummy games

2. Pure NE: local maximum of potential function

3. Best-response dynamics converge to NE

4. For continuous utility functions

$$\frac{\partial P}{\partial a_i} = \frac{\partial u_i}{\partial a_i}\ , \qquad \frac{\partial^2 P}{\partial a_i \partial a_j} = \frac{\partial^2 u_i}{\partial a_i \partial a_j} = \frac{\partial^2 u_j}{\partial a_i \partial a_j}$$

for every $i, j \in N$

5. Efficiency guarantees

Coordination

3,3	0,2
2,0	2,2

+

Dummy

4,3	1,3
4,2	1,2

=

7,6	1,5
6,2	3,4

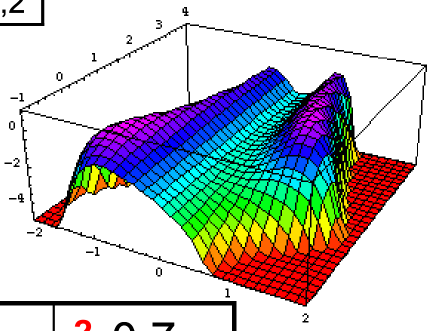

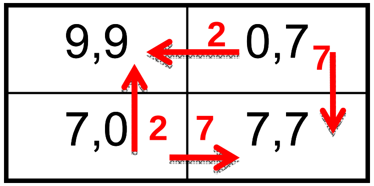

9,9	0,7
7,0	7,7

$$V(a) = \sum_{i \in N} \int_0^1 \frac{\partial u_i}{\partial a}(x(t))\, x_i^{'}(t)\, dt$$

See examples on following slides...

Multicast Game: Edge Sharing

- directed graph $G = (V, E)$, k players
- edge costs $c_e \geq 0$
- If x players share edge e, they each pay c_e/x
- goal: connect s_j to t_j with minimal cost

Example:

1	2	1 pays	2 pays
outer	outer	4	8
outer	middle	4	5 + 1
middle	outer	5 + 1	8
middle	middle	5/2 + 1	5/2 + 1

NE!

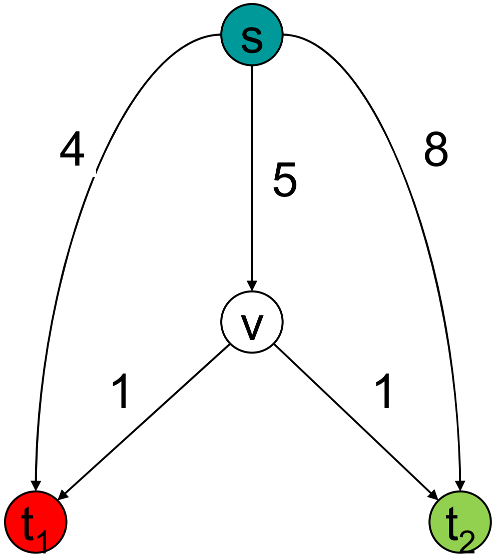

Edge Sharing: Efficiency?

Social optimum (OPT)
minimizes total costs of all players

Observation:
- there can be many NE
- NE not necessarily equal to OPT

Price of Anarchy = $\dfrac{\text{cost of } \textbf{worst} \text{ NE}}{\text{cost of OPT}}$

Price of Stability = $\dfrac{\text{cost of } \textbf{best} \text{ NE}}{\text{cost of OPT}}$

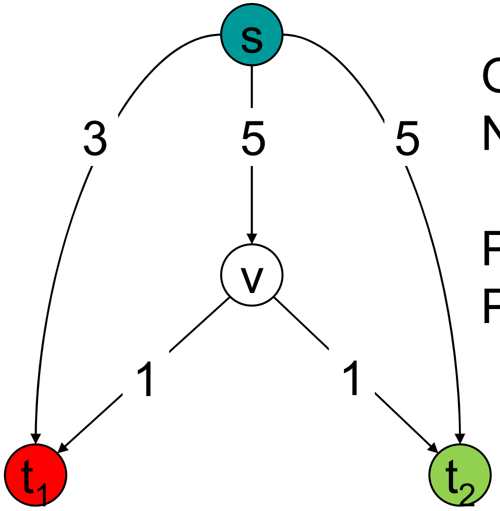

OPT = 7
NE = 8

PoA = 8/7
PoS = 8/7

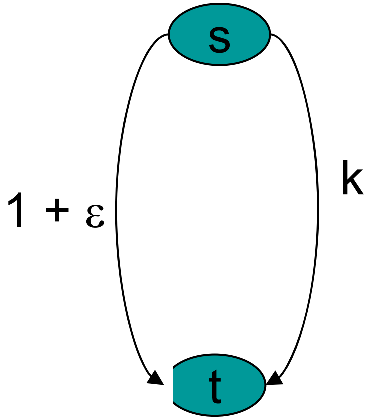

k players

OPT = 1+ ε
NE_1 = 1+ ε
NE_2 = k

PoA = k
PoS = 1

Edge Sharing: Bounds

Theorem: PoA $\le$ k

Proof:

- Let s be the worst NE
- Suppose by contradiction c(s) > k OPT
- Then, there exists a player i s.t. $c_i(s) > \text{OPT}$
- But i can deviate to OPT (by paying OPT alone)

  contradicts that N is a NE

Lower bound

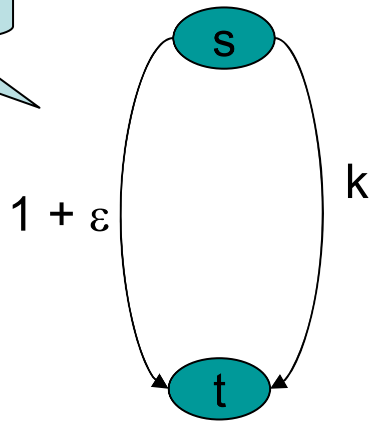

Theorem: PoS $\le$ O(log k)

Proof:

Potential Game

→ Potential function

→ best NE

Lower bound

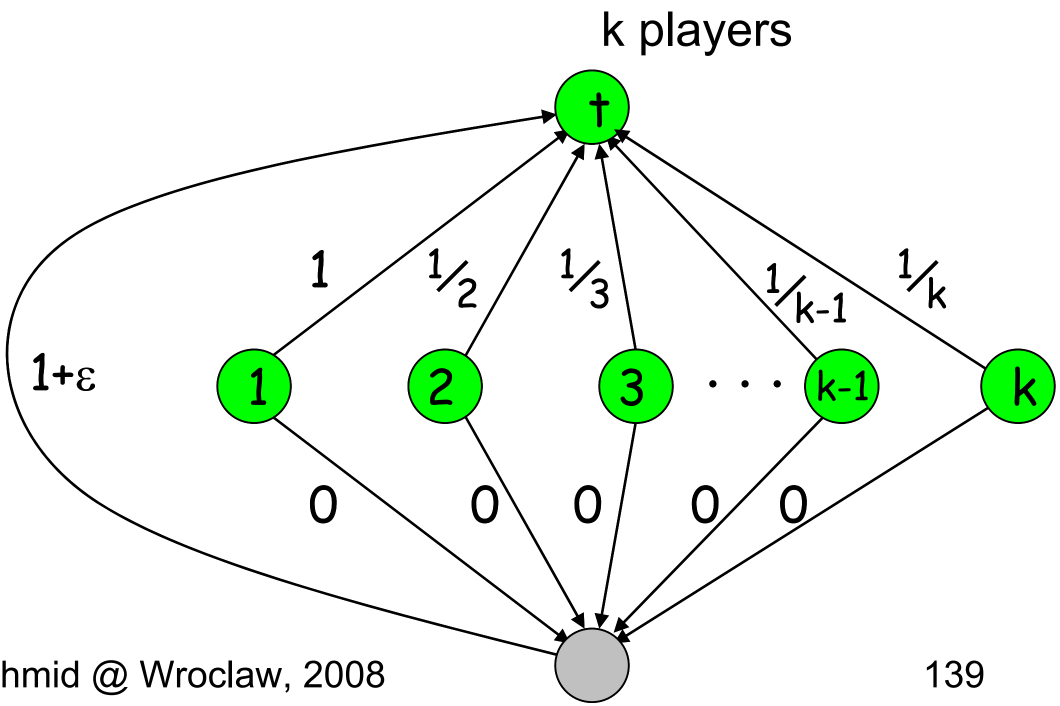

Routing Game

Congestion Game: cost depends on # users

- directed graph $G = (V, E)$, k players
- edge costs $c_e(x) \geq 0$
- If x players use edge e, they each pay $c_e(x)$
- goal: connect s_j to t_j with minimal cost (delay)

Example:

Player 1: $4 + 6 + x^2 + 8 = 22$

Player 2: $1/3 + 6 + x^2 + 1 = 11\ 1/3$

Efficiency? personal goal vs global goal

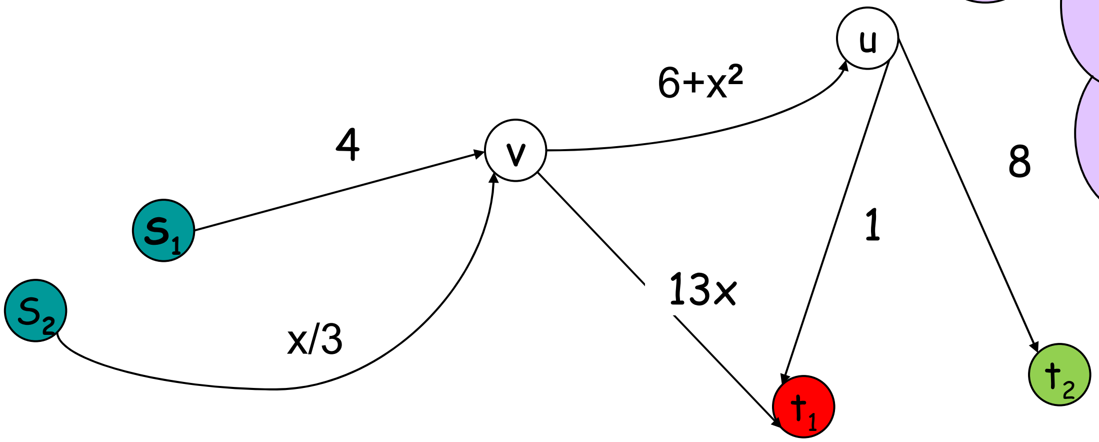

Routing Game: Efficiency?

Social optimum (OPT)
minimizes total costs of all players

Observation:

- there can be many NE
- NE not necessarily equal to OPT

Price of Anarchy = $\frac{\text{cost of **worst** NE}}{\text{cost of OPT}}$

Price of Stability = $\frac{\text{cost of **best** NE}}{\text{cost of OPT}}$

2 players

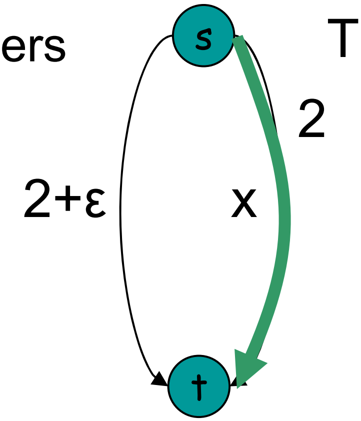

Total delay cost: 4

What happens?

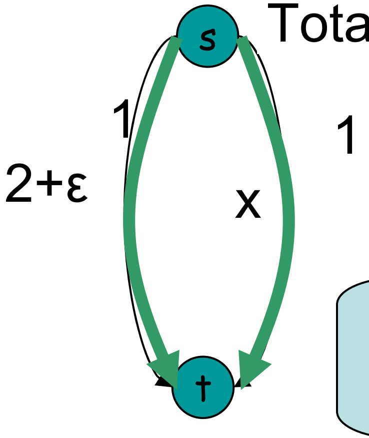

Total delay cost: 3+ε

Better, but no NE!

OPT = 3+ ε
NE = 4

PoA = 4/3
PoS = 4/3

Routing Game: Results [Roughgarden&Tardos]

PoA and PoS depend on delay function

always on 2-edge graphs!

- $c_e(x) = ax + b \rightarrow$ PoA $< 4/3$
- $c_e(x) = ax^d + \ldots \rightarrow$ PoA in $O(d/\log d)$
- continuous, nondecreasing delay functions:

rate = load/s

cost of Nash with rates r_i for all i $\leq$ cost of OPT with rates $2r_i$ for all i

Morale for the Internet: build for double rate

Repeated Games

Performance of Games

Game

- Interactions between players
- Strategies
- Utilities

	C_2	D_2
C_1	6,6	2,7
D_1	2,7	0,0

Convergence

- Existence, uniqueness of NE
- Conditions

Efficiency

- Social Welfare
- Pareto Efficiency
- Price of Anarchy

Fairness

- Ressources/utility shared equally?

Extensive Form Games: Tree Representation

So far: players choose actions simultaneously.

Does order of player moves influence outcome?

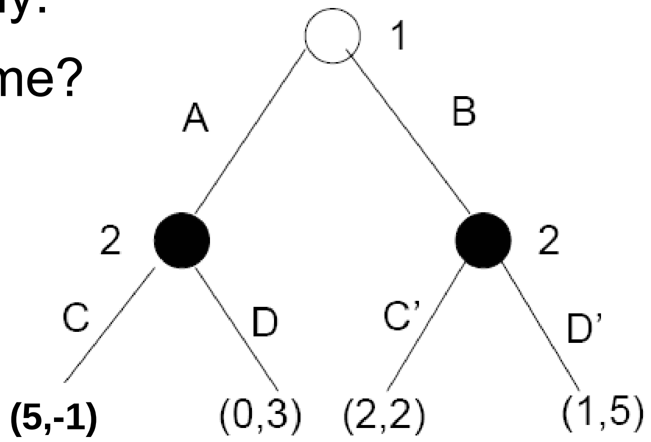

Components
- nodes: history of moves
- edges: actions taken
- payoffs based on history
- knowledge of nodes

Equivalence with games in strategic form?

- Player 1: {A,B}
- Player 2: {(C,C'),(C,D'),(D,C'),(D,D')}

	C,C'	C,D'	D,C'	D,D'
A	5,-1	2,-1	0,3	0,3
B	2,2	1,5	2,2	1,5

Any game can be represented as a tree !

Strategy

A strategy for a player is a complete plan of actions

It specifies a feasible action for the player in every contingency that the player might encounter.

Matched pennies game example:

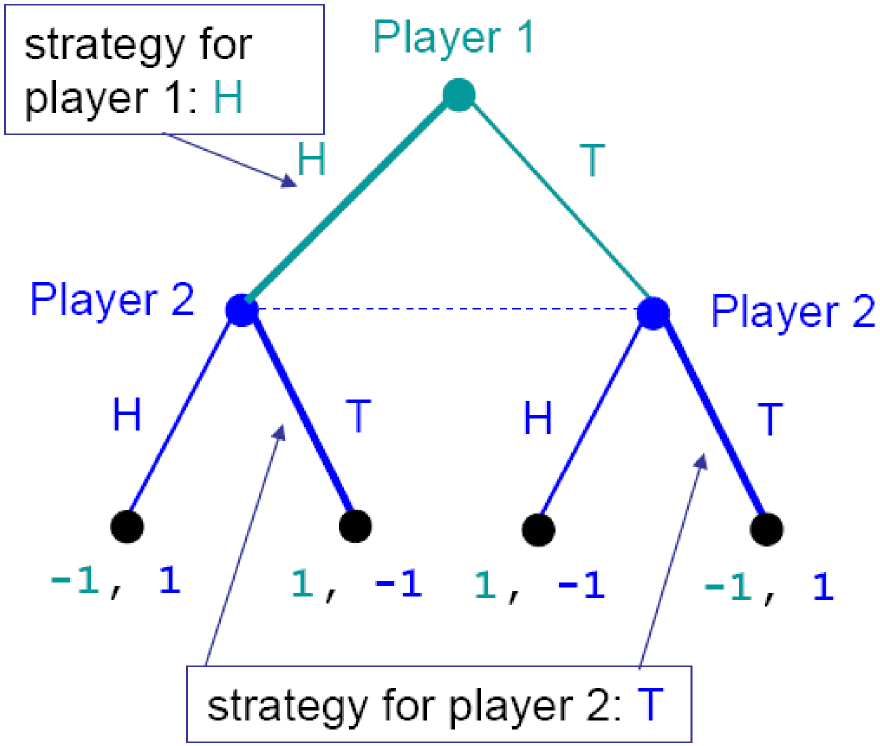

Repeated Games

Players condition their actions on opponents previous play

=> New strategies and equilibria!

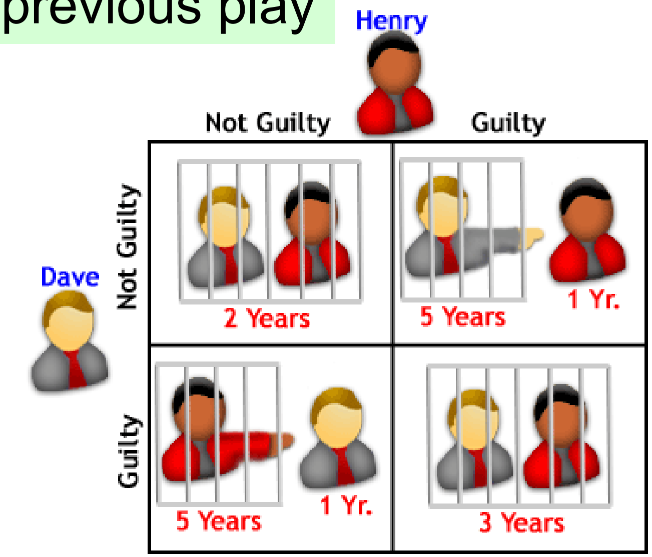

Example Strategies

- Always defect
- Always cooperate
- Tit-for-Tat (do what other did)
- Trigger (cooperate as long as other does)
- Generous Tit-for-Tat

  (after defect, cooperate with prob.)

Payoffs $g_i(a^t)$ depend on current actions, discount future payoffs: discount factor δ

Cumulative payoff (normalized)

$$\frac{1-\delta}{1-\delta^{T+1}} \sum_{t-0}^{T} \delta^{t} g_i\left(a^{t}\right)$$

Repeated Games

Trigger (cooperate as long as other does, afterwards defect) is a NE for T infinite.

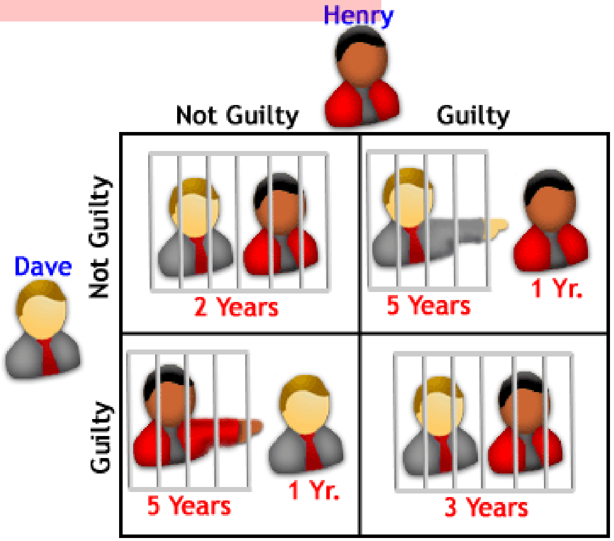

Two cases:

a) both play trigger

b) one defects at time t

Payoff in case a)

$(1-\delta)(2+2\delta+2\delta^2+\ldots)=2$

2 years in prison

Payoff in case b)

$(1-\delta)(2+2\delta+2\delta^2+\ldots+2\delta^{t-1}+\delta^t+3\delta^{t+1}+3\delta^{t+2}+\ldots)$

~3 years in prison

Cumulative payoff

$$\frac{1-\delta}{1-\delta^{T+1}}\sum_{t-0}^{T}\delta^t g_i\left(a^t\right)$$

Repeated Games

Trigger (cooperate as long as other does, afterwards defect) is a NE for T infinite and $\delta > 1/2$.

Two cases:

a) both play trigger

b) one defects at time t

1,1	-2,2
2,-2	0,0

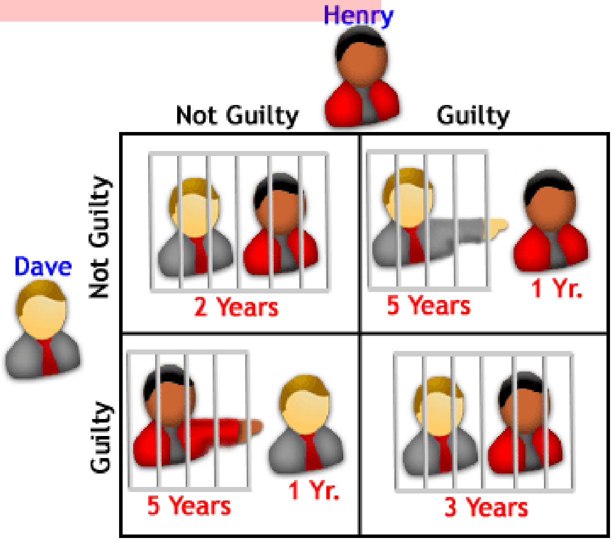

Payoff in case a)

$(1-\delta)(1+\delta+\delta^2+\ldots)=1$

Payoff in case b)

$(1-\delta)(1+1\delta+1\delta^2+\ldots+1\delta^{t-1}+2\delta^t+0)=1+\delta^t(1-2\delta)<1$

Cumulative payoff

$$\frac{1-\delta}{1-\delta^{T+1}}\sum_{t-0}^{T}\delta^t g_i\left(a^t\right)$$

Folk Theorem of Game Theory

Definitions:

- minimax condition: player minimizes the maximum possible loss
- feasible outcome: minimax condition satisfied for all players

Theorem:
In repeated games, any outcome is a feasible solution concept, if under that outcome the players' minimax conditions are satisfied.

In other words:

If players are sufficiently patient, then any feasible, individually rational payoffs can be enforced as an equilibrium.

Proof idea*: when players are patient, any finite one-period gain from deviation is outweighed by even a small loss in utility in every future period.*

Wireless Game Theory

Wireless Networks Challenges

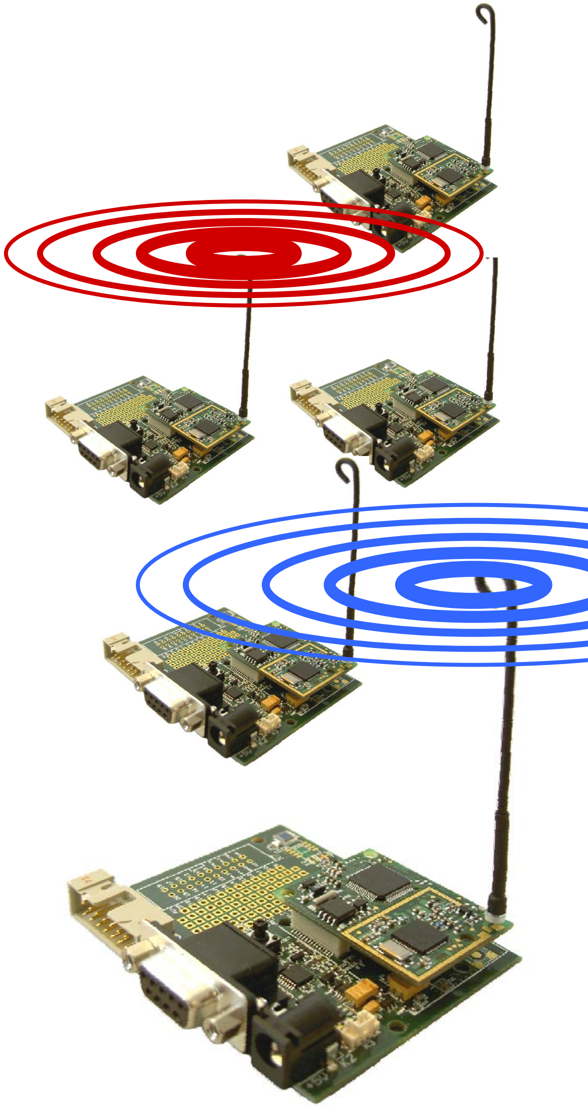

- **Rate of reliable data transmission limited:**
- noise (receiver & background)
- path losses (spatial diffusion & shadowing)
- multipath (fading & dispersion)
- interference (multiple-access & co-channel)
- dynamics (mobility, random-access & bursty traffic)
- limited transmitter power

- **No centralized solutions**

  (only for downlink of cellular networks)
- **Heterogeneity**

=> Game Theory and Mechanism Design

Wireless Networks Games

Goal: Avoid Interference

Resource allocation problem:

- Channel allocation (frequency, time, codes)
- Transmission power assignment
- Rate assignment
- Route selection

Strategies

Performance measure:

– BER (Bit Error Rate)

– SINR (Signal to Interference + Noise Ratio)

– Energy consumption

– Delay

– Throughput

Utility

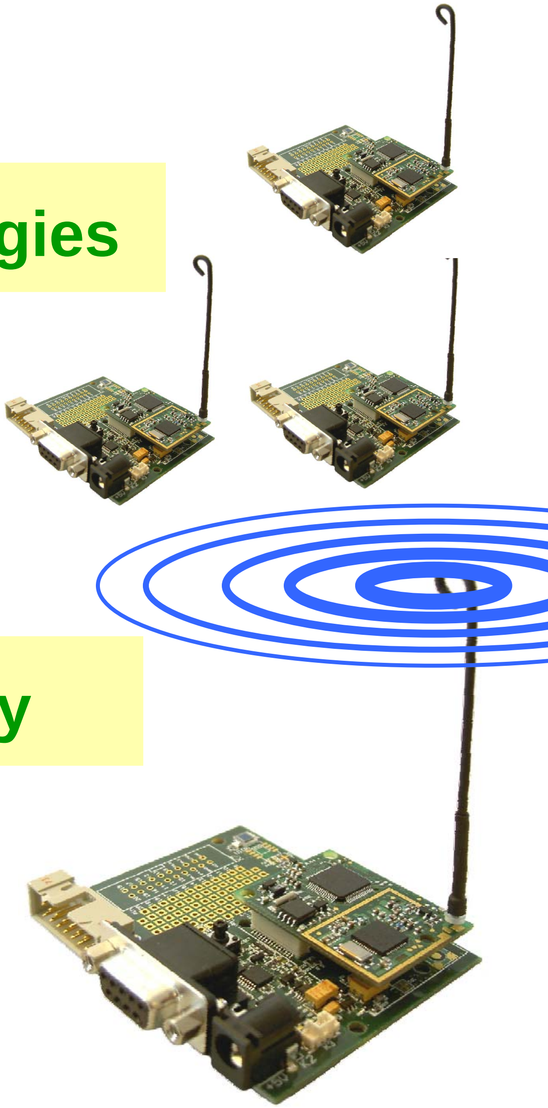

Physical Model

- Let us look at the signal-to-noise-plus-interference (SINR) ratio!
- Message arrives if SINR is larger than β at receiver

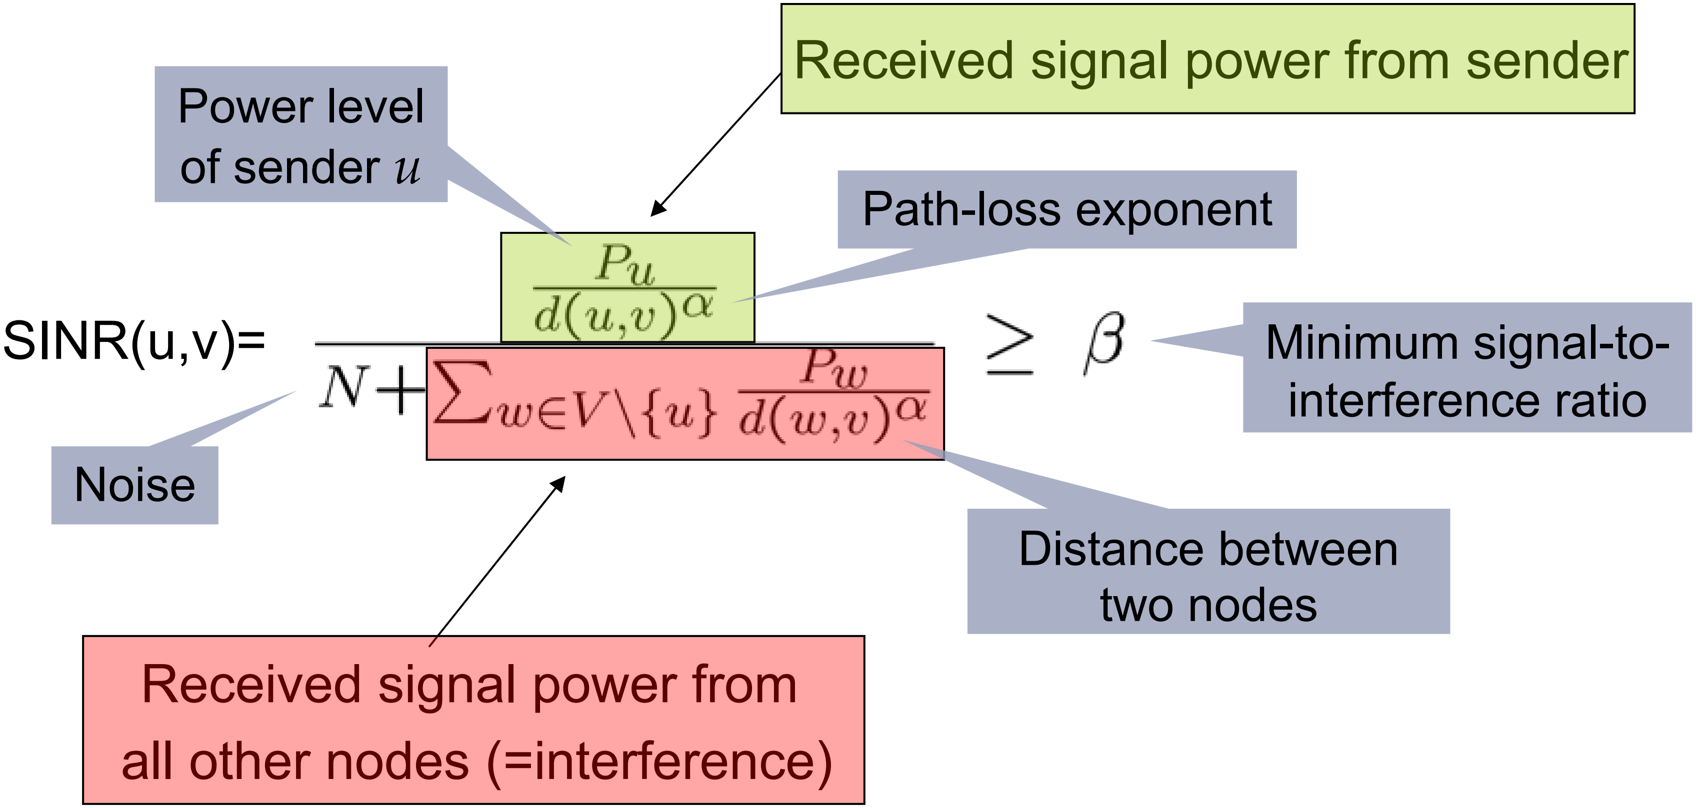

$$\text{SINR(u,v)=}\ \frac{\frac{P_u}{d(u,v)^{\alpha}}}{N+\sum_{w\in V\setminus\{u\}}\frac{P_w}{d(w,v)^{\alpha}}} \geq \beta$$

Received signal power from sender

Power level of sender u

Path-loss exponent

Minimum signal-to-interference ratio

Noise

Distance between two nodes

Received signal power from all other nodes (=interference)

Power Control

Ideal Case:

Select power level to meet target SINR exactly

Why?

- SINR < β ⇔ transmission not successful

need retransmission,
not energy-efficient

- SINR > β ⇔ power too high

battery drain

interference

How can we find a GOOD power level
in a distributed setting?

Simple Interference Games

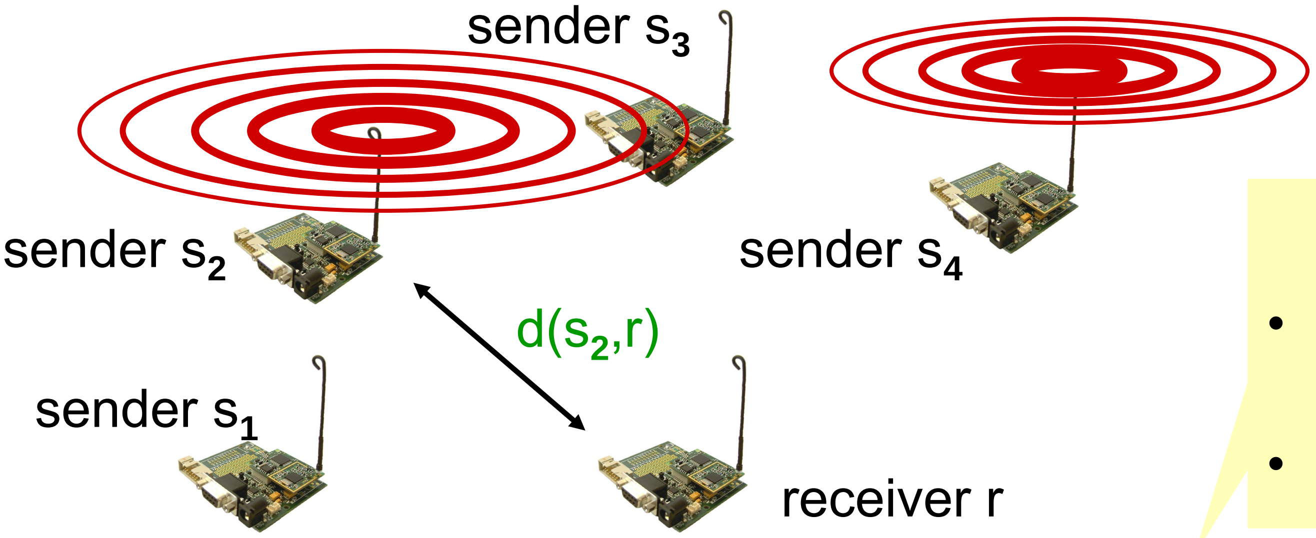

Performance ?

Convergence

- Existence, uniqueness of NE
- Conditions

Interaction:

- s_i chooses power p_i
- transmission i successful if SINR(s_i ,r) > β

Payoff:

- Successful: v – p_i NOT successful: – p_i

Efficiency

- Social Welfare
- Pareto Efficiency
- Price of Anarchy

Fairness

- Ressources/utility shared equally?

Simple Interference Games

Example: 2 players, 3 power levels

Power	0	1	2
0	0,0	0,v-1	0,v-2
1	v-1,0	-1,-1	-1,v-2
2	v-2,0	v-2,-1	-2,-2

Performance ?

Convergence

- No pure NE for # power levels > 2
- Cycle? (0, 1), (2, 1), (2, 0), (1, 0), (1, 2), (0, 2), (0, 1)
- Mixed NE $p_0=p_1=1/v$
- Corr. E?

Fairness

Ratio: min(payoff)/max(payoff)

- Cycle fairness =1
- Mixed NE fairness = 1

Social Welfare

- OPT: v-1
- Cycle: v-2
- Mixed NE: 0
- Correlated E: $v-+v/(v^2-2)$

Simple Interference Games, more general results

2 players, k power levels

- fair mixed NE, social welfare 0
- fair corr. E, social welfare max(0,v-2k+1), fair
- odd k: fair corr. E, social welfare v-k
  fair cycle, social welfare v-k
- even k: unfair mixed NE, social welfare v-k
  2 cycles, social welfare v-k,
  fairness ((2k-2)v-2k^2+2)/((2k+6)v-k^2-1)

n players, k power levels

- fair mixed NE, social welfare 0
- even k: unfair mixed NE, social welfare v-k

[„Interference Games in Wireless Networks“, Auletta et al., WINE08]

More Models

So far:
very simple model, more realistic assumptions?

1

Wireless data QoS → different from voice

- Different utility functions:

Utility vs Power (constant interference)

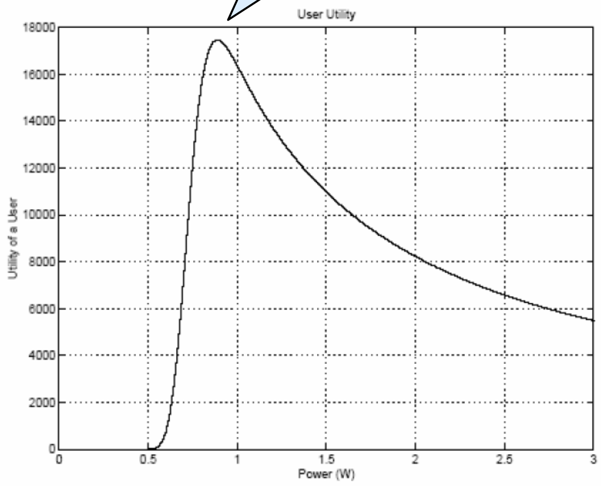

Figure 1.2 Utility function of a data user for fixed interference

2

Utility $U_i(p_i) = E_i R_i f(SINR(p_i)/ p_i$

battery energy

transmission rate

Efficiency function

Exact Potential Games Properties

Properties:

1. Sum of coordination and dummy games
2. Pure NE: local maximum of potential function
3. Best-response dynamics converge to NE
4. For continuous utility functions

$$\frac{\partial P}{\partial a_i}=\frac{\partial u_i}{\partial a_i}\ , \qquad \frac{\partial^2 P}{\partial a_i \partial a_j}=\frac{\partial^2 u_i}{\partial a_i \partial a_j}=\frac{\partial^2 u_j}{\partial a_i \partial a_j}$$

for every *i*, *j* ∈ *N*

5. Efficiency guarantees

3,3	0,2
2,0	2,2

Coordination

\+

4,3	1,3
4,2	1,2

Dummy

=

7,6	1,5
6,2	3,4

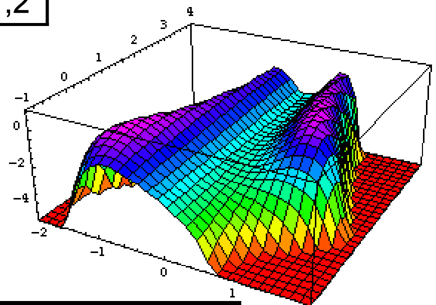

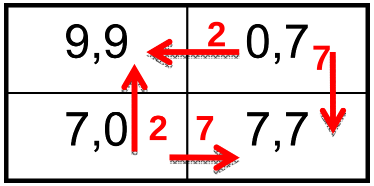

9,9	0,7
7,0	7,7

$$V(a)=\sum_{i\in N}\int_0^1 \frac{\partial u_i}{\partial a}(x(t))\,x_i^{'}(t)\,dt$$

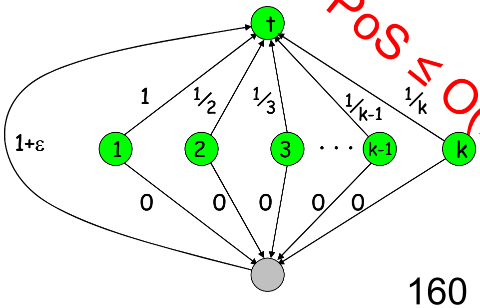

PoS ≤ O(log k)

Potential Game Approach

How do nodes know best response?

$u_i(p_i) = 1-(\beta-p_i/d_i^{\alpha} +\sum_{j!=i} p_i/d_j^{\alpha}\}+N)^2$

verify

$$\frac{\partial^2 u_i}{\partial p_i \partial p_j} = \frac{\partial^2 u_j}{\partial p_j \partial p_i}$$

converges to pure strategy NE (best response)

$P(p_1,\ldots,p_n) = 2 \sum_j(\beta p_i/d_i^{\alpha} - p_i^2/d_i^{2\alpha})+2\sum_j \sum_{j=i+1} p_i\, p_j/d_i^{\alpha} /d_j^{\alpha}$

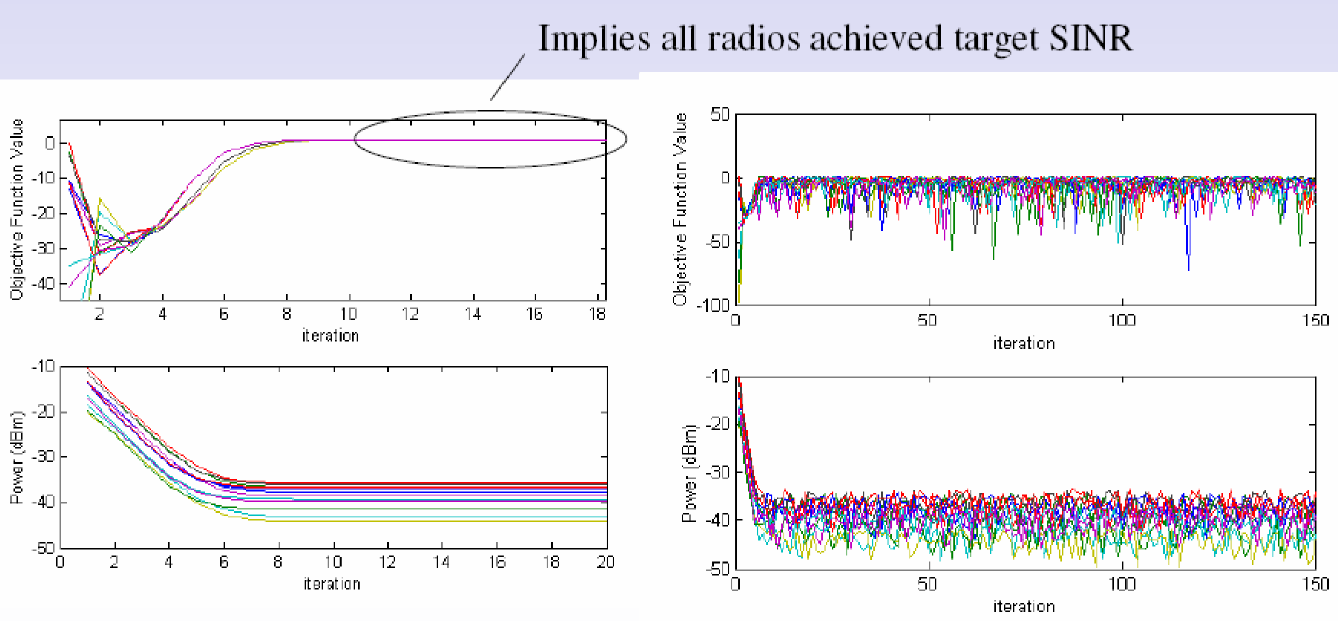

Noiseless Best Response

Noisy Best Response

[Neel, Reed, Gilles, "Convergence of cognitive radio networks,". WCNC04]

NP-hardness Locality Game

Complexity of Nash Equilibrium

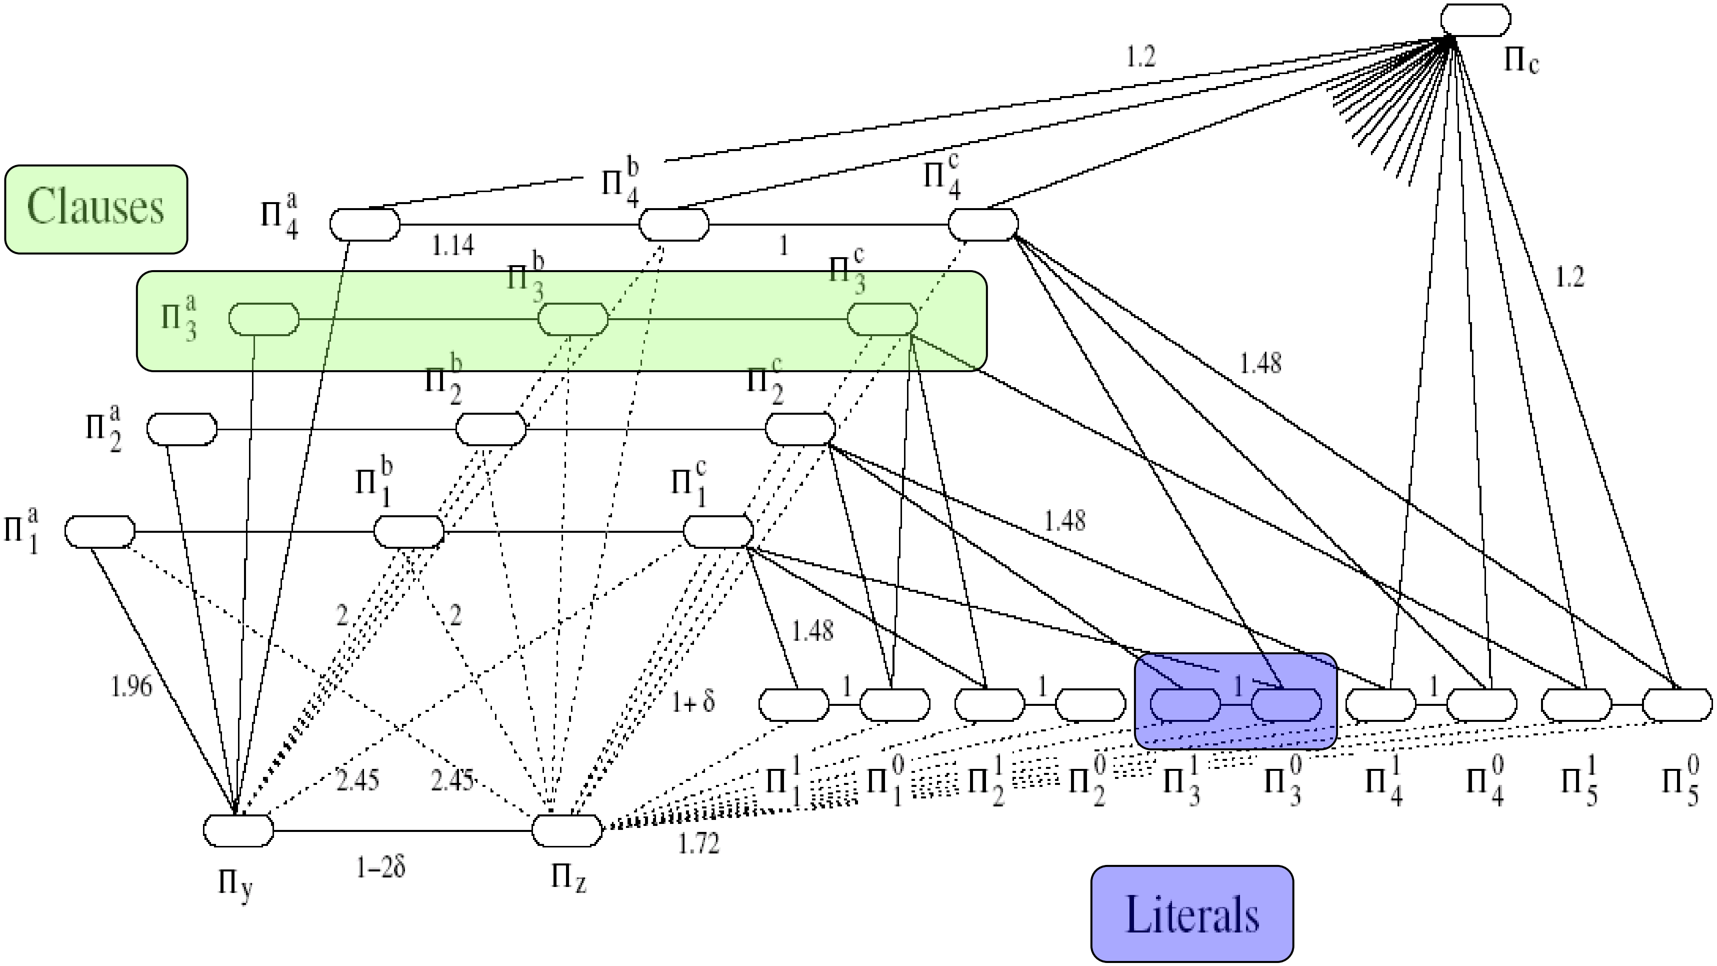

Complexity of Nash Equilibrium

- It can be shown: In any Nash equilibrium, these links must exist…

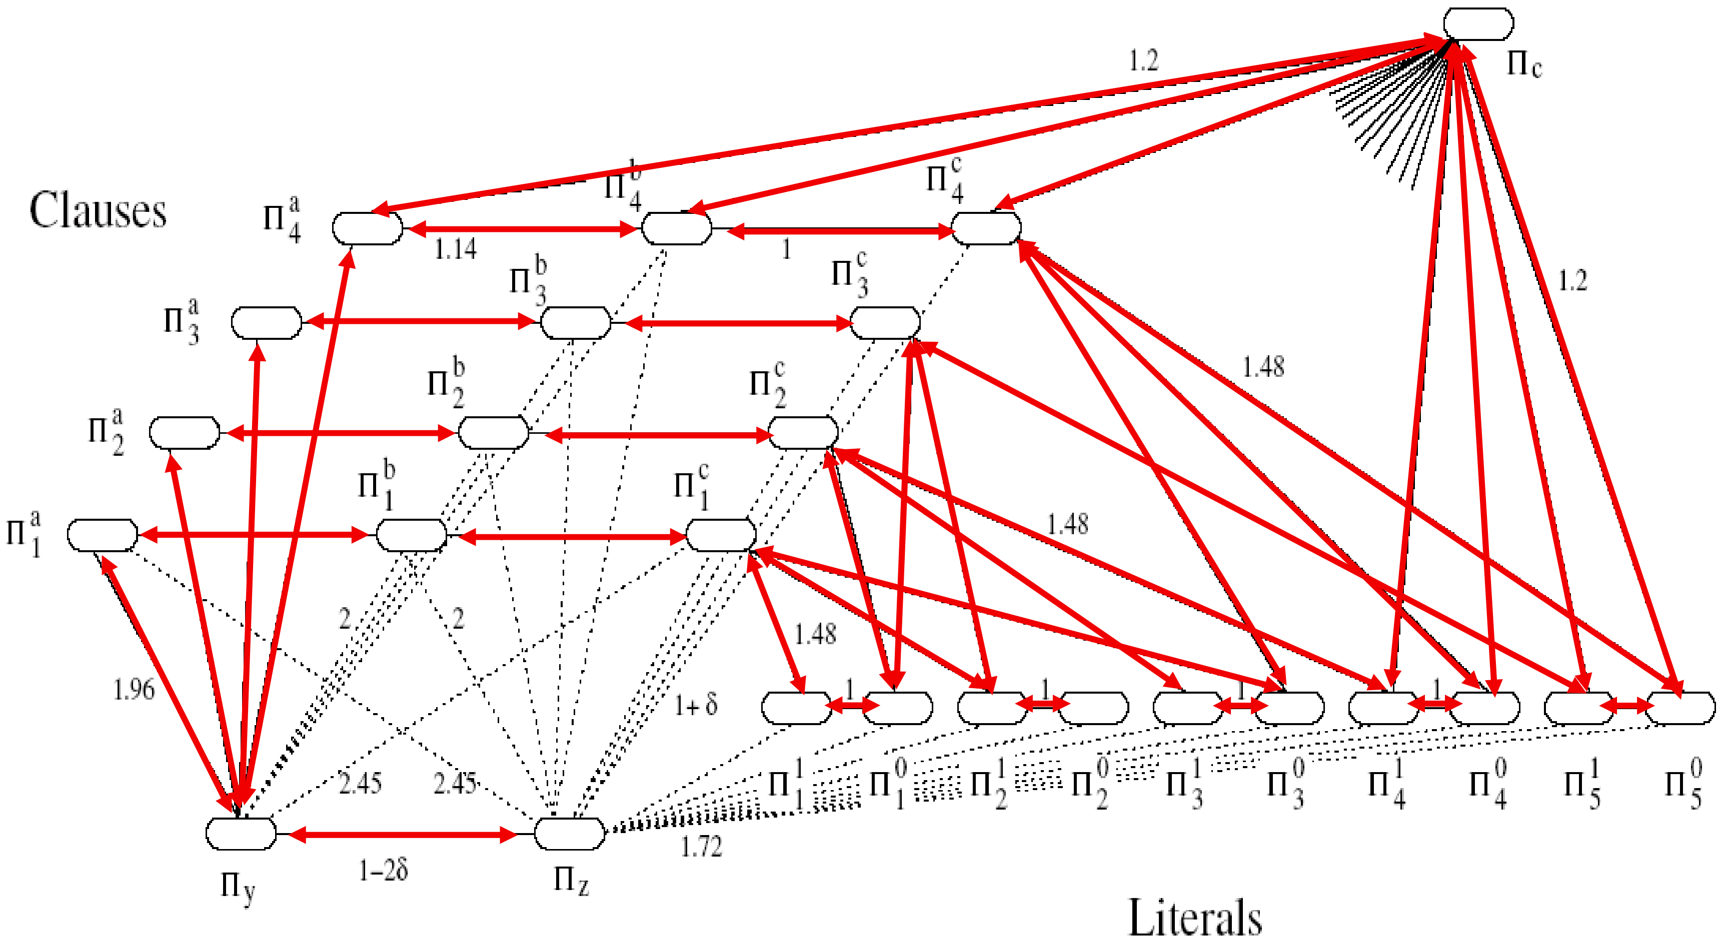

Complexity of Nash Equilibrium

Special SAT: Each variable in at most 3 clauses!

- Additionally, Π_z has exactly one link to one literal of each variable!
  - Defines the “assignment” of the variables for the formula.
  - If it’s the one appearing in the clause, this clause is stable!

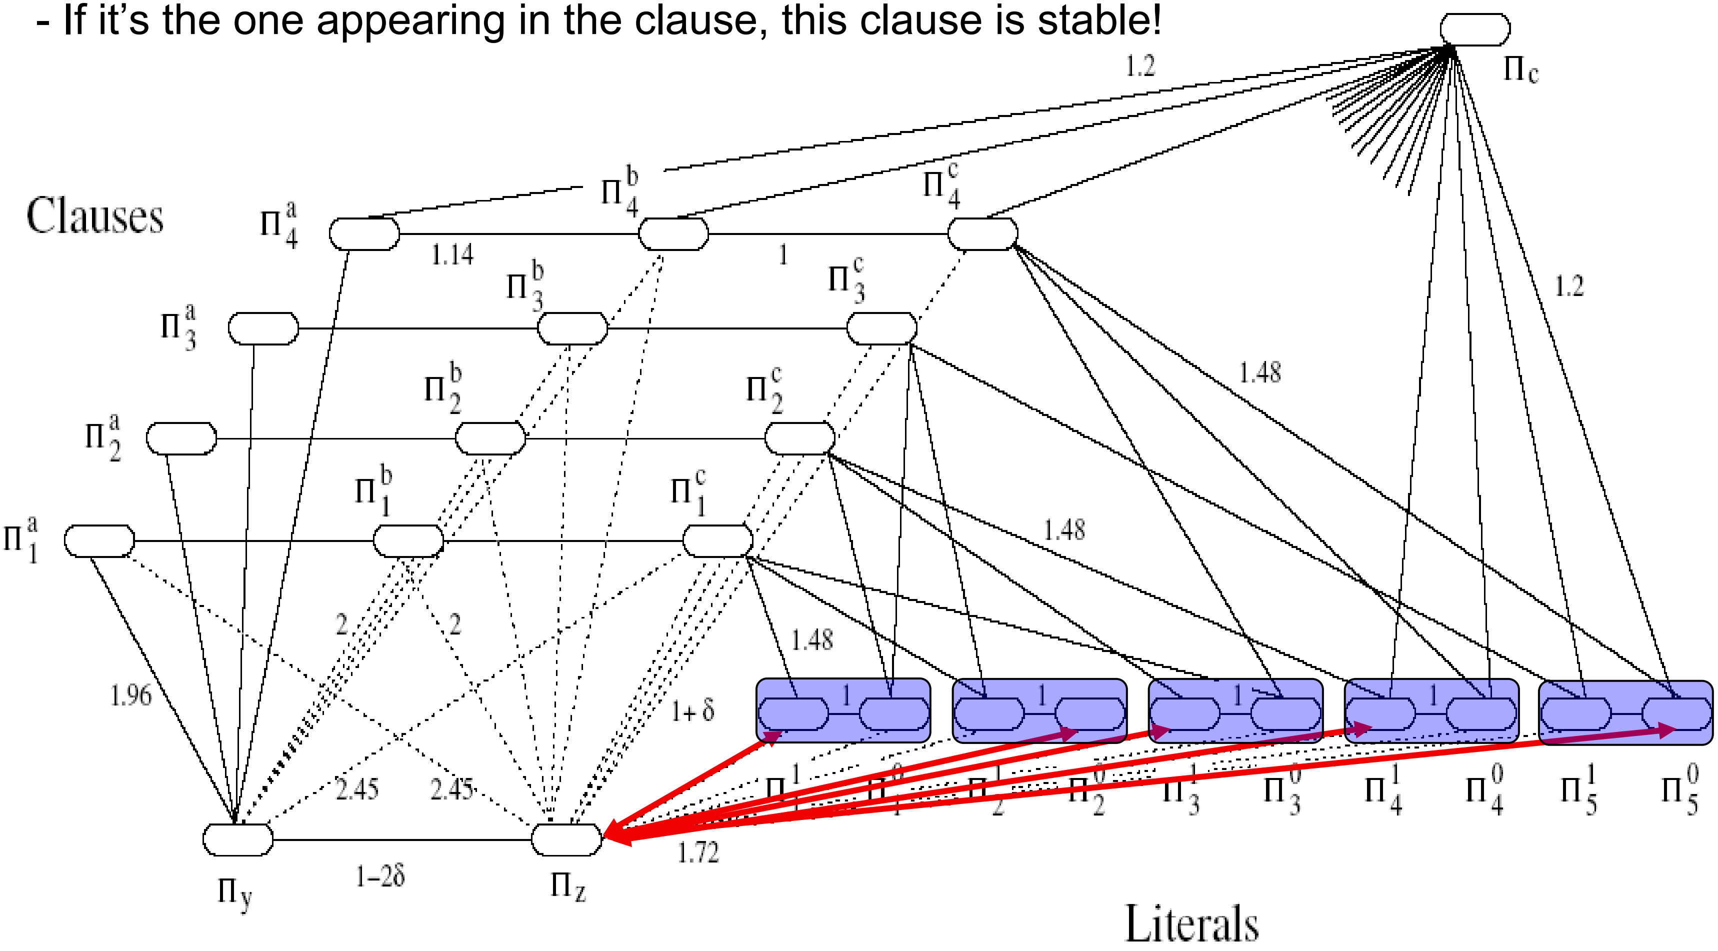

Complexity of Nash Equilibrium

- Such a subgraph (Π_y, Π_z, Clause) does not converge by itself…

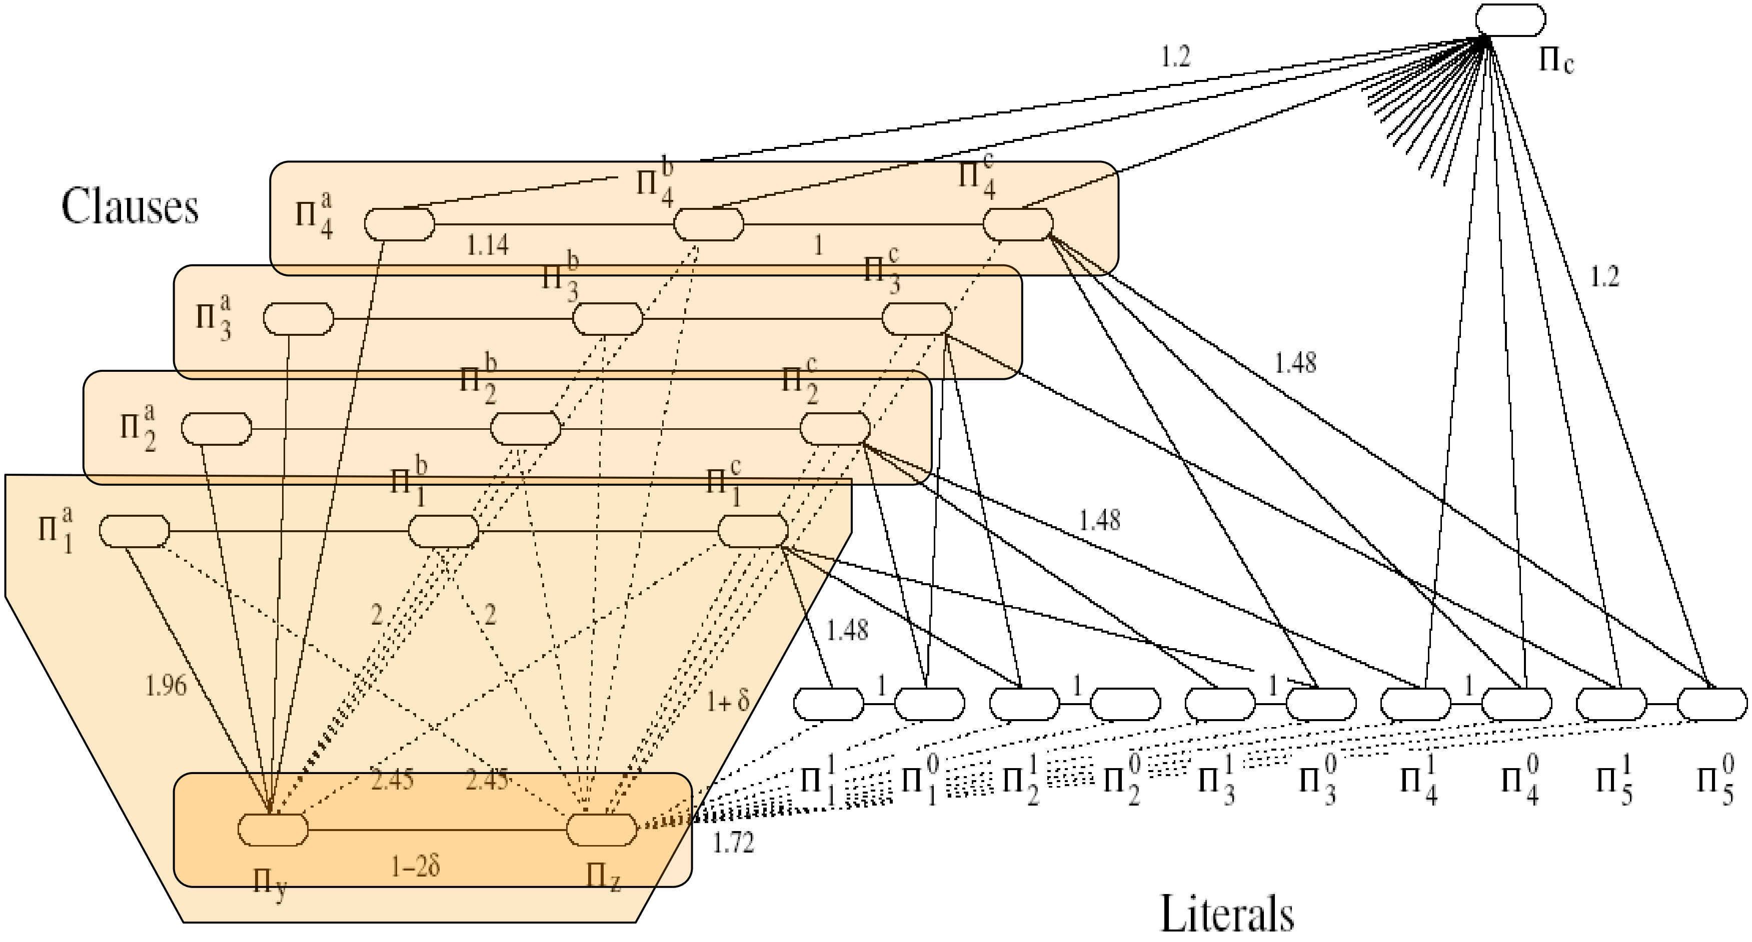

Complexity of Nash Equilibrium

- In NE, each node-set Π^c is connected to those literals that are in the clause (not to other!)

→ if Π_z has link to not(x1),
there is a “short-cut” to such clause-nodes, and C_2 is stable

→ But not to other clauses (e.g., C_1 = x_1 v x_2 v not(x_3)): literal x_1 does not appear in C_1…

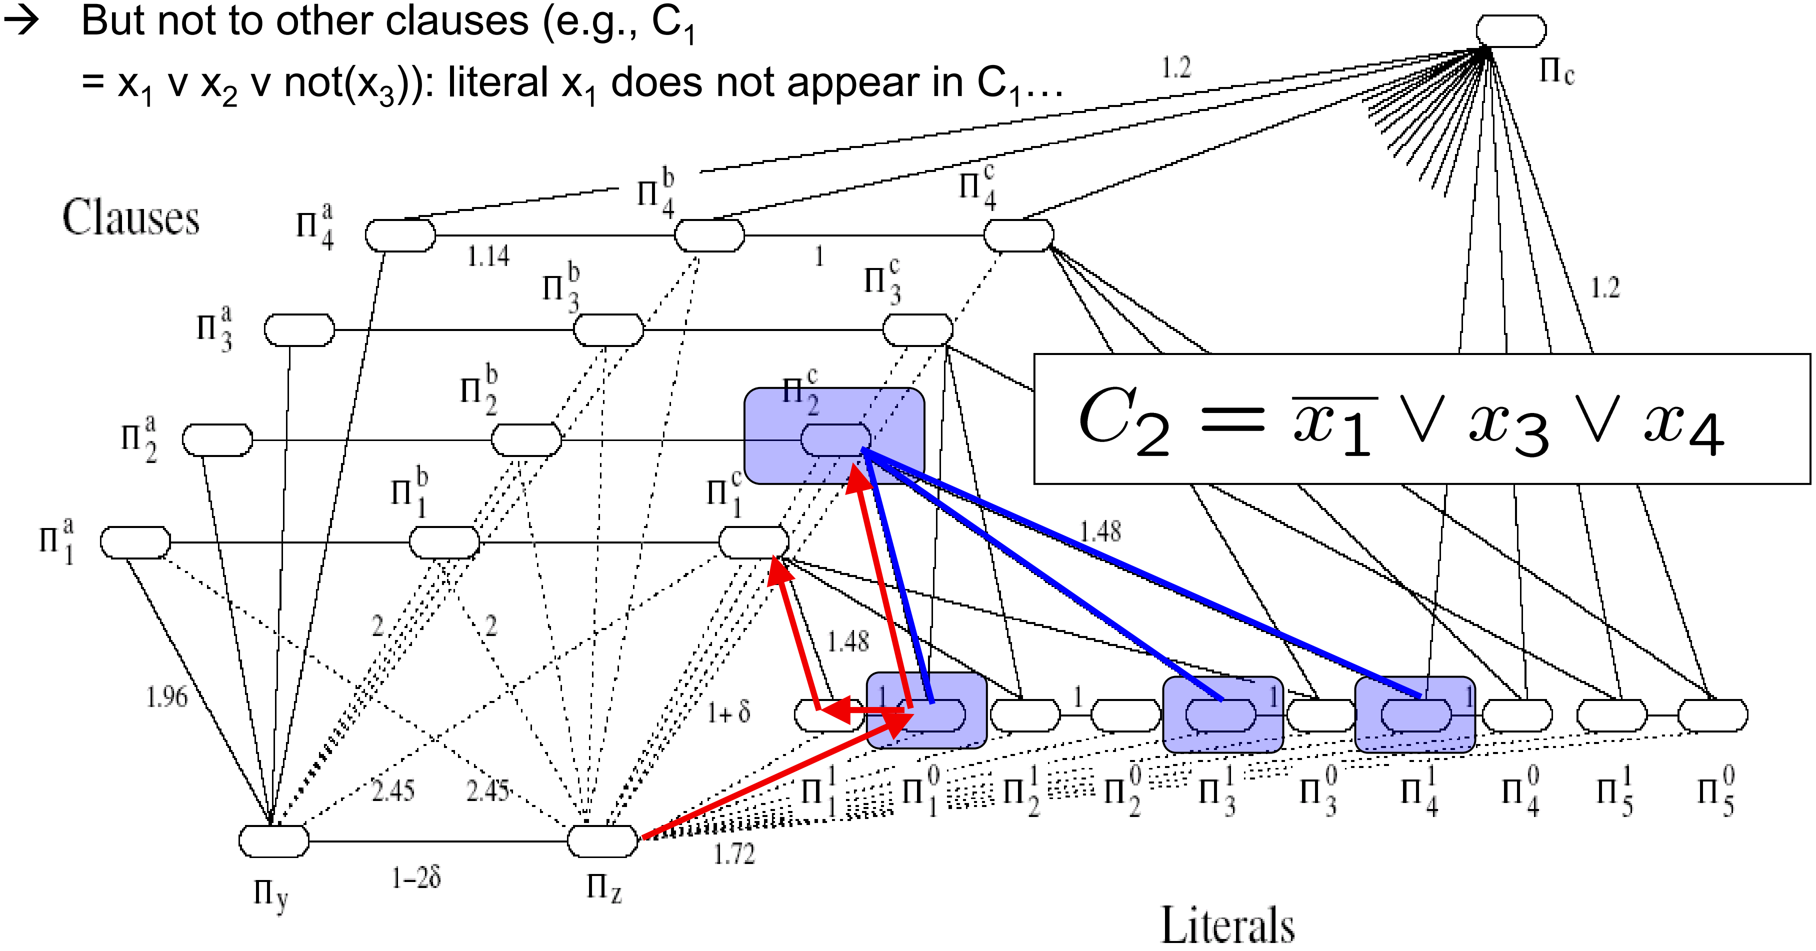

$$C_2 = \overline{x_1} \vee x_3 \vee x_4$$

Complexity of Nash Equilibrium

- A clause to which Π_z has a “short-cut” via a literal in this clause becomes stable! (Nash eq.)

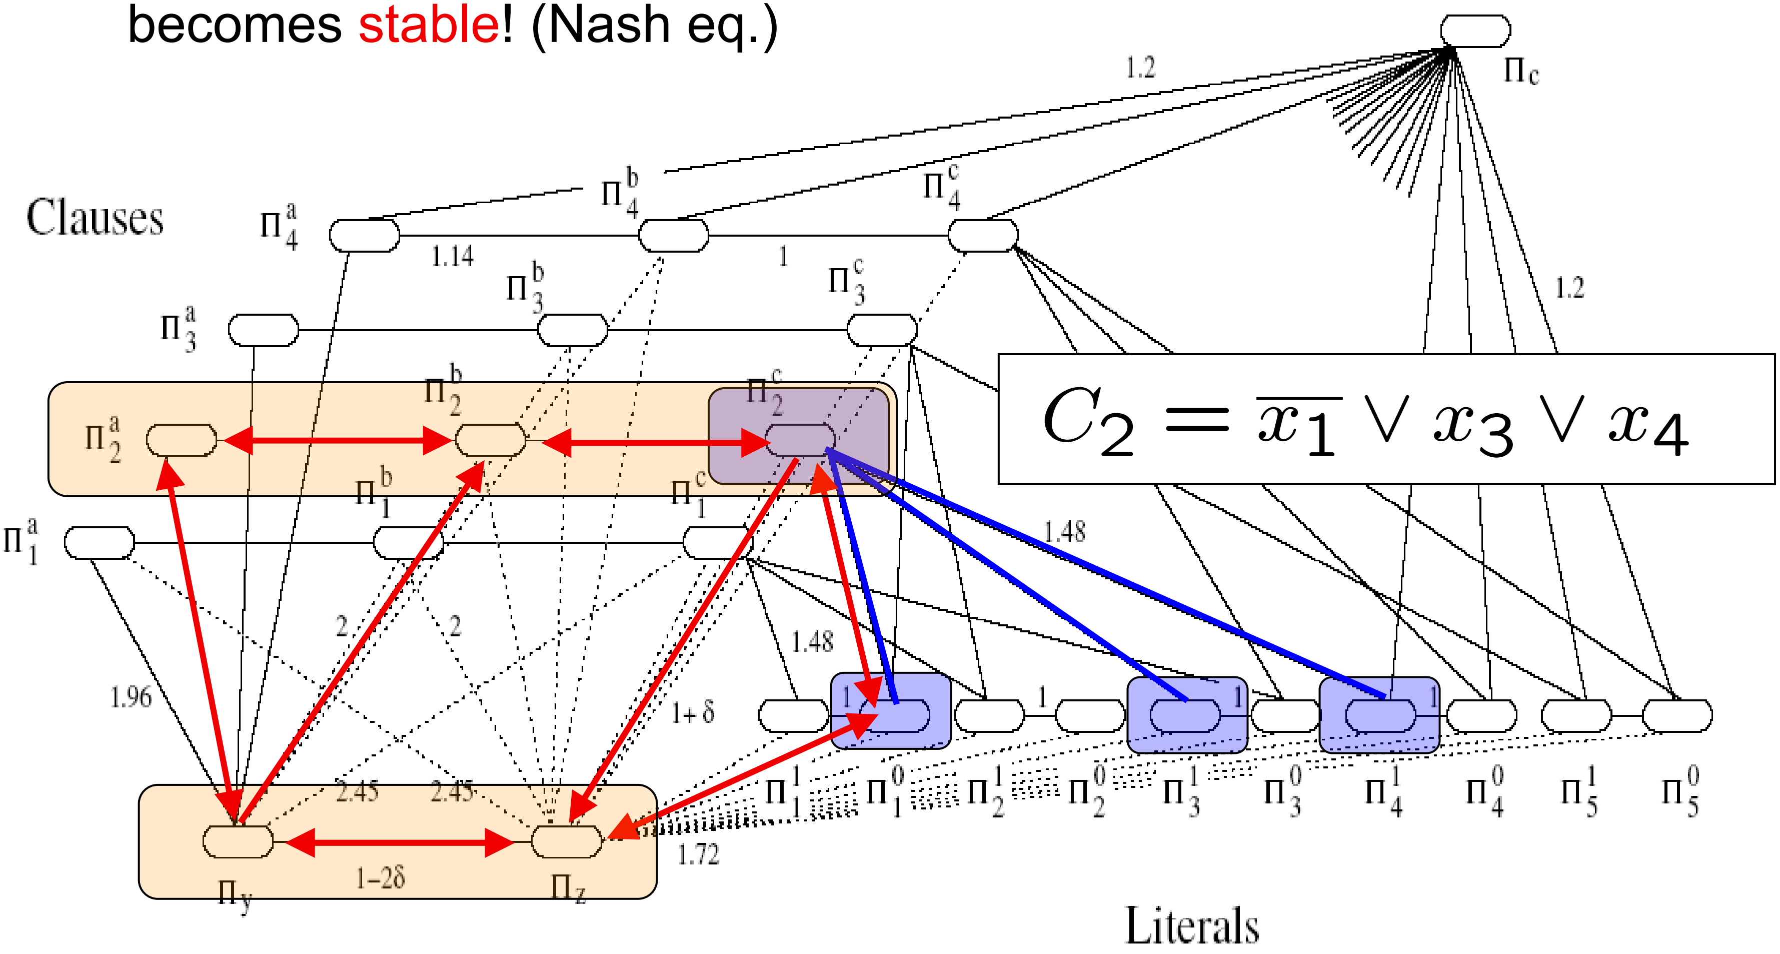

$$C_2 = \overline{x_1} \vee x_3 \vee x_4$$

Complexity of Nash Equilibrium

- If there is no such "short-cut" to a clause, the clause remains instable!
- Lemma: not satisfiable -> instable / no pure NE

  (contradiction over NE's properties)

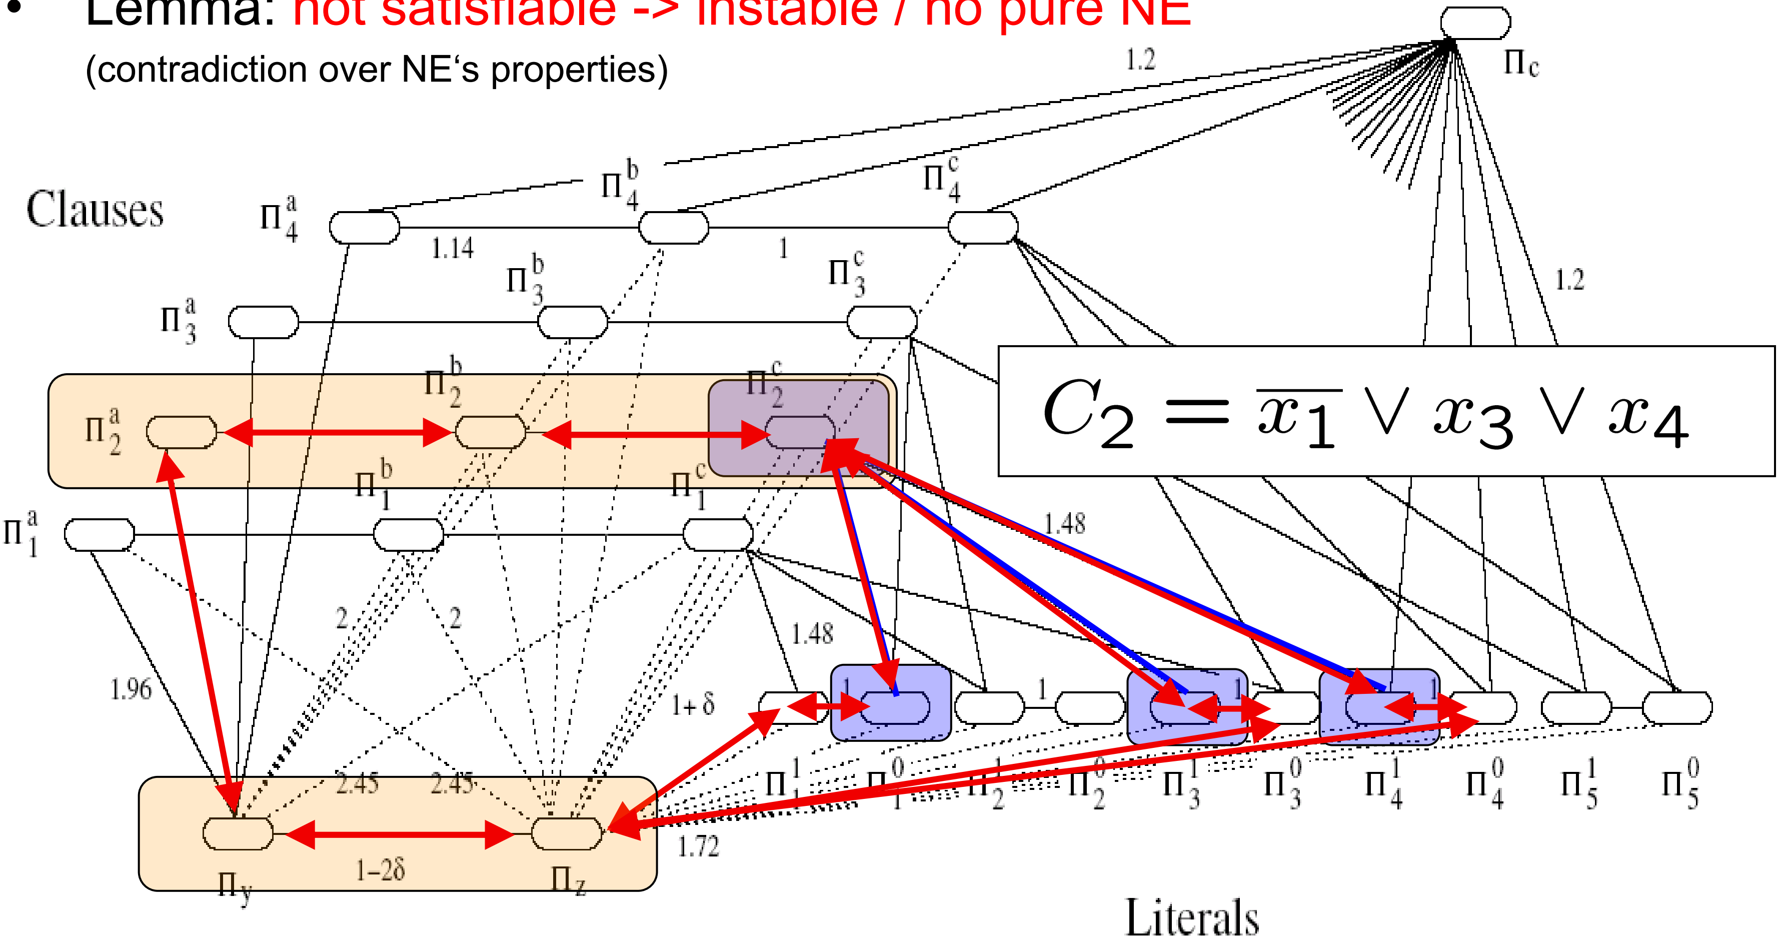

$$C_2 = \overline{x_1} \vee x_3 \vee x_4$$

Complexity of Nash Equilibrium

- Example: satisfiable assignment -> all clauses stable -> pure NE

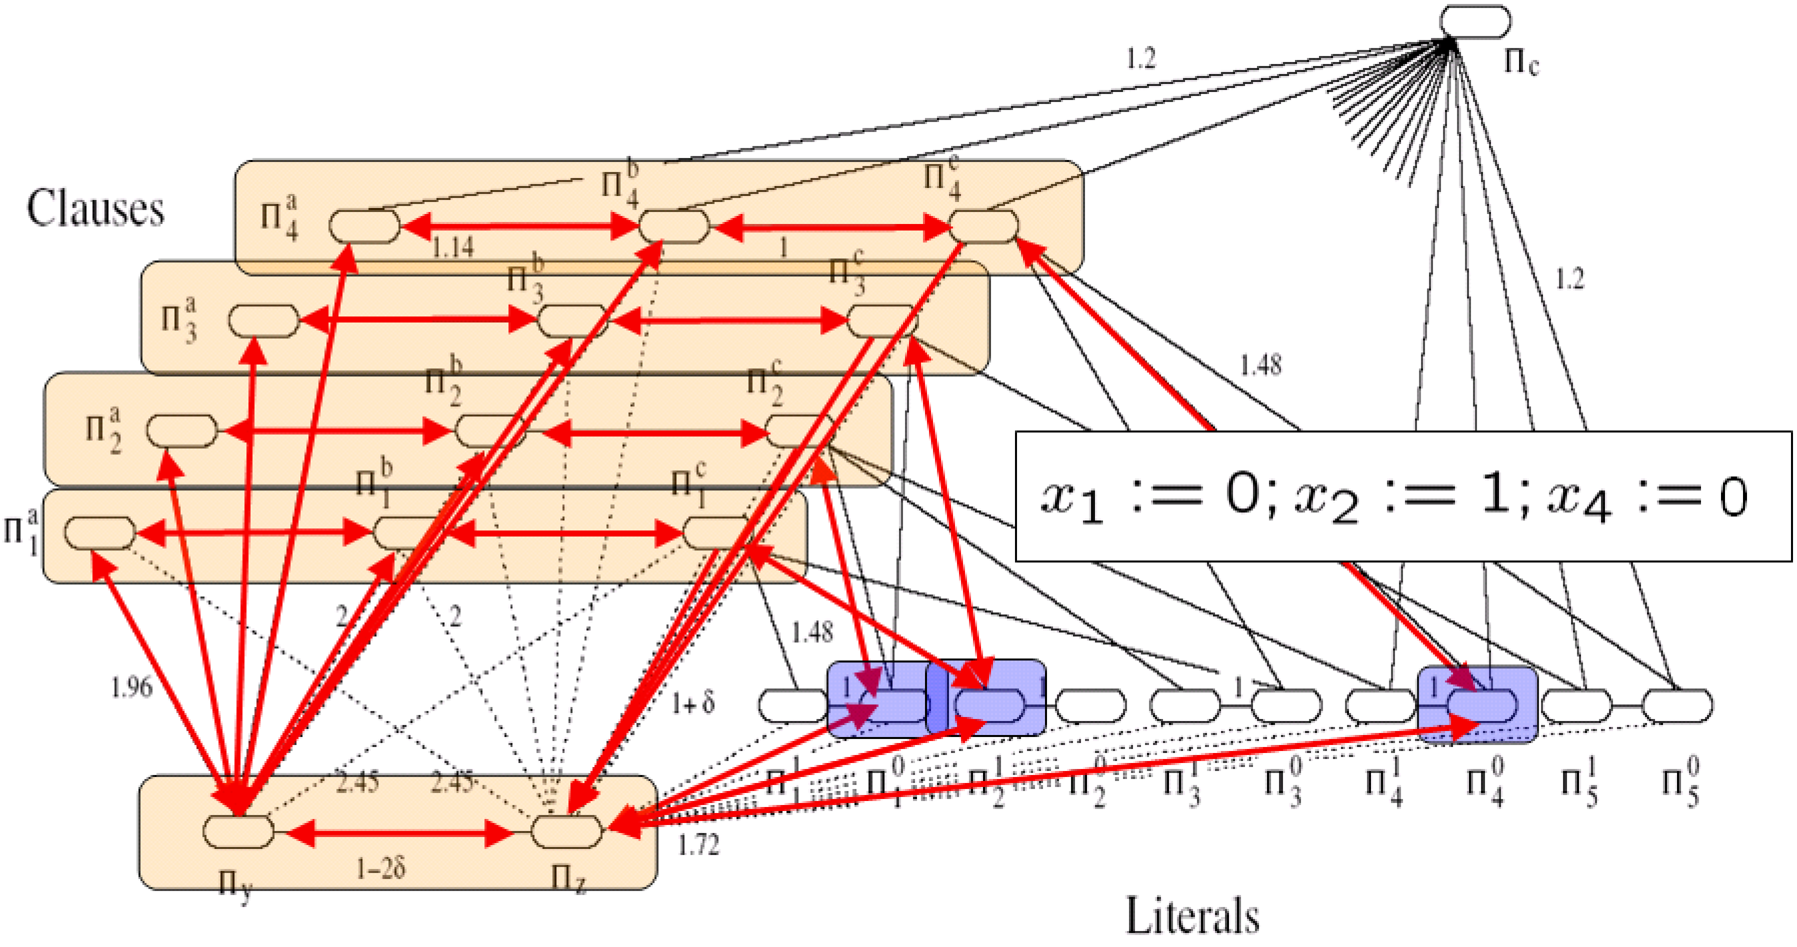

$x_1 := 0; x_2 := 1; x_4 := 0$