On the Impact of Malicious Players in Distributed Systems

Stefan Schmid

When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game

Thomas Moscibroda
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
moscitho@tik.ee.ethz.ch

Stefan Schmid
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
schmiste@tik.ee.ethz.ch

Roger Wattenhofer
Computer Engineering and Networks Laboratory
ETH Zurich
8092 Zurich, Switzerland
wattenhofer@tik.ee.ethz.ch

ABSTRACT

Over the last years, game theory has provided great insights into the behavior of distributed systems by modeling the players as utility-maximizing agents. In particular, it has been shown that selfishness causes many systems to perform in a globally suboptimal fashion. Such systems are said to have a large Price of Anarchy. In this paper, we extend this active field of research by allowing some players to be malicious or Byzantine rather than selfish. We ask: What is the impact of Byzantine players on the system's efficiency compared to purely selfish environments or compared to the social optimum? In particular, we introduce the Price of Malice which captures this efficiency degradation. As an example, we analyze the Price of Malice of a game which models the containment of the spread of viruses. In this game, each node can choose whether or not to install anti-virus software. Then, a virus starts from a random node and iteratively infects all neighboring nodes which are not inoculated. We establish various results about this game. For instance, we quantify how much the presence of Byzantine players can deteriorate or—in case of highly risk-averse selfish players—improve the social welfare of the distributed system.

Categories and Subject Descriptors

DYNAMO'07:

1st Workshop on Dynamic Networks

(Salerno, Italy, May 2007)

Talk is based on PODC'06 paper, joint work with
Dr. Thomas Moscibroda, MS Research
Prof. Dr. R. Wattenhofer, ETH Zurich

Distributed Computing Group

ETH Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Systems…

Wireless Mesh Networks

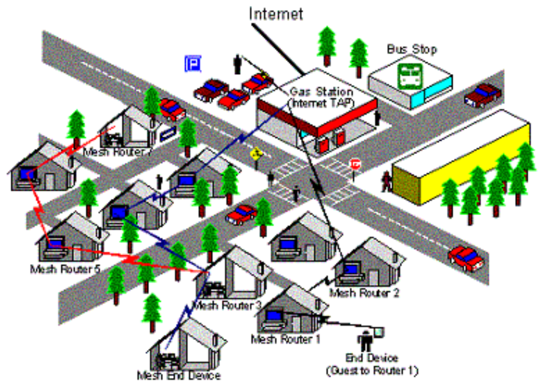

Internet

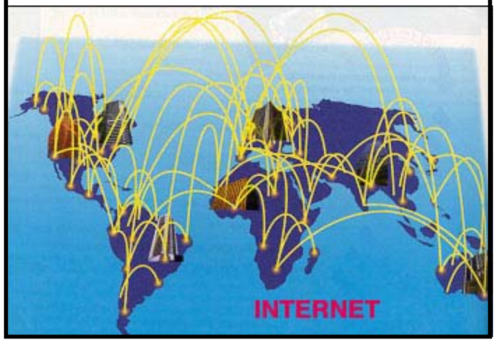

P2P Networks

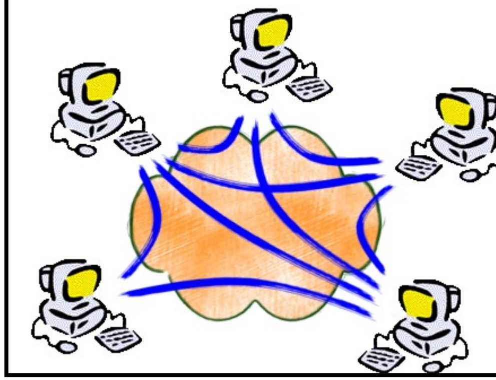

(Electronic) Markets, Society,…?

Wireless Sensor Networks

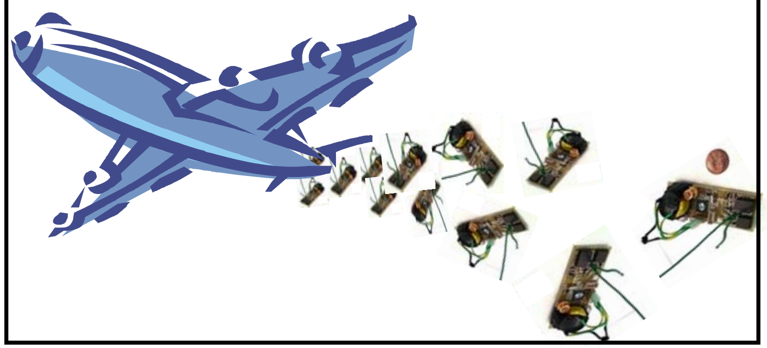

Modeling Participants of Distributed Systems

- One possibility to model a distributed system:

  all participants are benevolent!

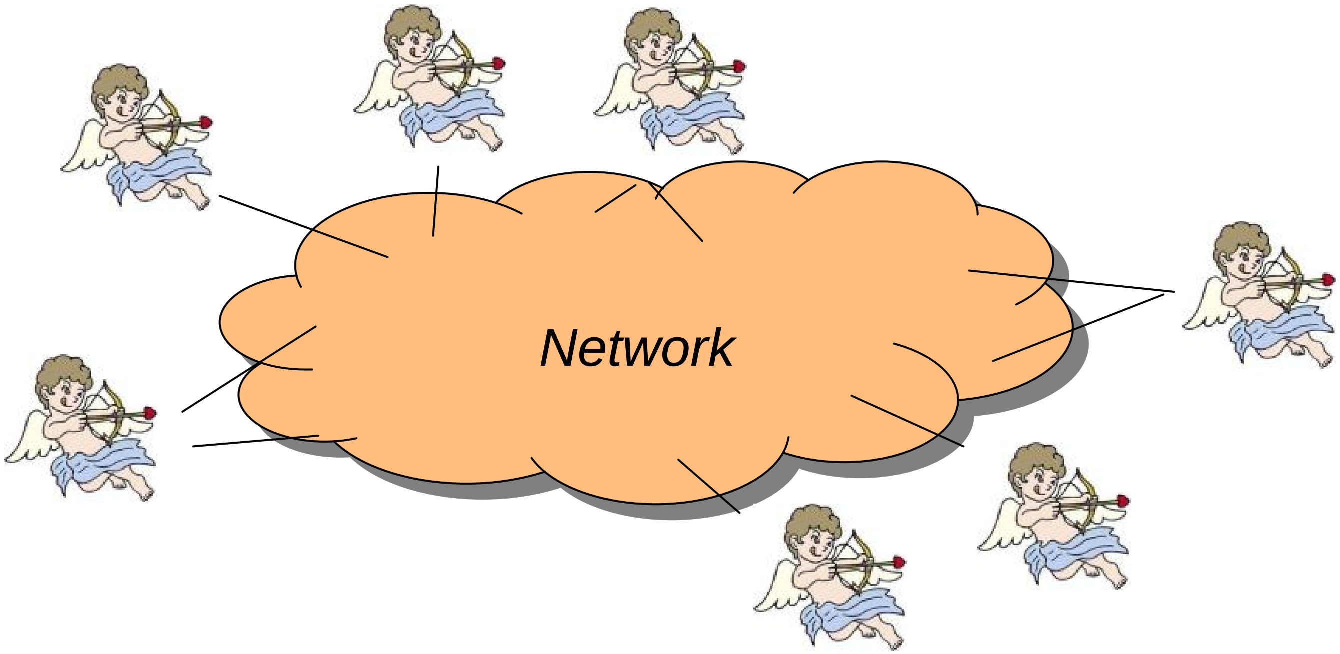

Selfishness in Networks

- Alternative: Model all participants as selfish

  → e.g. impact on congestion, or on p2p topologies, etc.

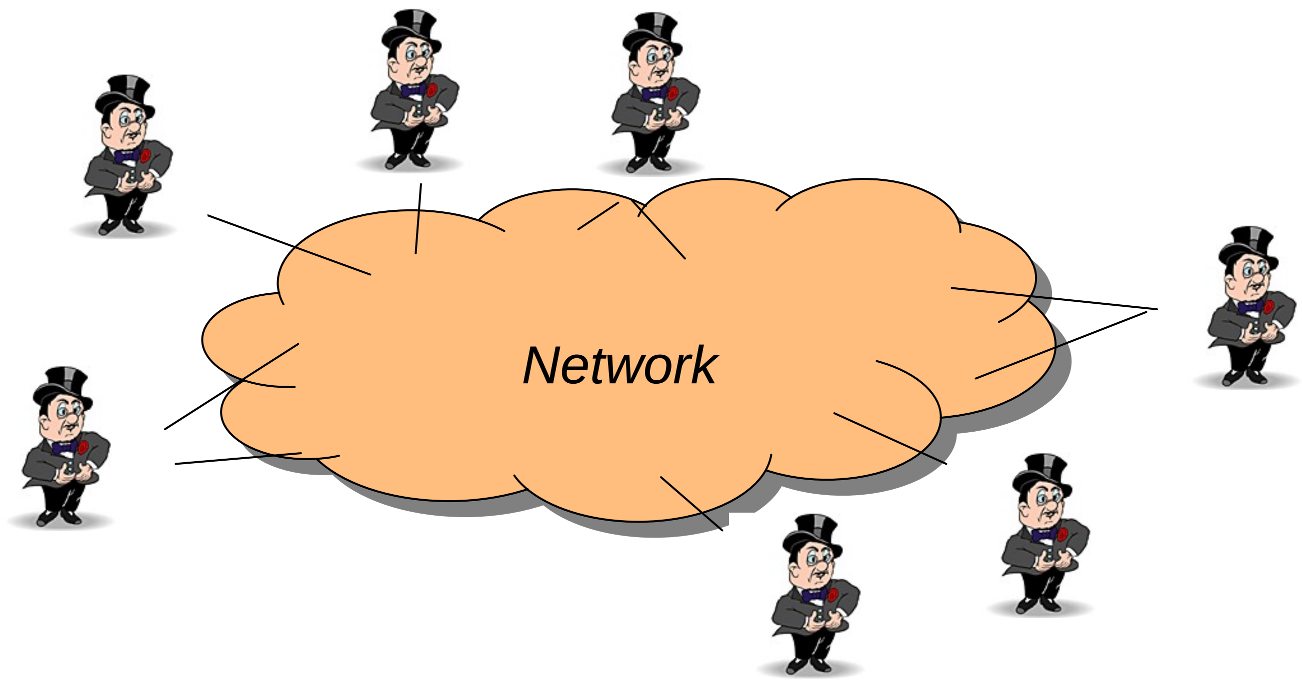

Classic game theory: What is the impact of selfishness on network performance…? (=> Notion of price of anarchy, etc.)

When Selfish meets Evil…

- But selfishness is not the only challenge in distributed systems!

  → Malicious attacks on systems consisting of selfish agents

Hackers, Polluters, Viruses, DOS attacks

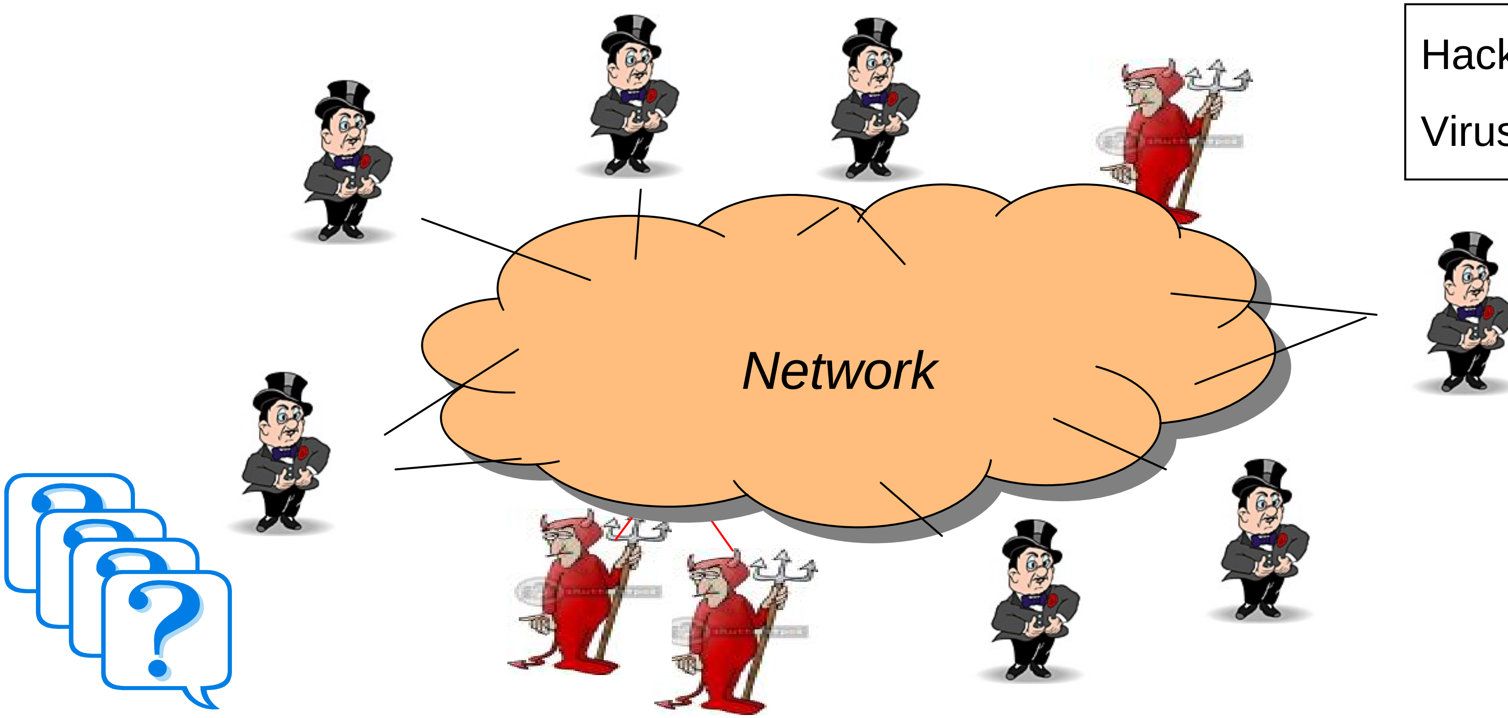

What is the impact of malicious players on selfish systems…?

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!

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

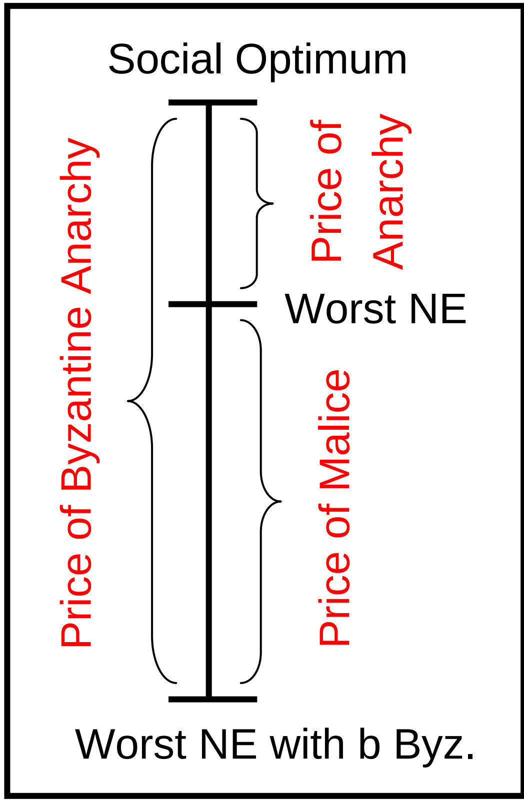

The Price of Malice captures the degradation of a system consisting of selfish agents due to malicious participants!

Remark on “Byzantine Game Theory”

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

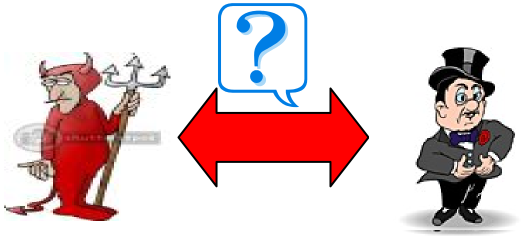

- Theoretically, malicious players are also selfish... .... just with a different utility function!

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!

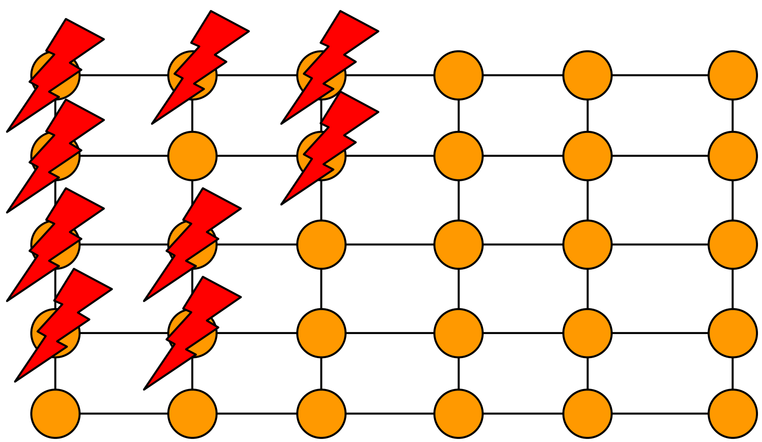

Related Work:

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

- General Graphs
- No malicious players

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

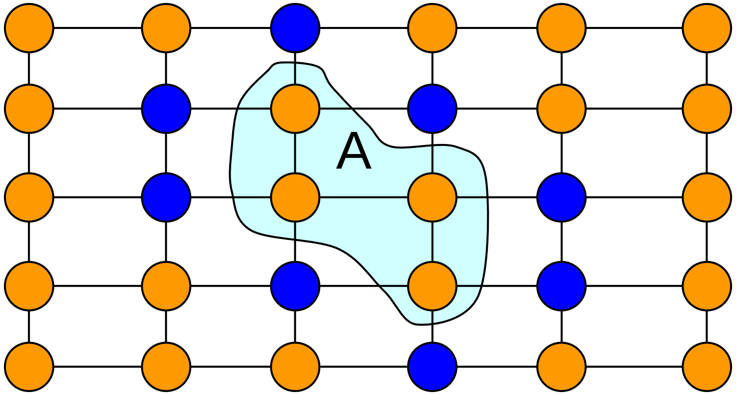

Total infection cost: $$Cost_{inf} = \frac{L}{n}\sum_i k_i^2$$ ← k_i: insecure nodes in the ith component

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

  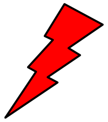

  (worst-case: malicious player cannot be trusted and may say s.th. but do s.th. else…)
- What is the Price of Malice…?

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

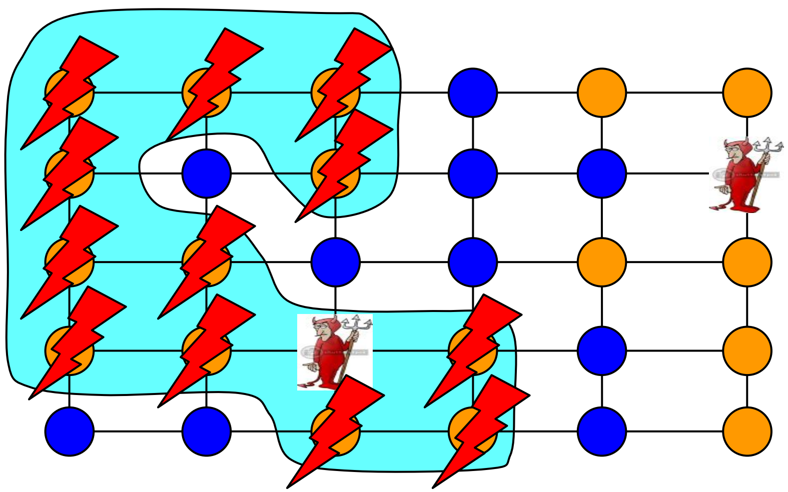

Distinguish between:
- Oblivious model
- Non-oblivious model
  - Risk-averse

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)

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

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

l_i: #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!)

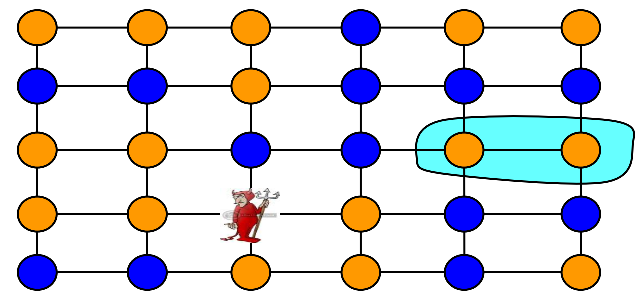

- On the other hand: a component i with $b_i>0$ malicious nodes: $\sum_i b_i = b$

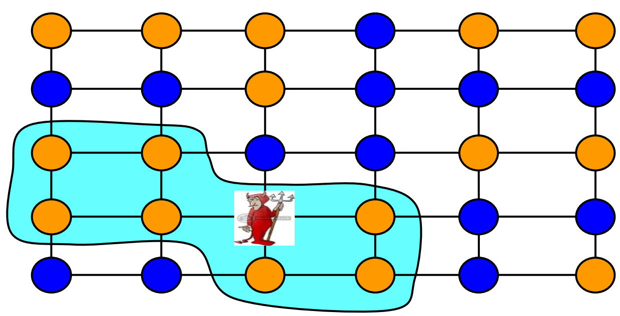

- In any non-Byz NE, the size of an attack component is at most n/L, so

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

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

for b<L/2

(for other case see paper)

- Hence, the Price of Byzantine Anarchy is at most

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

- The Price of Malice is at most

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

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

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)

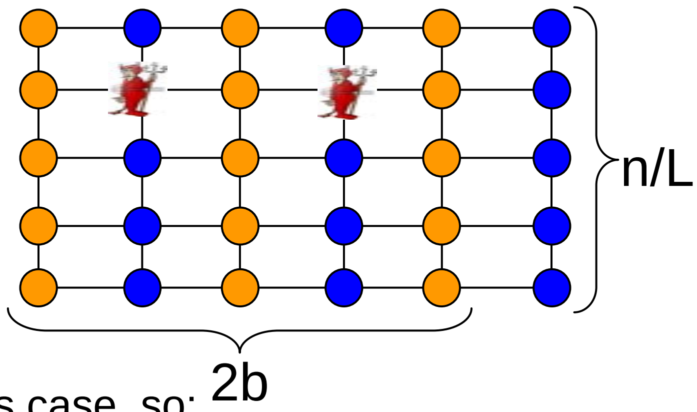

→ this scenario where every second column is

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 ≥ 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)

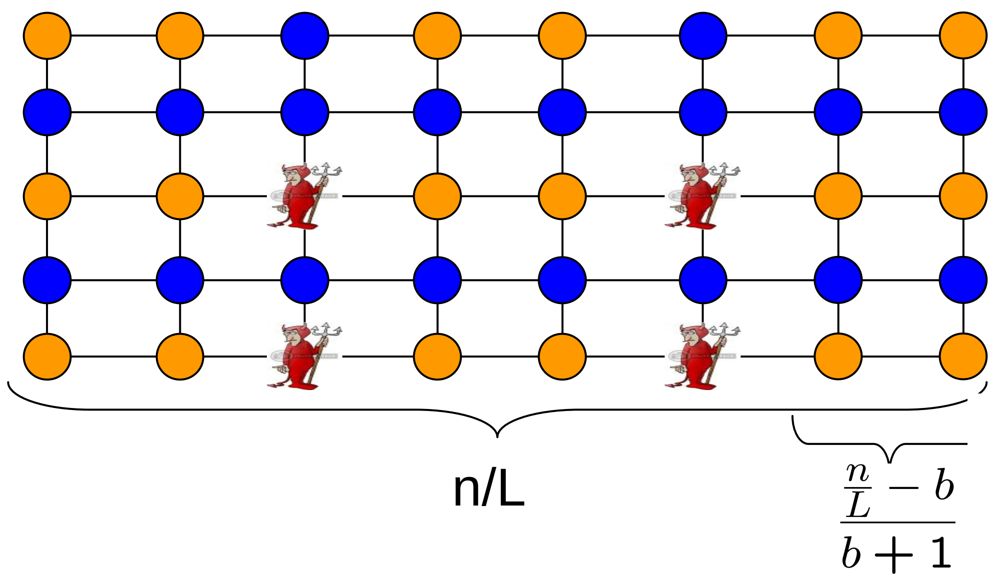

This constitutes a Byzantine Nash equilibrium!

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)

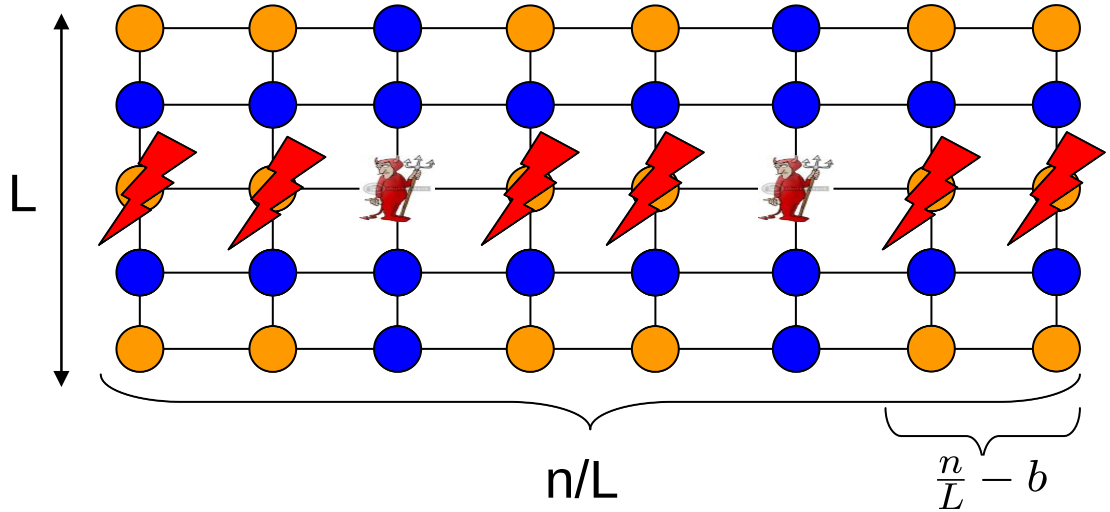

$$\frac{\frac{n}{L}-b}{b+1}$$

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:

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

Infection cost of selfish nodes in other rows…

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

↑ number of insure 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:

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

- Price of Malice is:

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

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

$\rightarrow$ 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

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

⇓

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

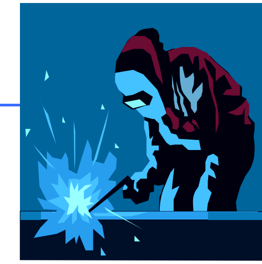

- What is the implication in practical networking…?
- If Price of Anarchy is high

  → 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

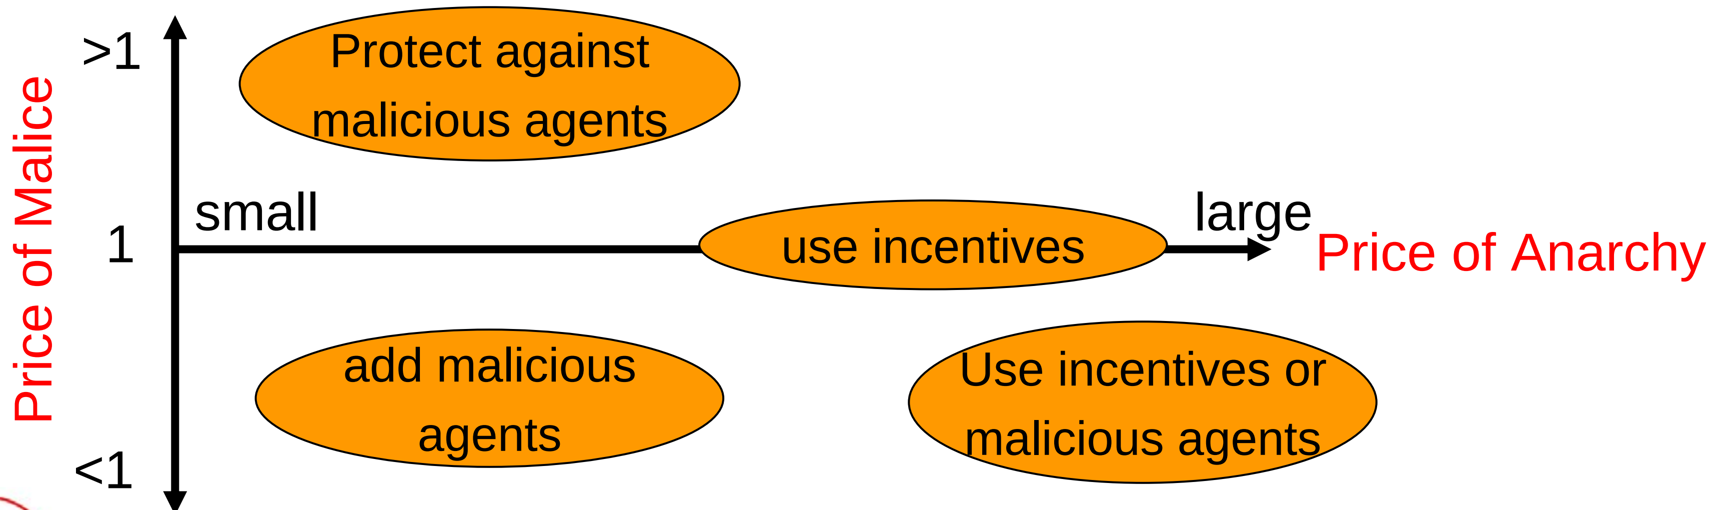

Reasoning about the Fear Factor

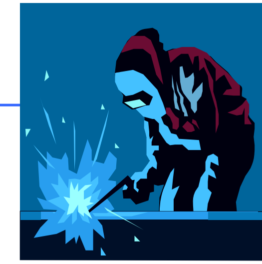

- What is the implication in practical networking…?
- Fear-Factor can improve network performance of selfish systems!

  (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)

Future Work and Open Questions

THANK YOU !

- Plenty of open questions and future work!
- <u>Virus Inoculation Game</u>

  → The Price of Malice in more realistic network graphs

  → High-dimensional grids, small-world graphs, general graphs,…

  → How about other perceived-cost models…? (other than risk-averse)

  → How about probabilistic models…?

- <u>The Price of Malice in other scenarios and games</u>

  → Routing, caching, etc…

  → Fear-Factor in other systems…?

  → Can we use Fear-Factor to improve networking…?

Recent study of congestion games:
„Congestion Games with Malicious Players“ by
M. Babaioff, R. Kleinberg, C. Papadimitriou (EC’07)