Network Games with Friends and Foes

Stefan Schmid
T-Labs / TU Berlin

Who are the participants in the Internet?

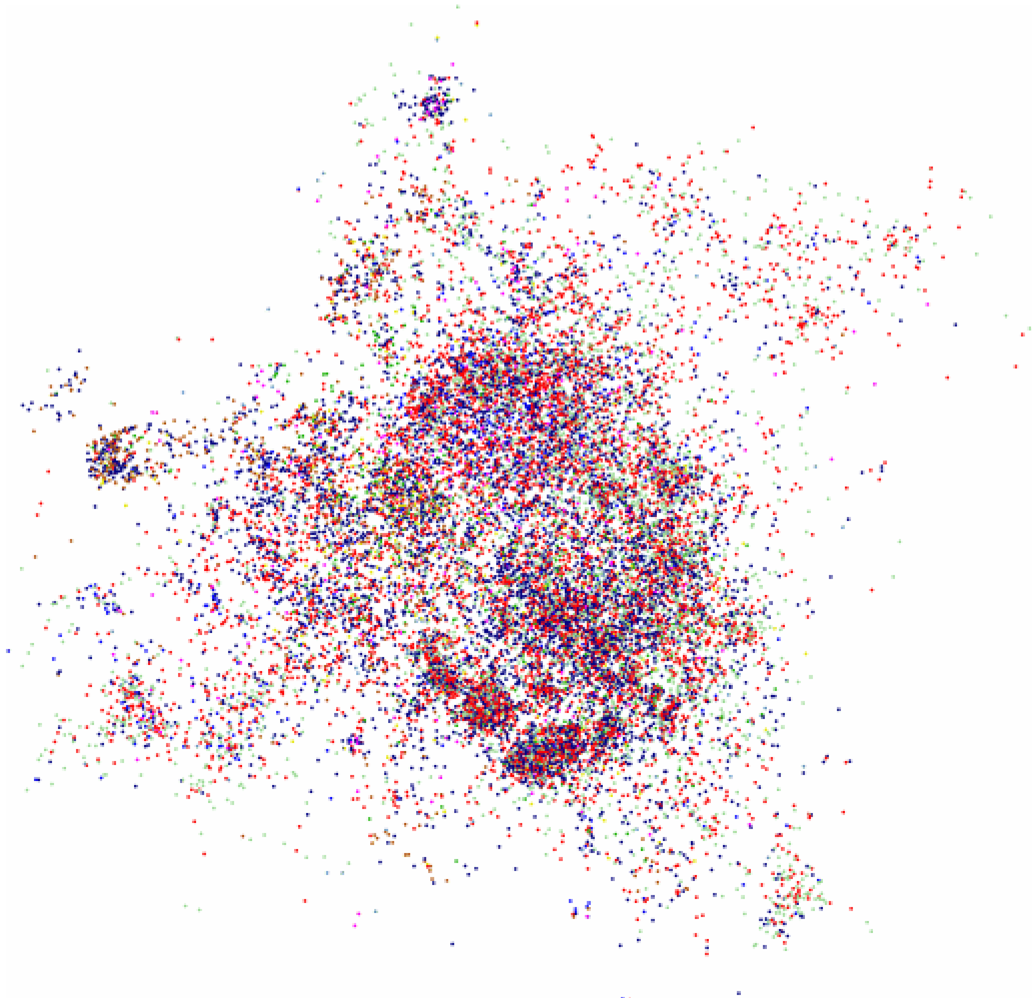

How to Model the Internet?

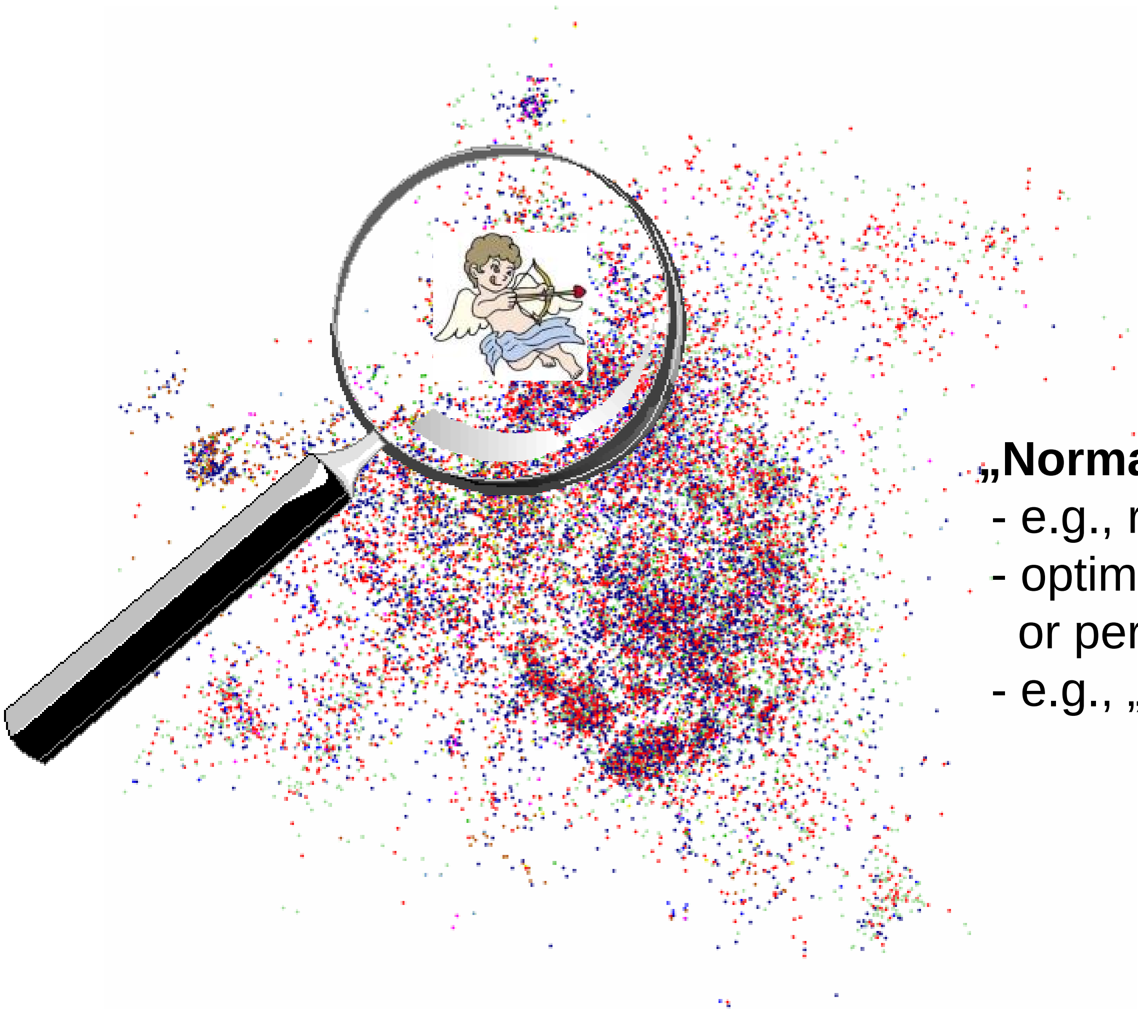

„Normal participants“:

- e.g., running „standard protocols“
- optimized for overall welfare (?) or perfectly „collaborating“
- e.g., „peer-to-peer“ systems?

How to Model the Internet?

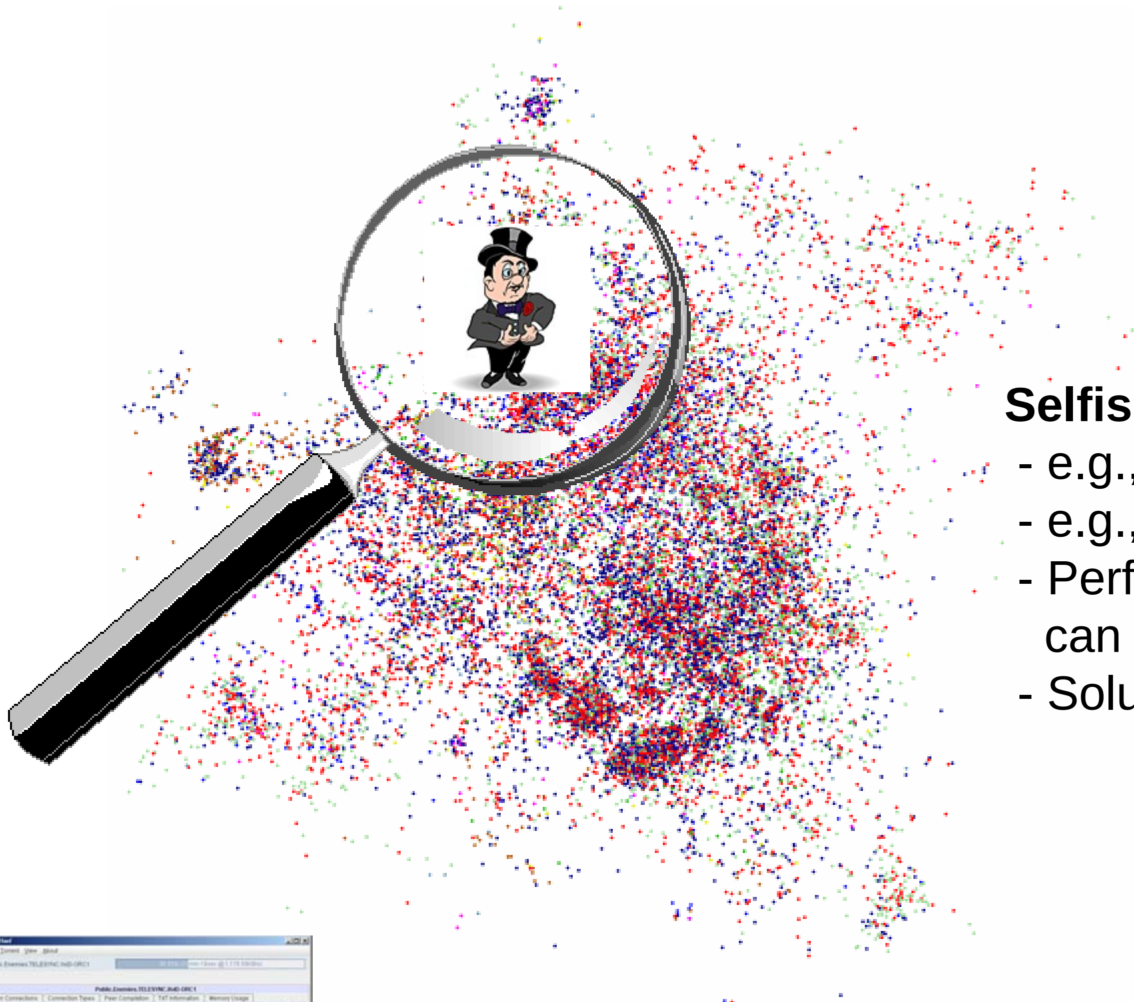

Selfish players à la Game Theory...:
- e.g., profit-maximizing company?
- e.g., BitThief BitTorrent client*?
- Performance / Nash equilibrium can be sub-optimal
- Solution: mechanism design?

*** BitThief: Downloading without reciprocating (i.e., uploading), see http://dcg.ethz.ch/projects/bitthief/**

How to Model the Internet?

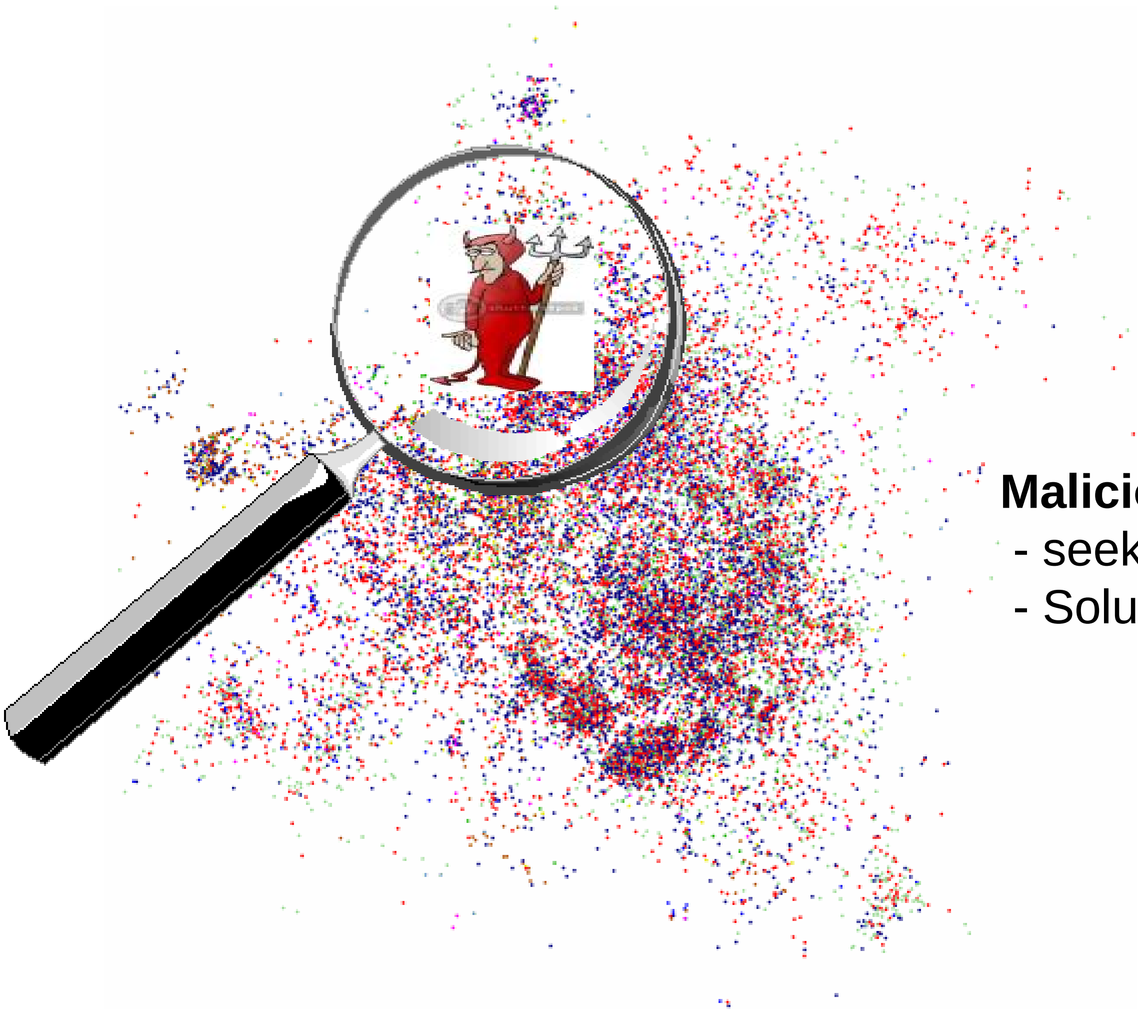

Malicious players:
- seek to harm the system (e.g., p2p)
- Solution: access control?

How to Model the Internet?

Social networks:

- Player selfish...
- ... but supports friends!

Socio-Economic Complexity (Papadi‘ 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.“

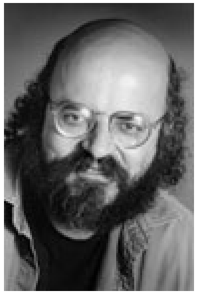

Game Theory Reveals Interesting Phenomena

Example: **Windfall of Malice**

The presence of malicious players can yield better outcomes!

„When Selfish Meets Evil“ / „Price of Malice“

Moscibroda, Schmid, Wattenhofer (PODC 2006 + IM 2010)

When selfish players are afraid of malicious players, they play more safely.
This can lead to better outcomes.

„Congestion Games with Malicious Players“

Babaioff, Kleinberg, Papadimitriou (EC 2007 + GEB 2009)

When malicious players need to route traffic through the network as well,
inefficiencies of Braess paradoxes may be avoided, yielding a better outcome.

Example:

The Virus Inoculation Game

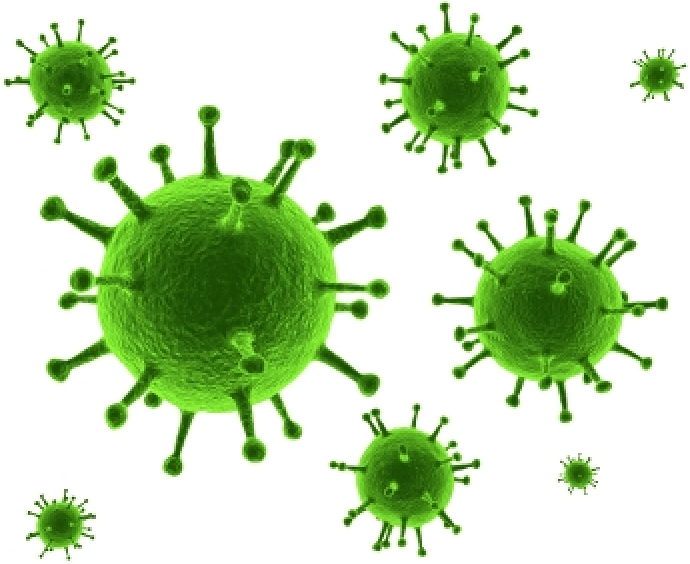

Virus Game: Setting

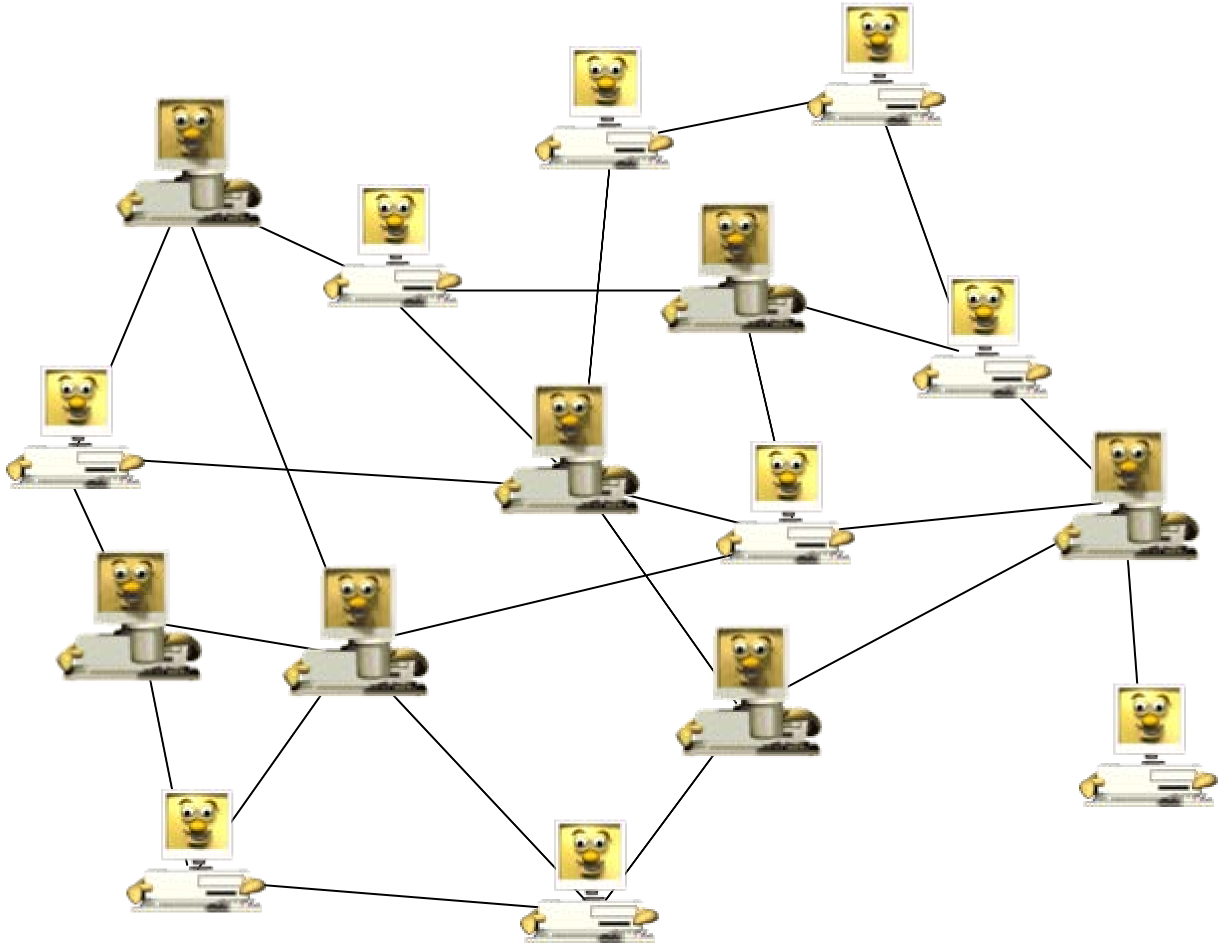

Virus Game: Players May Inoculate

Inoculation costs: C
Virus damage costs: L

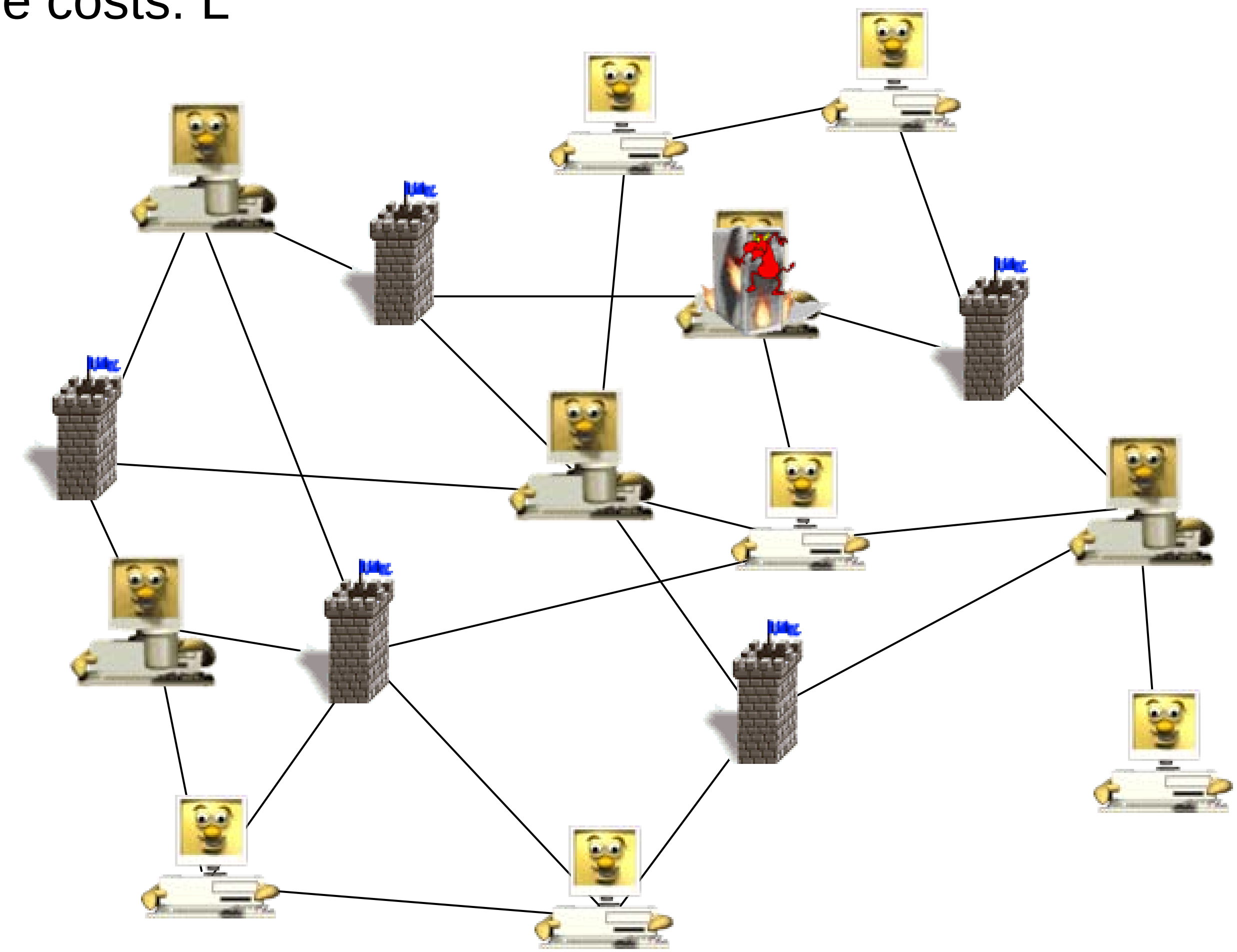

Virus Game: Virus Propagates...

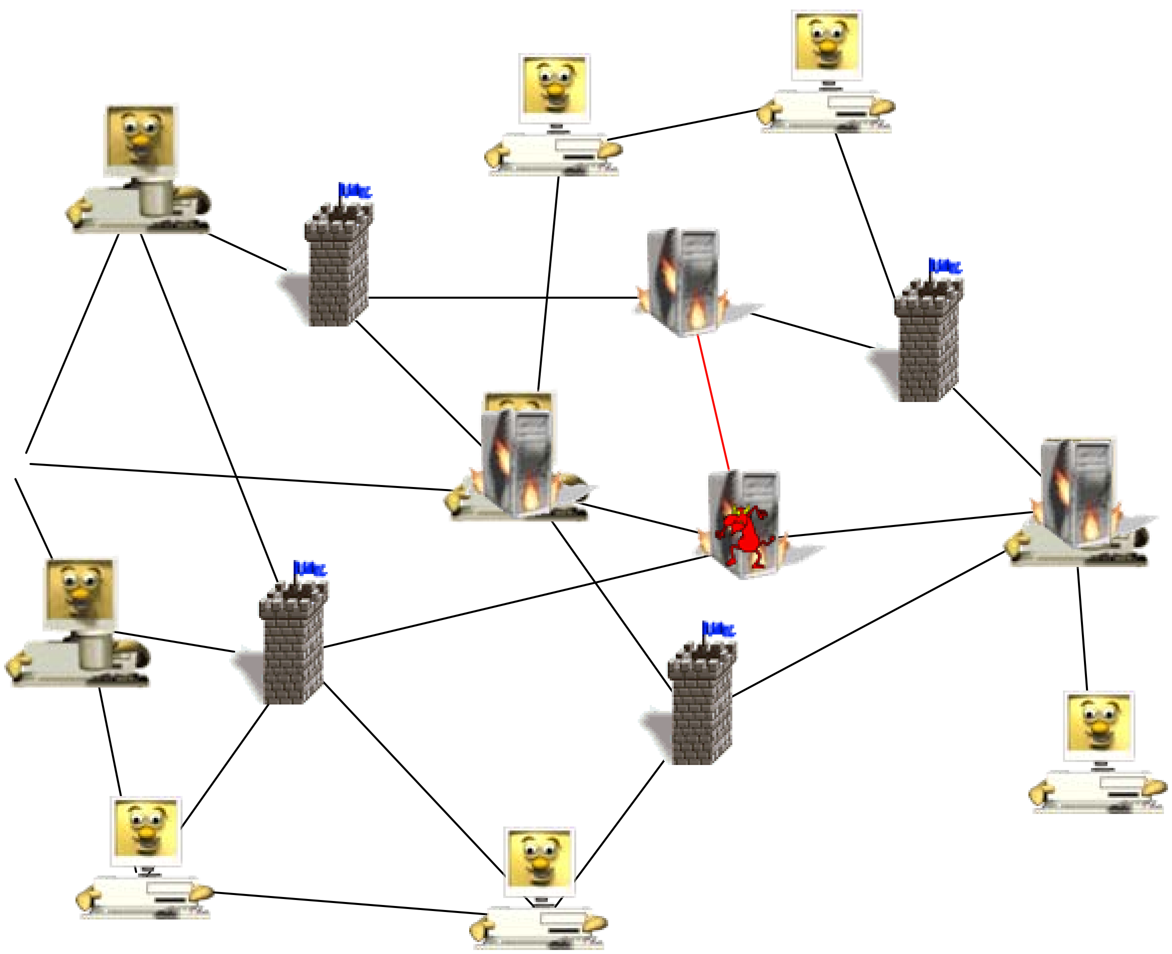

Virus Game: ... in Attack Component

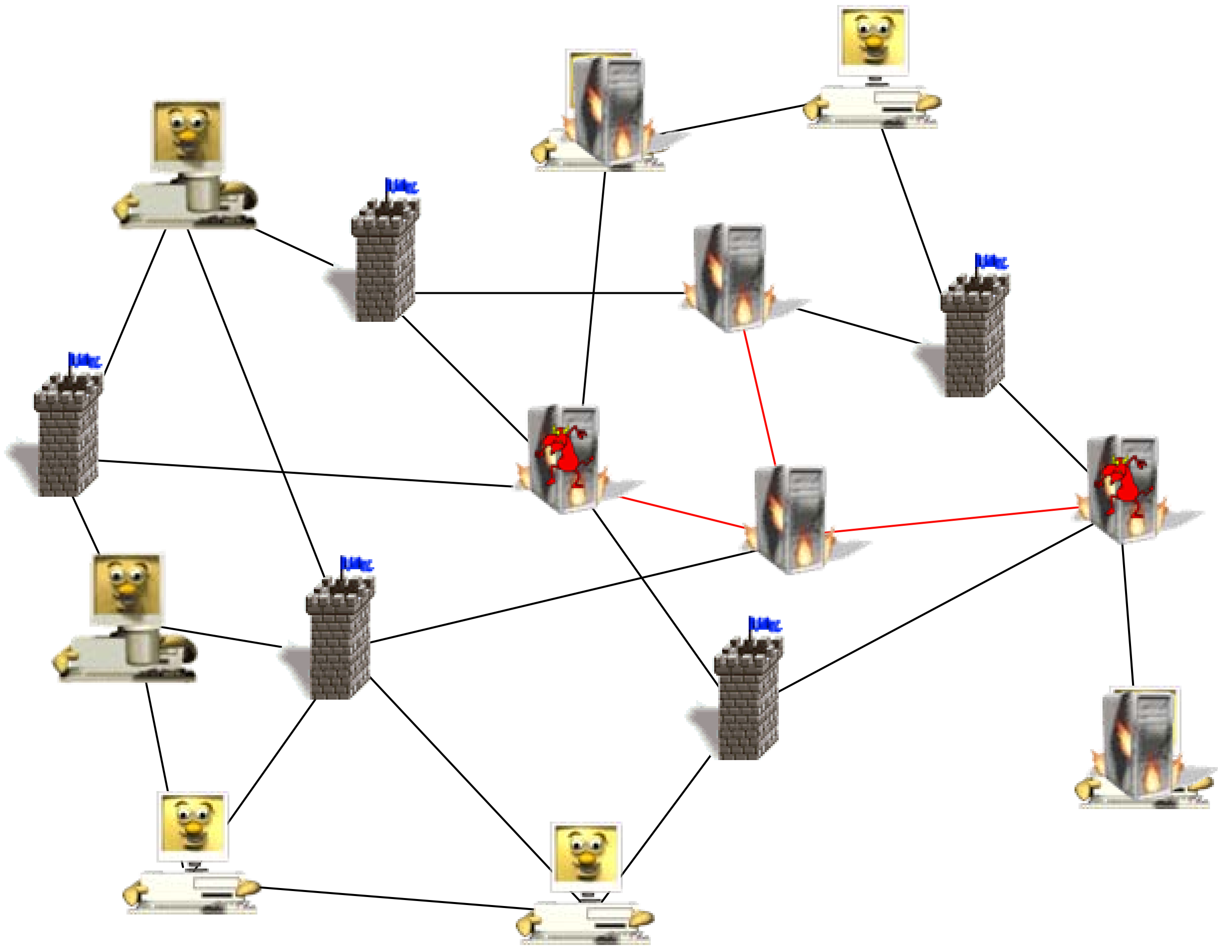

Virus Game

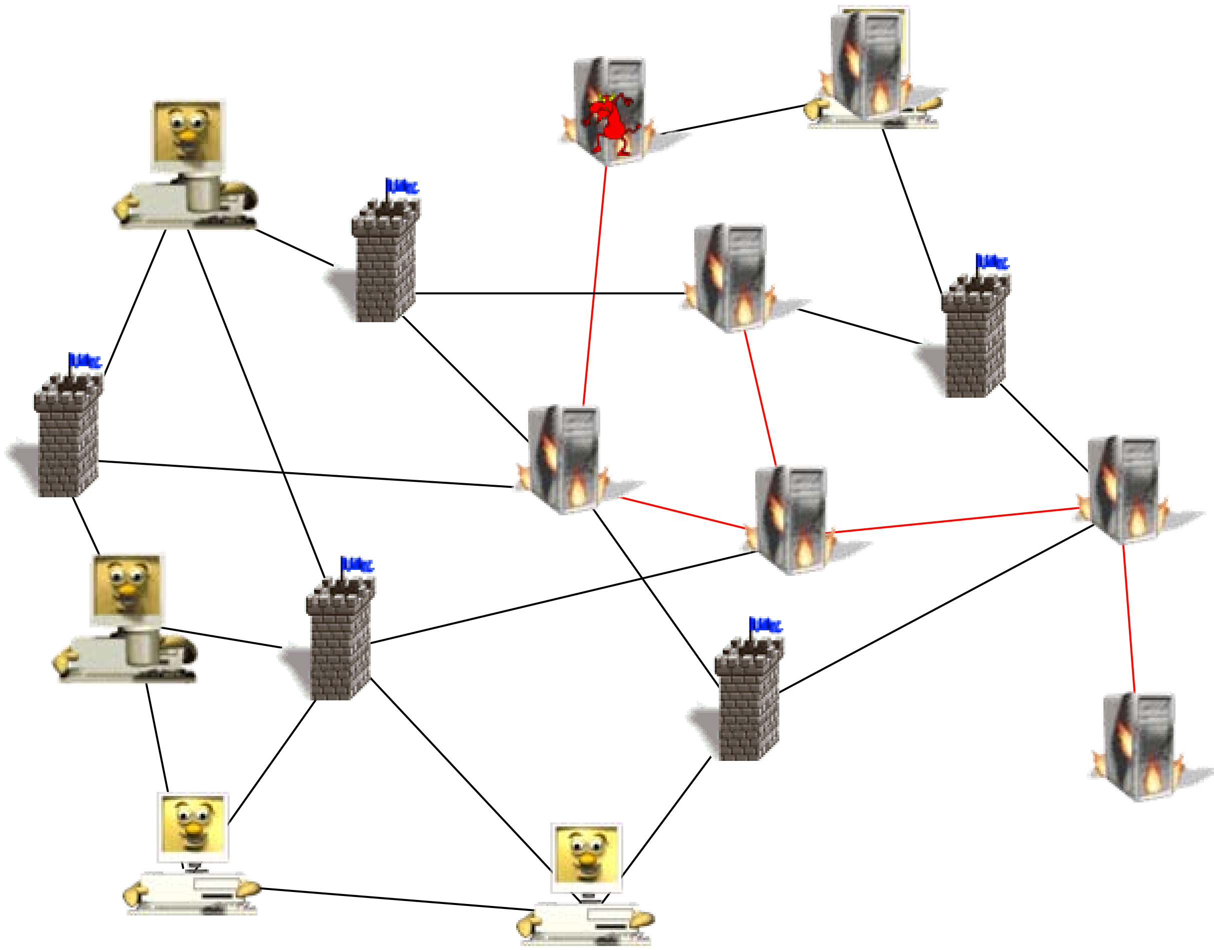

A Social Network Model

- 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

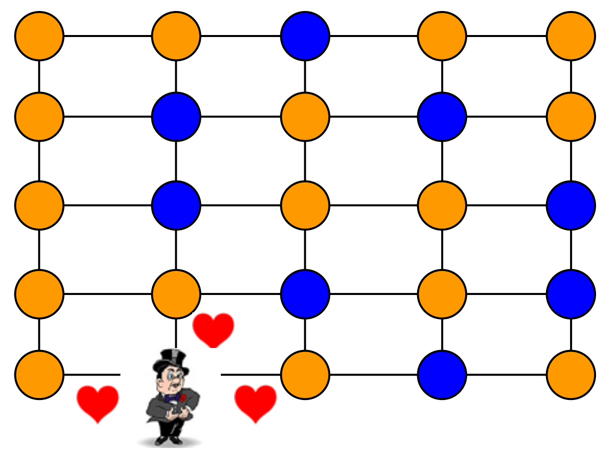

Some Definitions

- Quantify effects of social behavior vs purely selfish behavior?

- Social costs
  - Sum over all players‘ actual costs

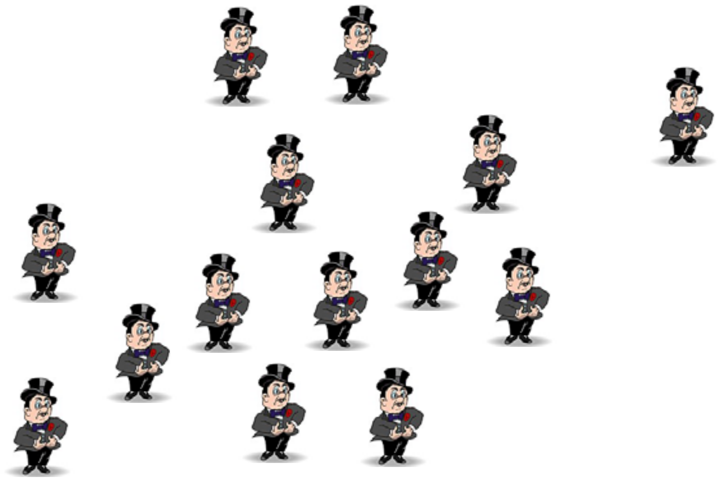

- Typical Assumption: Selfish player end up in an equilibrium (if it exists)

- Nash equilibrium
  - 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!

Windfall of Friendship

- 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")

Results based on

„On the Windfall of Friendship"
Meier, Oswald, Schmid, Wattenhofer (EC 2008)

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

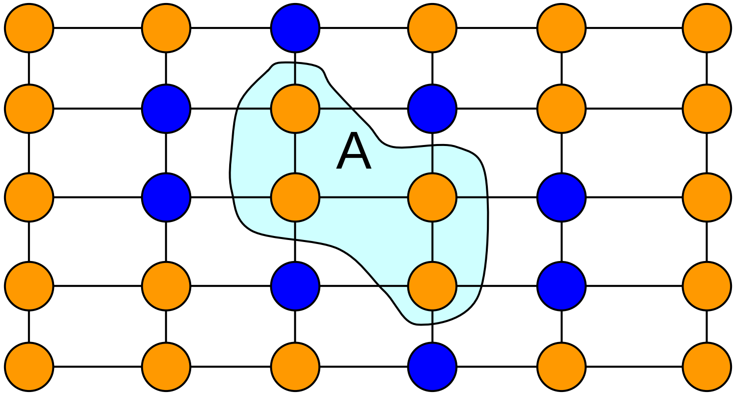

- Secure player: if she became insecure, she would be in attack component of size at least Cn/L
  - same argument: otherwise it‘s worthwhile to change strategies

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

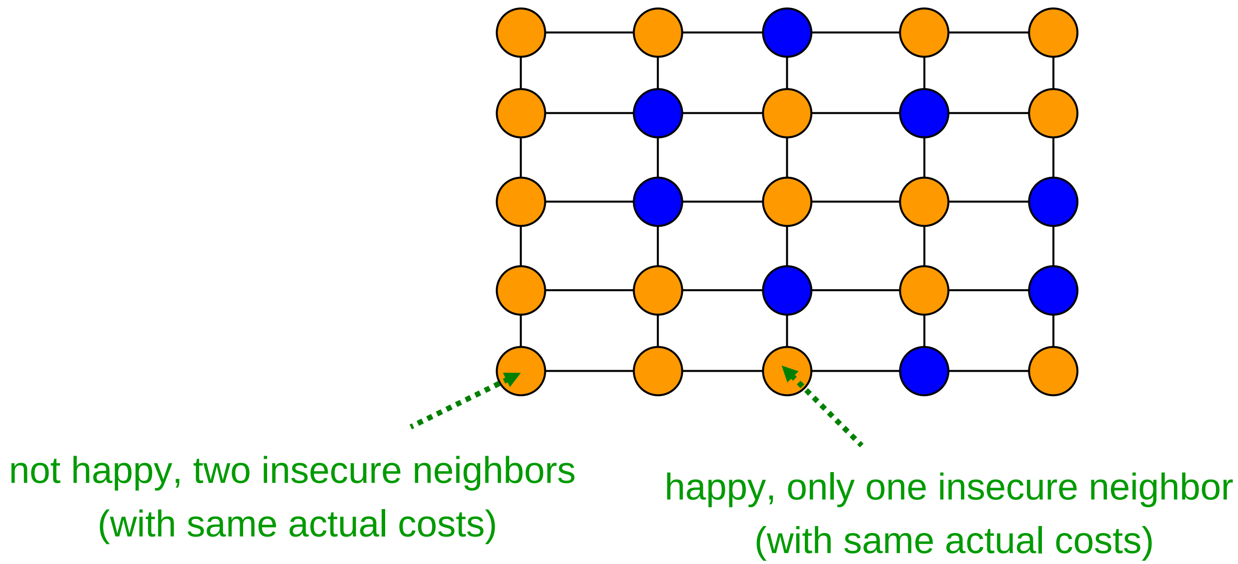

Bounds for the Windfall

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):

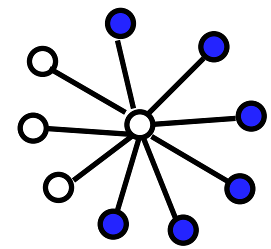

Social cost = (n/L)/n * n/L * L + (n-n/L) ~ n

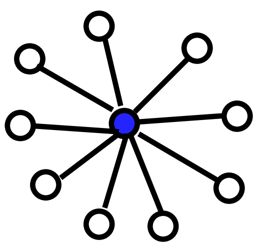

- In friendship Nash equilibrium, there are situations where center *must* inoculate, yielding optimal social costs of (for C=1):

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: players 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

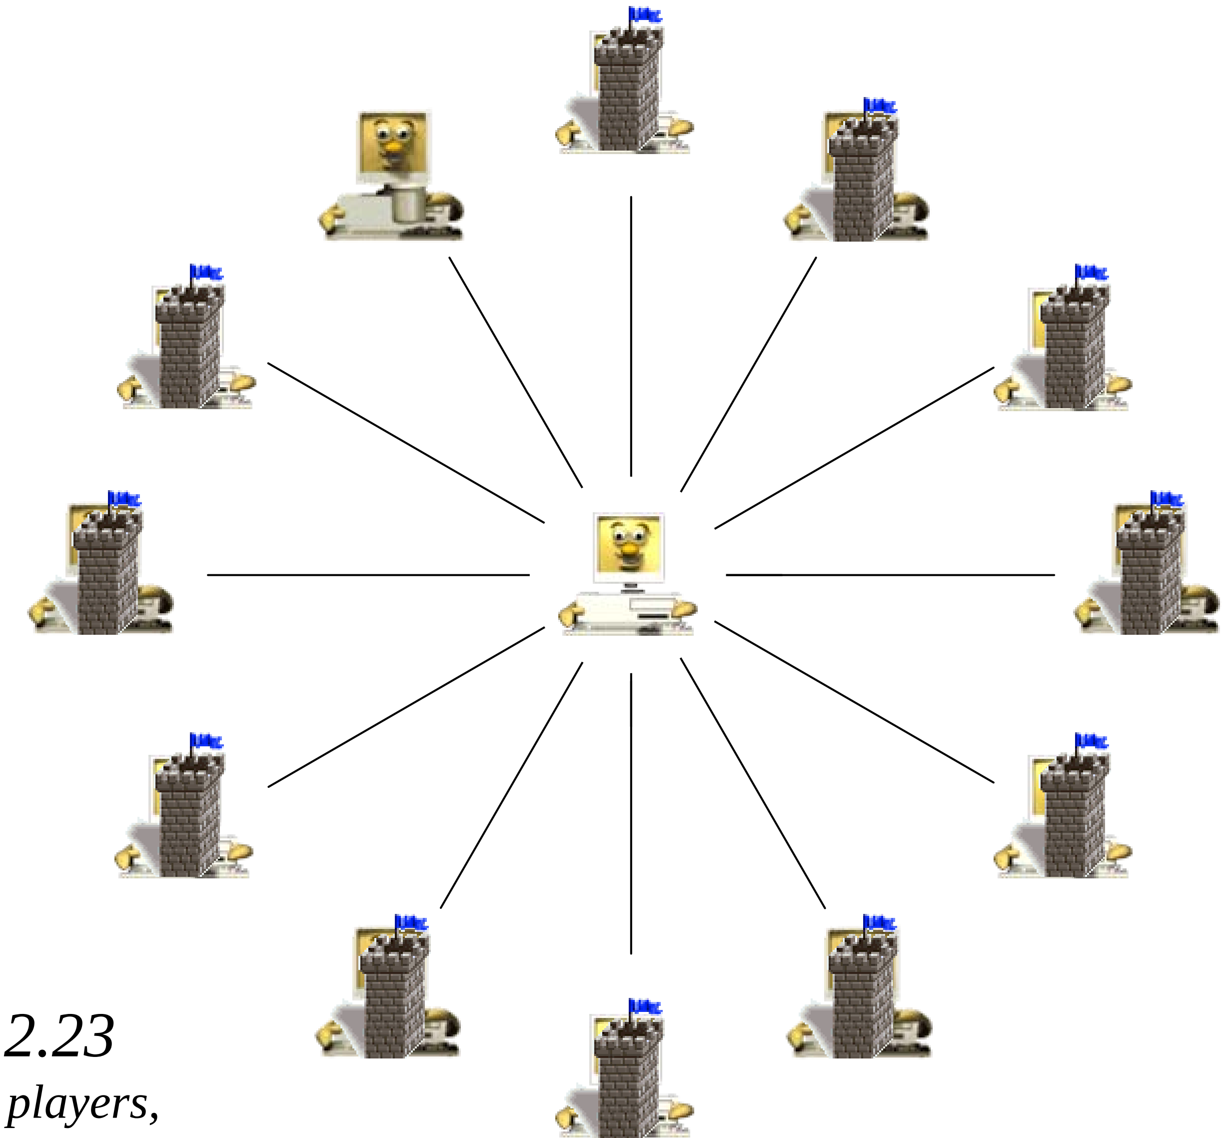

total cost = 12.23
(many inoculated players, attack component size two)

Monotonicity: Counterexample

n = 13
C = 1
L = 4
F = 0.1

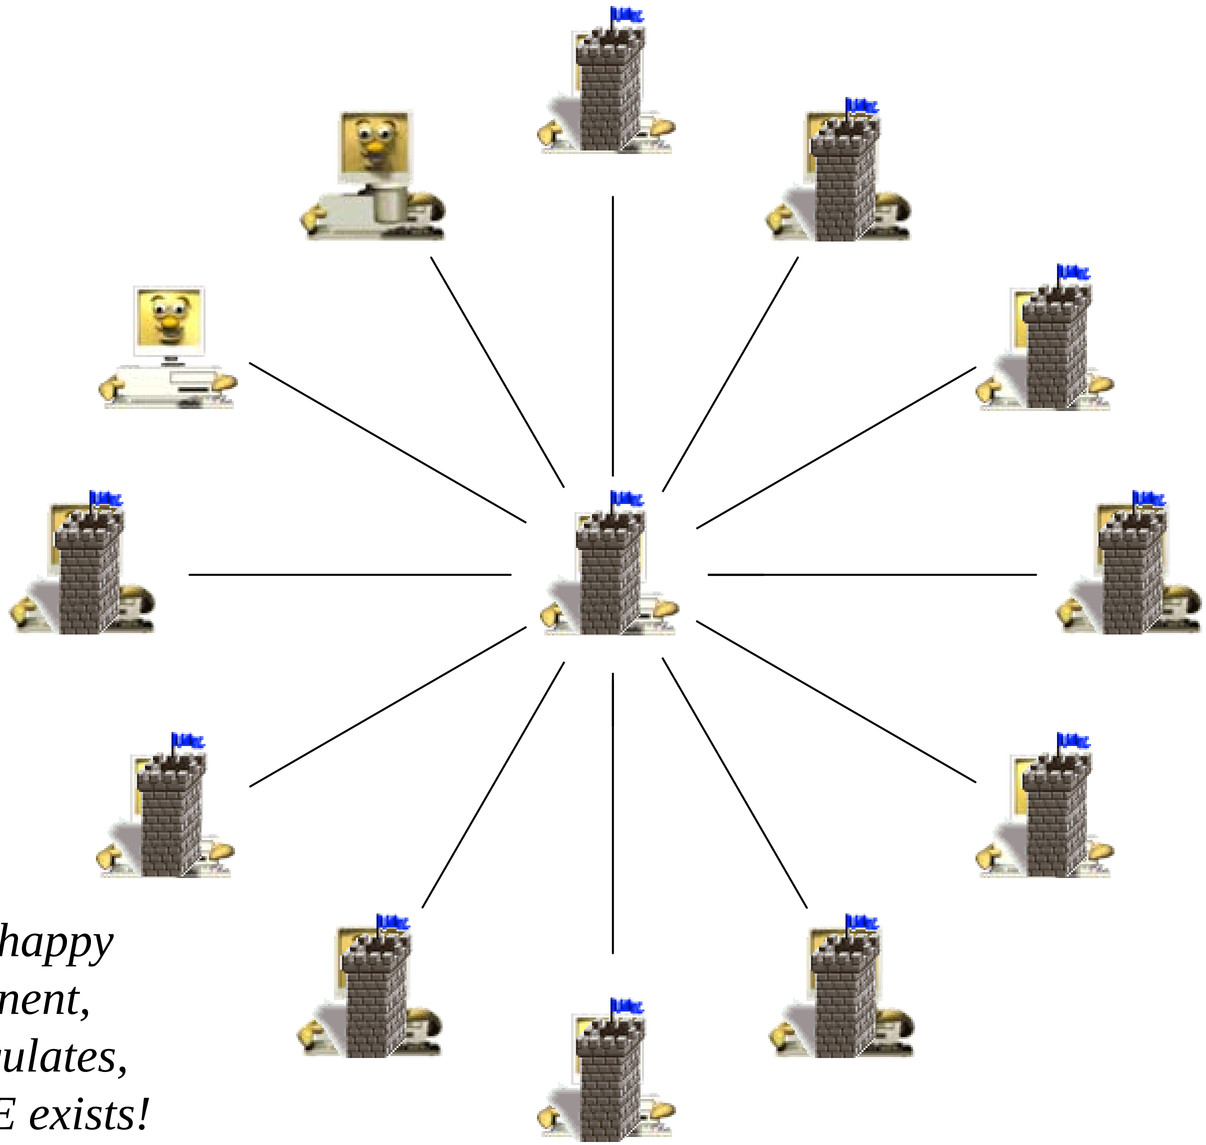

Boundary players happy with larger component, center always inoculates, thus: only this FNE exists!
total cost = 4.69

Towards Socially Heterogeneous Networks

The Social Range Matrix

f_{ij} = How much does player i care about player j?

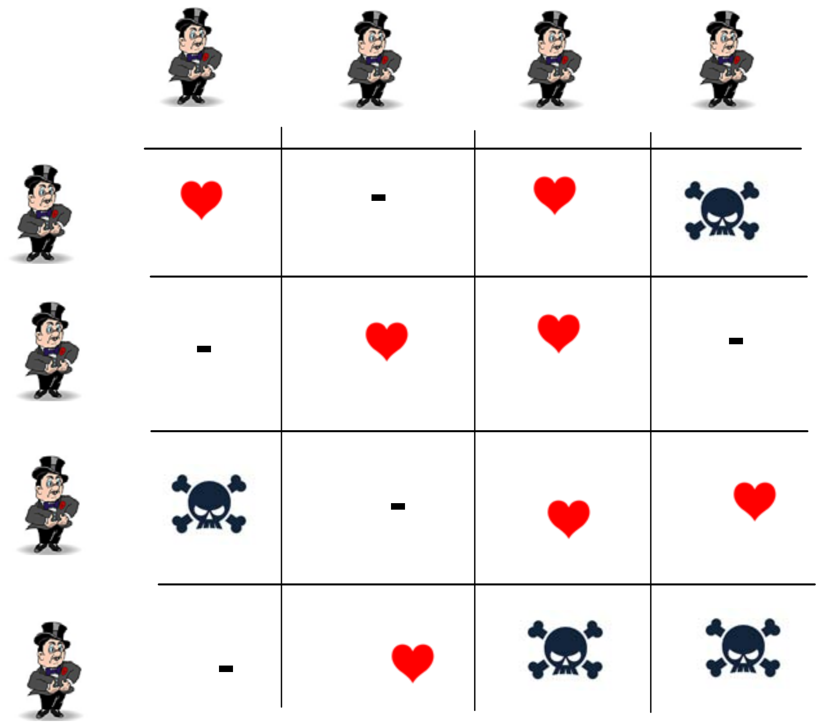

Costs and Equilibria

c_a(i,s) : actual cost of player *i* in profile *s*

c_p(i,s) : perceived costs of player i in profile s

$$c_p(i,s) = \sum_i f_{i,j}\, c_a(j,s)$$

Social equilibrium: No player has can reduce perceived costs by changing the strategy, given the other players‘ strategies.

The Social Range Matrix: Examples

Classic game theory: Anarchy

1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

The Social Range Matrix: Examples

Altruisitic setting:

1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1

Note: Social optimum is an equilibrium!

The Social Range Matrix: Examples

Monarchy setting:

ε	1	0	0
0	1	0	0
0	1	ε	0
0	1	0	ε

Arbitrarily small ε > 0

The Social Range Matrix: Examples

Selfish and a bad guy (seeks to minimize system performance):

	1	0	0	0
	0	1	0	0
	0	0	1	0
	-1	-1	-1	-1

The Social Range Matrix: Properties

Matrix can be „scaled“:

1	1	2	0	:2	1/2	1/2	1	0
3	3	3	3	:3	1	1	1	1
1	0	1	0	:1	1	0	1	0
4	1	2	4	:4	1	1/4	1/2	1

⇔

Multiplying a row by the same factor does not change equilibria or convergence!

Example:

Network Creation Game

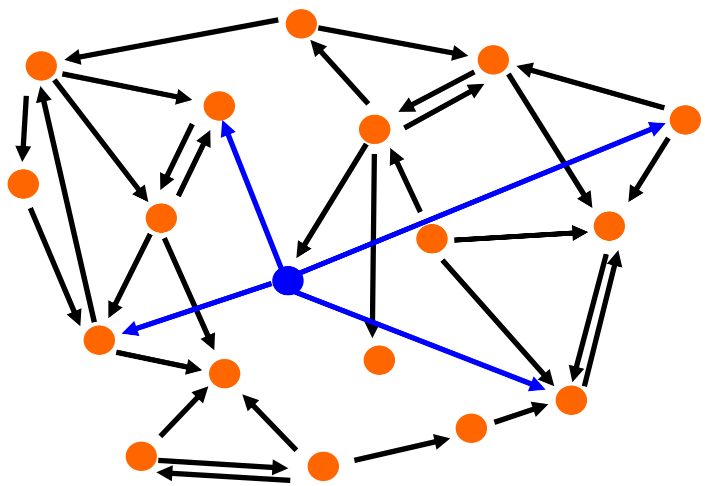

Lookup in Unstructured P2P Network

R=0:

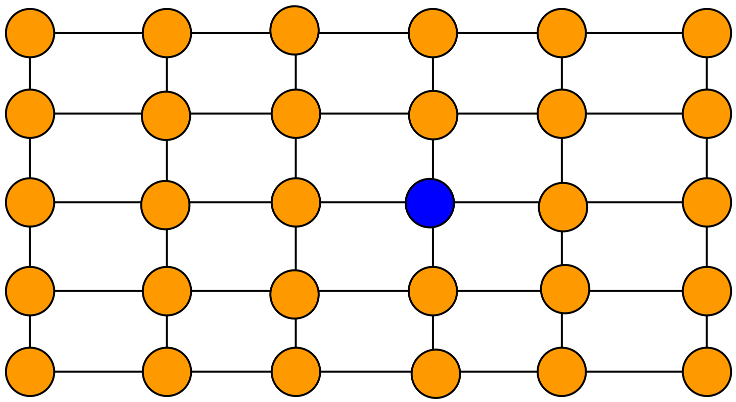

R=1:

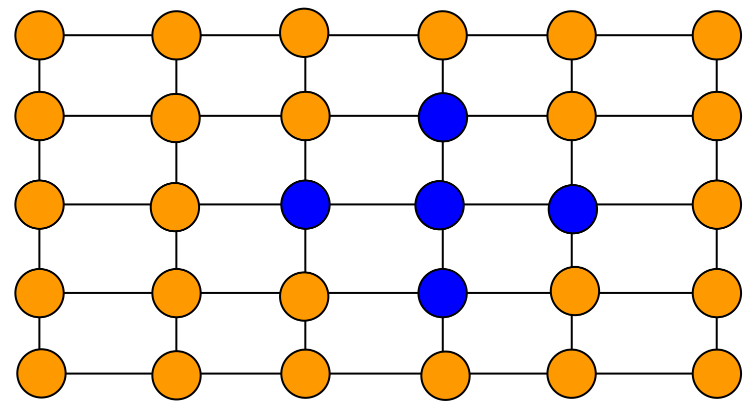

R=2:

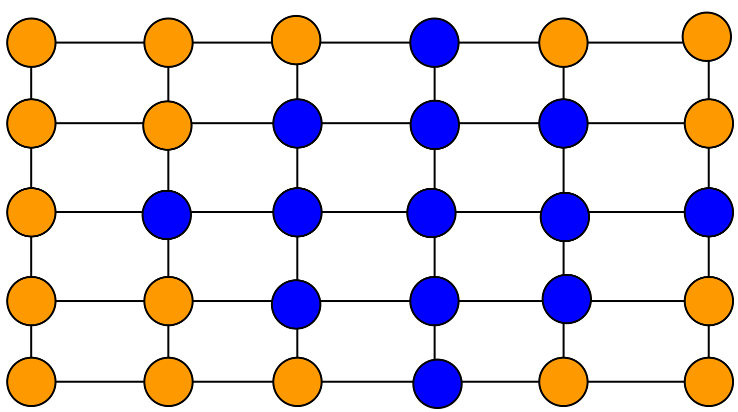

Gnutella, e.g., R=5.

The Game

Undirected graph where one end pays for the connection:

$$c_a(i,s) = \alpha \cdot s_i - g\left(\sum_{j=1}^{R} |\Gamma^j(i,s)|\right)$$

utility increase as a function of R-hop neighborhood

parameter

number of connections payed for

set of neighbors at hop distance j

Typically, g is *concave*: The marginal benefit declines (e.g., chance of new data).

Some Results

For a complete set of results, see:

„Towards Network Games with Social Preferences“
Kuznetsov, Schmid (SIROCCO 2010)

Theorem: For R=1, any social range matrix has a pure equilibrium, for any α.

Theorem: If $g(x+1)-g(x) > \alpha/2$, then the social optimum is either the clique (if R=1) or a tree of diameter at most R.

Theorem: Anarchy can have a higher social welfare than monarchy (if α large), and vice versa!

Social Relationships and Topology

Intuitively, there is a tight relationship between the social range matrix and equilibrium topologies, e.g., for R=1:

Players tend to connect to friends more than to foes.

Theorem: If $f_{ij} \in \{0,1\}$, $f_{ii}=1$, and $1<\alpha<2$:
There is a Nash equilibrium whose adjacency matrix corresponds to the social range matrix.

Social Relationships and Topology

Theorem: If $f_{ij} \in \{0,1\}$, $f_{ii}=1$, and $1<\alpha<2$:
There is a Nash equilibrium whose adjacency matrix corresponds to the social range matrix.

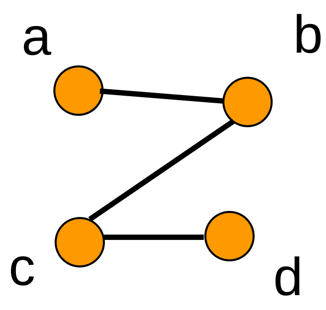

F_{ij}:

	a	b	c	d
a	**1**	**1**	**0**	**0**
b	**1**	**1**	**0**	**0**
c	**0**	**1**	**1**	**0**
d	**0**	**0**	**1**	**1**

A_{ij}:

	a	b	c	d
a	**0**	**1**	**0**	**0**
b	**1**	**0**	**1**	**0**
c	**0**	**1**	**0**	**1**
d	**0**	**0**	**1**	**0**

Windfall of Friendship

Theorem: Society can only benefit from additional „1“-entries in social range matrix!
(Worst and best equilibrium are not worse.)

F_{ij}:

	a	b	c	d
a	**1**	**1**	**0**	**0**
b	**1**	**1**	**0**	**0**
c	**0**	**1**	**1**	**0**
d	**0**	**0**	**1**	**1**

☹

F_{ij}:

	a	b	c	d
a	**1**	**1**	**1**	**0**
b	**1**	**1**	**0**	**0**
c	**0**	**1**	**1**	**0**
d	**1**	**0**	**1**	**1**

☺

There is an equilibrium with at least as many connections
(yielding higher social welfare).

Price of Ill-Will

Theorem: The worst and best equilibrium can only be better if -1 entries are turned to 0 entries.

F_{ij}:

	a	b	c	d
a	**1**	**-1**	**0**	**0**
b	**-1**	**1**	**0**	**0**
c	**0**	**-1**	**1**	**0**
d	**0**	**0**	**-1**	**1**

☹

F_{ij}:

	a	b	c	d
a	**1**	**0**	**0**	**0**
b	**-1**	**1**	**0**	**0**
c	**0**	**0**	**1**	**0**
d	**0**	**0**	**-1**	**1**

☺

Conclusion

1. Models that capture socio-economic complexity of today‘s distributed systems?

2. Social range matrix as a further step.

3. Interesting phenomena „Windfall of Malice“ or „Price of Friendship“
   - depend on game

4. Network topologies reflect social relationships

5. Many open questions!

Thanks!

Thanks to my collaborators:




Petr Kuznetsov, Thomas Moscibroda, Yvonne Anne Pignolet, Roger Wattenhofer & Dominic Meier

Questions?

Papers online:

http://www.net.t-labs.tu-berlin.de/~stefan/