Dynamics and Cooperation

Algorithmic Challenges in Peer-to-Peer Computing

PhD Defense

Stefan Schmid

Distributed Computing Group
ETH Zurich

Peer-to-Peer Technology

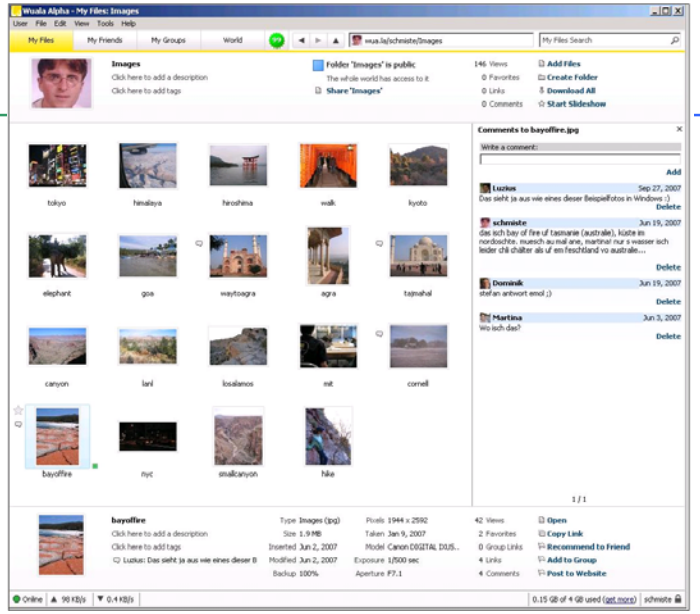

- Well-known p2p systems
  - Internet telephony: *Skype*
  - File sharing: *BitTorrent, eMule, ...*
  - Streaming: *Zattoo, Joost, ...*

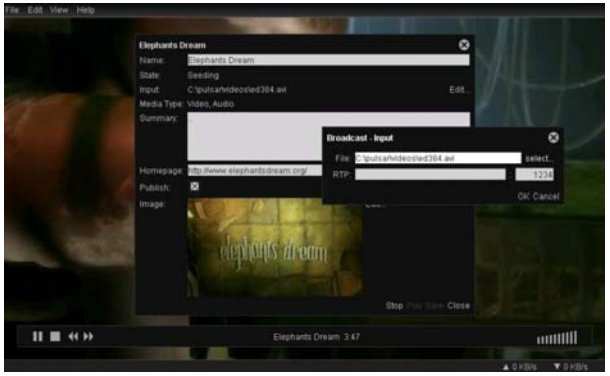

- Other (well-known?) systems
  - Pulsar streaming system (e.g., *tilllate* clips?)
  - Wuala online storage system

- Impact: Accounts for much Internet traffic!

(source: *cachelogic.com*)

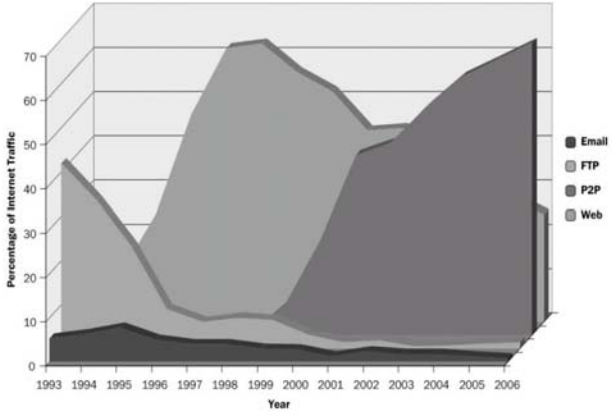

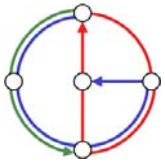

A Hot Topic...

- Hot topic in society / press / blogs / ...
  - start-ups receive media coverage (e.g., student project Wuala, 12,000 users, 3,000 in queue...)
  - new clients discussed in blogs (e.g., BitThief)

The Paradigm

- Key concepts
  - Machines (peers) in the network: consumer and producer of resources („all machines have equal role“)
  - Use of decentralized resources on the edge of the Internet (e.g., desktops)

- Benefits
  - Scalability: More resources in larger networks („cake grows“)
  - Robustness: No single point of failure
  - Can outperform server-based solutions
  - Cheap: start-up vs Google

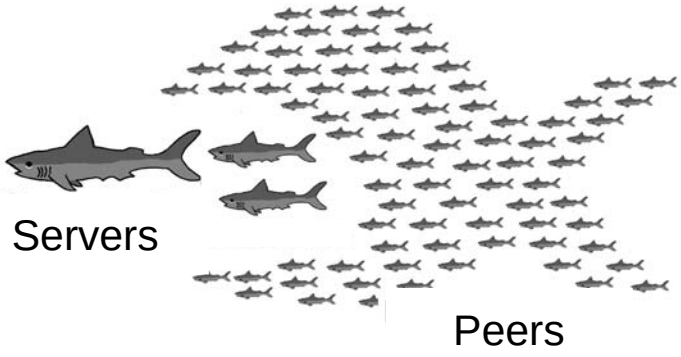

- Therefore:
  - No need for expensive infrastructure at content distributors
  - Democratic aspect: Anyone can publish media contents / speeches

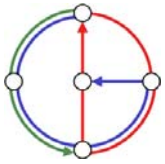

Implications and Challenges

- Implies that participating machines are often
  - under the control of (anonymous) individuals
  - dynamic (join only to download a file)
  - heterogeneous („more equal“, Internet connection, CPU power)
  - unreliable (desktop computers)
  - geographically distributed (latencies)
  - may be non-cooperative (no voluntary resource contributions)

- Some challenges
  - Discovery of dynamic resources which do not have fixed addresses?
  - Include weak and strong participants
  - Ensuring user anonymity
  - etc.

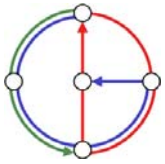

PhD Thesis (1)

- Focus on two challenges: **dynamics** and **cooperation**

Distributed Asymmetric Verification in Computational Grids
Michael Kuhn, Stefan Schmid, and Roger Wattenhofer.
22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS), Miami, Florida, USA, April 2008.
Documents: paper pdf meta bibtex

Structuring Unstructured Peer-to-Peer Networks
Stefan Schmid and Roger Wattenhofer.
14th Annual IEEE International Conference on High Performance Computing (HiPC), Goa, India, Springer LNCS 4873, December 2007.
Documents: paper pdf slides ppt meta bibtex

Manipulation in Games
Raphael Eidenbenz, Yvonne Anne Oswald, Stefan Schmid, and Roger Wattenhofer.
18th International Symposium on Algorithms and Computation (ISAAC), Sendai, Japan, Springer LNCS 4835, December 2007.
Documents: paper pdf slides ppt techreport pdf meta bibtex

Push-to-Pull Peer-to-Peer Live Streaming
Thomas Locher, Remo Meier, Stefan Schmid, and Roger Wattenhofer.
21st International Symposium on Distributed Computing (DISC), Lemesos, Cyprus, Springer LNCS 4731, September 2007.
Documents: paper pdf meta bibtex

Rescuing Tit-for-Tat with Source Coding
Thomas Locher, Stefan Schmid, and Roger Wattenhofer.
7th IEEE International Conference on Peer-to-Peer Computing (P2P), Galway, Ireland, September 2007.
Documents: paper pdf ps slides pdf meta bibtex

Mechanism Design by Creditability
Raphael Eidenbenz, Yvonne Anne Oswald, Stefan Schmid, and Roger Wattenhofer.
1st International Conference on Combinatorial Optimization and Applications (COCOA), Xi'an, Shaanxi, China, Springer LNCS 4616, August 2007.
Documents: paper pdf slides pdf techreport pdf meta bibtex

Dynamic Internet Congestion with Bursts
Stefan Schmid and Roger Wattenhofer.
13th Annual IEEE International Conference on High Performance Computing (HiPC), Bangalore, India, Springer LNCS 4297, December 2006.
Documents: paper pdf slides ppt meta bibtex

Free Riding in BitTorrent is Cheap
Thomas Locher, Patrick Moor, Stefan Schmid, and Roger Wattenhofer.
5th Workshop on Hot Topics in Networks (HotNets), Irvine, California, USA, November 2006.
Documents: paper pdf ps slides pdf meta bibtex

Cryptree: A Folder Tree Structure for Cryptographic File Systems
Dominik Grolimund, Luzius Meisser, Stefan Schmid, and Roger Wattenhofer.
25th IEEE Symposium on Reliable Distributed Systems (SRDS), Leeds, United Kingdom, October 2006.
Documents: paper pdf slides ppt meta bibtex

eQuus: A Provably Robust and Locality-Aware Peer-to-Peer System
Thomas Locher, Stefan Schmid, and Roger Wattenhofer.
6th IEEE International Conference on Peer-to-Peer Computing (P2P), Cambridge, United Kingdom, September 2006.
Documents: paper pdf ps slides pdf meta bibtex

A TCP with Guaranteed Performance in Networks with Dynamic Congestion and Random Wireless Losses
Stefan Schmid and Roger Wattenhofer.
2nd Annual International Wireless Internet Conference (WICON), Boston, Massachusetts, USA, August 2006.
Documents: paper pdf slides ppt meta bibtex

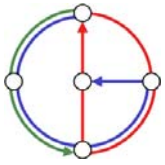

PhD Thesis (2)

- Focus on two challenges: **dynamics** and **cooperation**

When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game
Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer.
25th Annual Symposium on Principles of Distributed Computing (PODC), Denver, Colorado, USA, July 2006.
Documents: paper pdf ps meta bibtex

On the Topologies Formed by Selfish Peers
Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer.
25th Annual Symposium on Principles of Distributed Computing (PODC), Denver, Colorado, USA, July 2006.
Documents: paper pdf ps meta bibtex

A Blueprint for Constructing Peer-to-Peer Systems Robust to Dynamic Worst-Case Joins and Leaves
Fabian Kuhn, Stefan Schmid, Joest Smit, and Roger Wattenhofer.
14th IEEE International Workshop on Quality of Service (IWQoS), Yale University, New Haven, Connectitut, USA, June 2006.
Documents: paper pdf slides ppt meta bibtex

Havelaar: A Robust and Efficient Reputation System for Active Peer-to-Peer Systems
Dominik Grolimund, Luzius Meisser, Stefan Schmid, and Roger Wattenhofer.
1st Workshop on the Economics of Networked Systems (NetEcon), University of Michigan, Ann Arbor, Michigan, USA, June 2006.
Documents: paper pdf slides ppt meta bibtex

Algorithmic Models for Sensor Networks
Stefan Schmid and Roger Wattenhofer. (Invited paper)
14th International Workshop on Parallel and Distributed Real-Time Systems (WPDRTS), Island of Rhodes, Greece, April 2006.
Documents: paper pdf ps slides ppt meta bibtex

On the Topologies Formed by Selfish Peers
Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer.
5th International Workshop on Peer-to-Peer Systems (IPTPS), Santa Barbara, California, USA, February 2006.
Documents: paper pdf ps slides ppt meta bibtex

A Robust Interference Model for Wireless Ad-Hoc Networks
Pascal von Rickenbach, Stefan Schmid, Roger Wattenhofer, and Aaron Zollinger.
5th International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN), Denver, Colorado, USA, April 2005.
Documents: paper pdf ps slides pdf meta bibtex

A Self-Repairing Peer-to-Peer System Resilient to Dynamic Adversarial Churn
Fabian Kuhn, Stefan Schmid, and Roger Wattenhofer.
4th International Workshop on Peer-To-Peer Systems (IPTPS), Cornell University, Ithaca, New York, USA, Springer LNCS 3640, February 2005.
Documents: paper pdf slides ppt meta bibtex

Parallel Compilation of CMS Software
Shaun Ashby, Giulio Eulisse, Stefan Schmid, and Lassi Tuura.
Computing in High Energy and Nuclear Physics Conference (CHEP), Interlaken, Switzerland, September 2004.
Documents: paper pdf meta bibtex

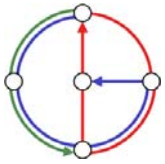

Some Questions

How to design a *BitTorrent* system which is provably robust to churn?

How to maximize throughput between two *eMule* peers if the available bandwidth changes in a worst-case manner?

Is it possible to free-ride in *BitTorrent*?

(Or: How to use BitTorrent without going to jail?)

How harmful is selfish behavior in *Gnutella*?

What is the effect on performance if there are malicious participants in a *swarm*?

How vulnerable is the *Skype* network to viruses which spread along the contact lists?

How to remove Simpsons from *Kad*?

Dynamics

Motivation

- Distributed resources
  - Under control of individual users (not one administrative domain)
  - Peers may only be online during their own „consumption time“

- Peer-to-peer paradigm relies an dynamic resources!
  - Unlike traditional multiprocessor architectures
  - Needs to cope with (or even exploit) the dynamics! (e.g., higher degree)

- A new topic...
  - Not many results known
  - Some basic principles such as the use of redundancy, local algorithms, etc.

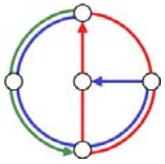

Resilience to Churn

- Our goal: Worst-case resilience to ongoing membership changes
  - i.e., resilience to adversarial topology changes

- Our reason:
  - Approach is pessimistic
  - However, it gives stronger guarantees
  - Includes a real adversary (worm along links, or crawler learnt topology)

- Our approach: Graph simulation

- Our results:
  - Despite ADV(log n, log n, 1), topology can be maintained
  - Peer degree and network diameter: O(log n)
  - Asymptotically optimal
  - Similarly for pancake graphs: replace log n by log n / loglog n

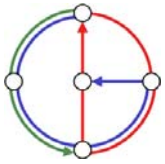

Recipe

1. Take a graph with desirable properties

2. Simulate the graph by representing each vertex by a set of peers

3. Find a token distribution algorithm on this graph

4. Find an algorithm to estimate the total number of peers in the system

5. Find an algorithm to adapt the graph‘s dimension

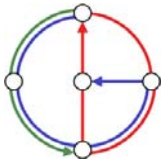

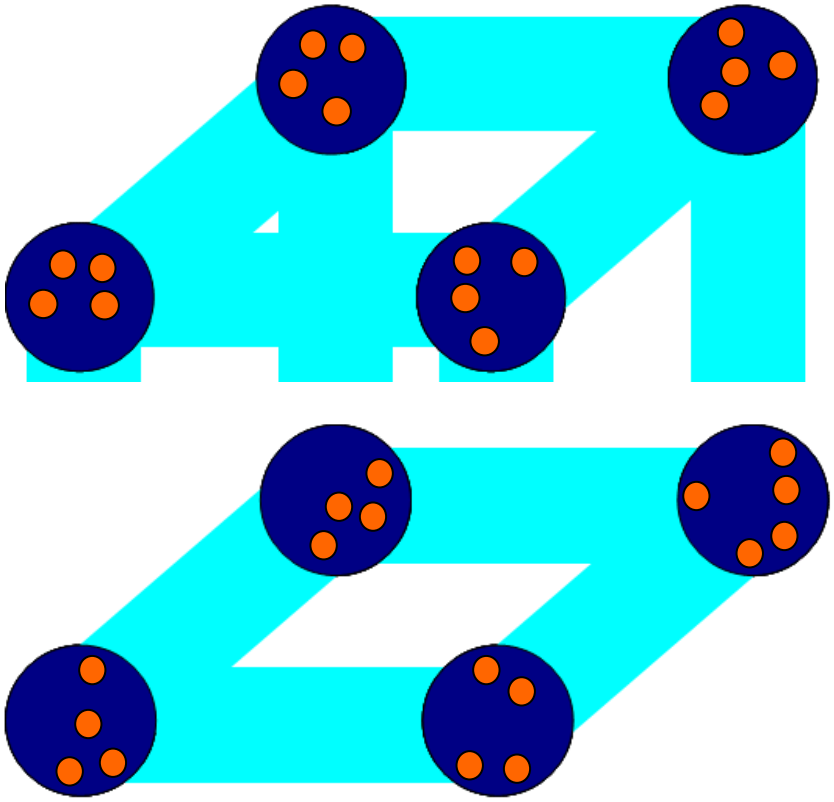

Cooperation

Motivation (1)

A good model for peer-to-peer networks?

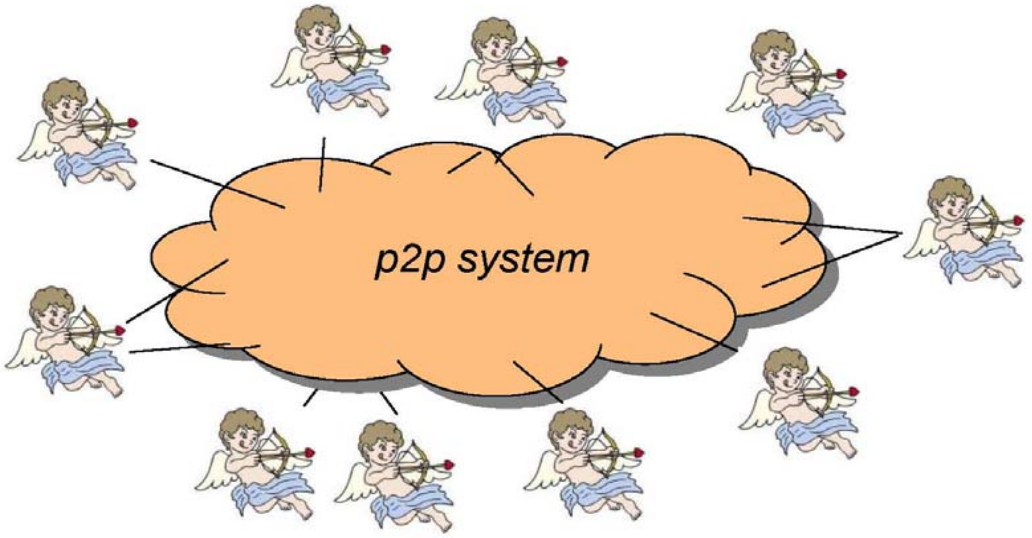

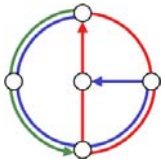

Motivation (2)

Participants may not be angels!

- May not volunarily contribute resources!
- Some participants may seek to harm the system!

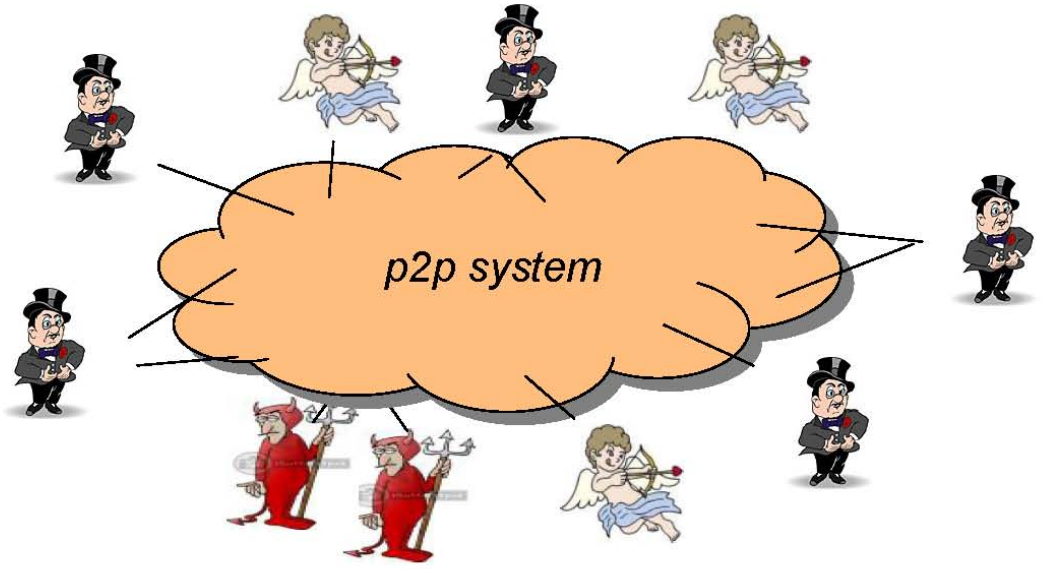

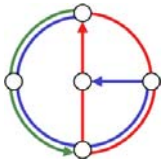

Free Riding Is Still Possible Today

- Case study BitTorrent
  - Our result: Free riding today is still possible
  - Even in the absence of seeders (optimistic unchoking mechanism!),
  - and despite the incentive mechanisms such as tit-for-tat policies!
  - Our client BitThief compartively fast if: many seeders (round robin), small files (large view exploit), or slow seeders.
  - Sharing communities can also be exploited

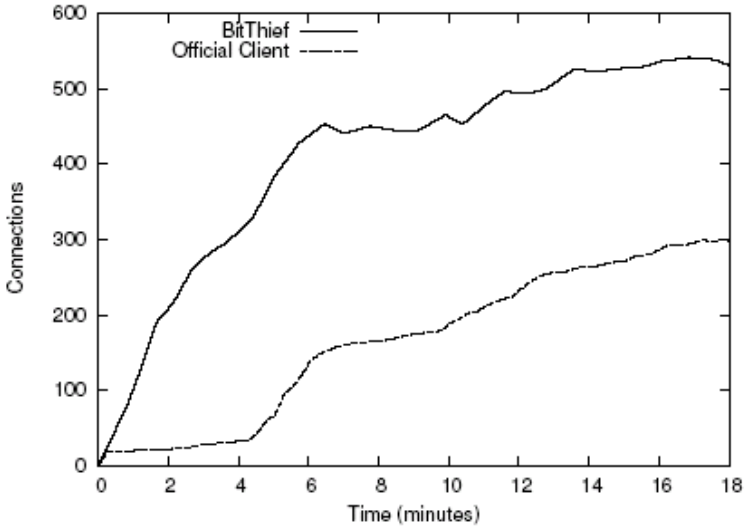

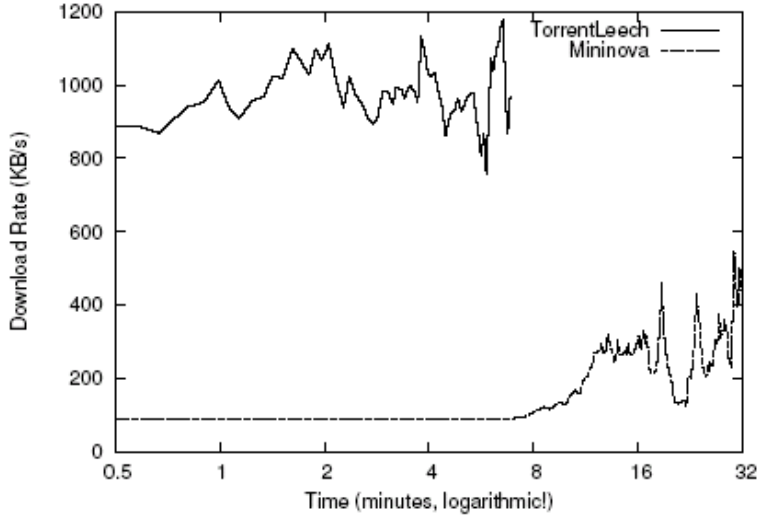

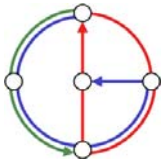

Impact of Selfish Behavior?

- Tools of game theory...
- Study of a network creation game
- Results
  - The price of anarchy can be high in large networks: $\Theta(\alpha,n)$ (tight in 1d Euclidean space)
  - Selfishness can destabilize a system even in absence of churn (no pure Nash equilibria)
  - Determining whether a given network can stabilize is NP-hard (reduction from special SAT)

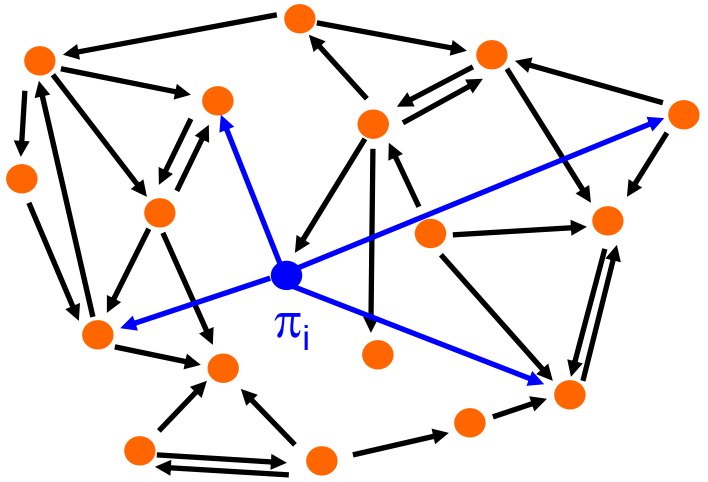

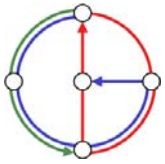

Impact of Malicious Behavior?

- Other forms of non-cooperative behavior
- Our framework allows to quantify impact
  - How resilient is system to malicious attacks?
- Results
  - Example: Virus inoculation game
  - Malicious participants are often harmful
  - If the other players are risk-averse, malicious players may actually improve the system performance!

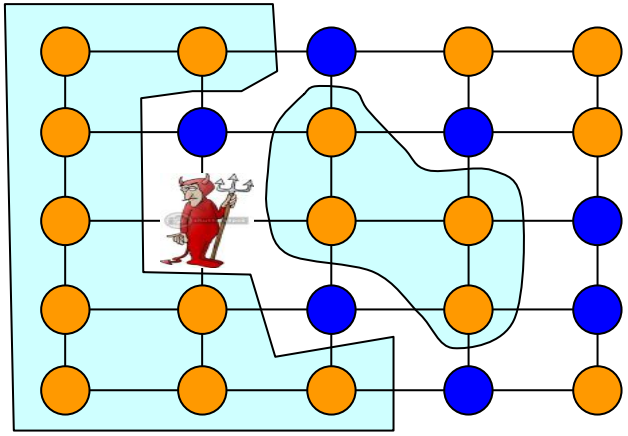

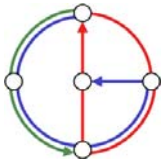

Impact of Social Behavior?

- Framework also applicable to social networks

- E.g., analysis of inoculation strategies if players care about their direct contacts

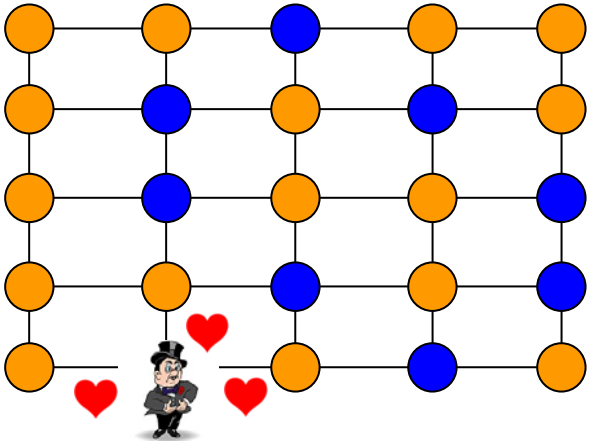

- Results
  - Example: Virus inoculation game
  - Equilibria are always at least as good as in purely selfish environments
  - Windfall increases non-monotonically in *F* (= extent to which players care about each other)
  - Wide spectrum already in simple graphs

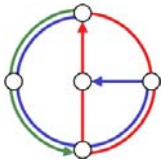

Case Study Kad (1)

- Real malicious attacks in Kad networks?
  - Large structured p2p application
  - Kad: > 1 mio simultaneous users

- Experiments with real users and real contents

- Example: Censoring contents in the Kad network

- Means:
  - Find data ID, insert node close to this ID (node insertion attack)
  - Fill up index tables of existing hosts (publish attack)
  - Plus: eclipse a peer, denial of service

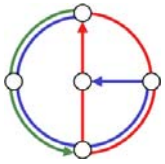

Case Study Kad (2)

- Results
  - Kad vulnerable to various attacks
  - Entire files can be removed without many resources

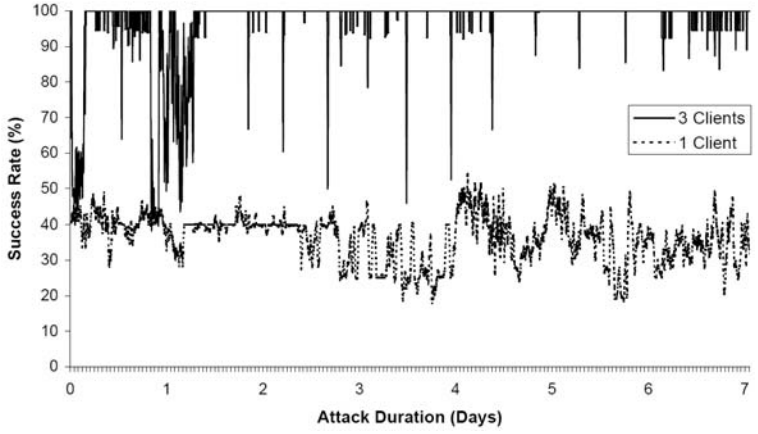

- Node insertion attack on keyword „Simpsons“
  - Also harms queries like „Simpsons Movie“ or „Simpsons Soundtrack“

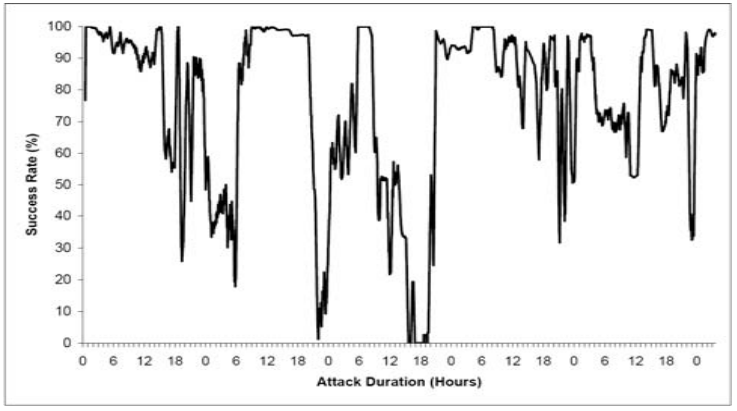

- Publish attack

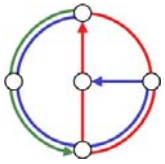

Conclusion

- Many important questions remain open!

- Dynamics
  - Efficiency?
  - Simplifications / models?
  - Self-stabilization / graceful degradation

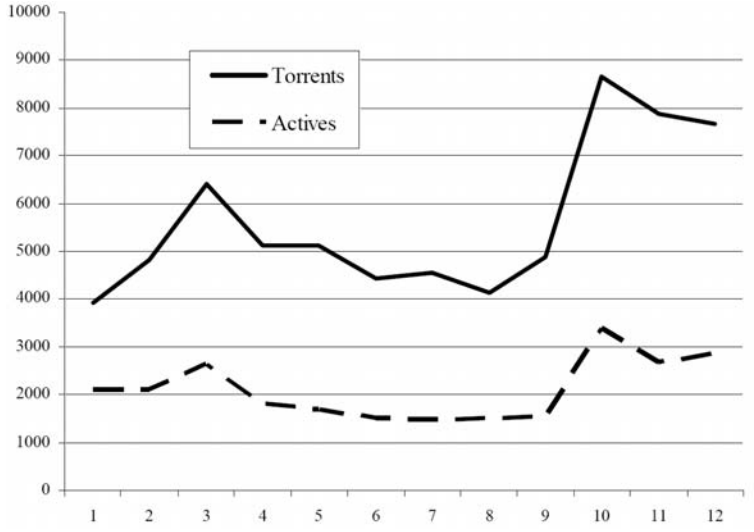

- Cooperation
  - How to solve the bootstrap problem?
  - Incentive mechanisms for multiple utility functions?
  - Are peers really selfish? BitThief stats per month...

    (BitThief: GUI, collects data, not much publicity, etc.)

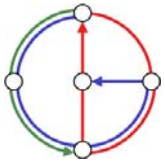

Thank you for your attention!