


Economics of P2P Computing

Stefan Schmid




ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Introduction


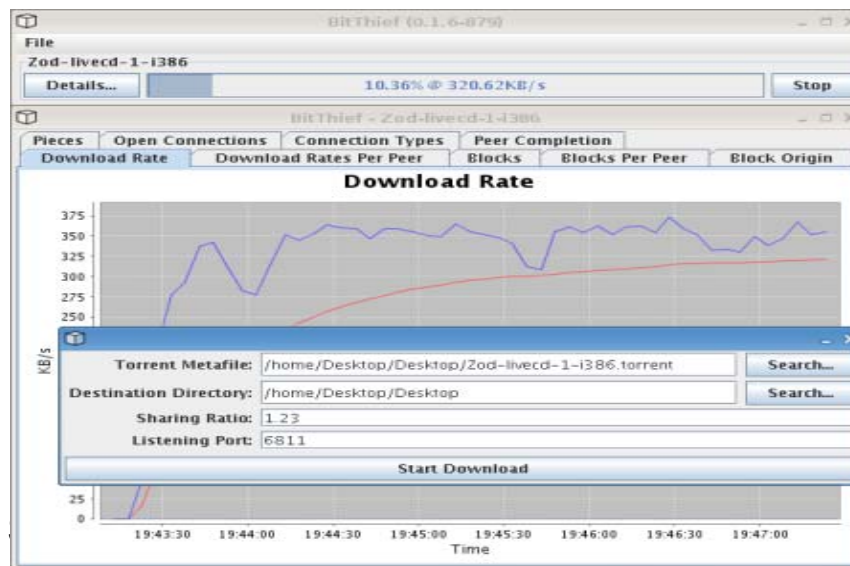
THANK YOU !

- Thank you for the invitation!! 😊
- Myself:
 - MSc in CS at **ETH Zurich**, Switzerland
 - **3rd year PhD** student of the Distributed Computing Group of Prof. Roger Wattenhofer
 - For more details, see <http://dcg.ethz.ch/members/stefan.html>
- Opportunity to meet **ECS Group!**



Why do we care?

- Several „human“ or „rational“ participants involved in P2P systems (rather than computers): **Content distributors, Users, ISPs, etc.**
- We believe that the understanding and design of **(fair) P2P economies** is relevant beyond the academic world!
- Example: Reactions to our **BitThief** client circumvents fairness mechanism of BitTorrent!
 - Tricks: see paper (HotNets'06); see also work by Shneidman, Parkes, Massoulié at SigComm'04



“Reactions”: Many People Interested in the Topic...

- > 15,000 downloads only in January, > 3,000 downloads of paper
- Much feedback...

"Anyhow, bitthief is a client which I've been waiting for so long, I mean.. bitcomet bent the rules but never really broke any of them.. that much Bitthief is an interesting client in that it openly says "fuck you, and fuck your swarm" to the torrent community. I wonder how fast this will get banned at every tracker alive. As others have said, this makes bittyrant look like a sunday school boy."

A fan!

-----Original Message-----

From: Warren Henning [mailto:warren.henning@gmail.com]

Sent: Friday, January 12, 2007 3:03 PM

To: lochert@tik.ee.ethz.ch; schmiste@tik.ee.ethz.ch; wattenhofer@tik.ee.ethz.ch

Subject: Stop distributing BitThief, you jerks!

BitTorrent is a beautiful thing and you are intentionally fucking it up by distributing software that is apparently specifically designed to attack the entire basis of the function of BitTorrent, software that serves no legitimate purpose.

Luckily it apparently requires having a JRE installed right now, and the knuckle-dragging numbskulls you've worked so hard to cater to are probably too lazy to install that.

You people piss me off.

Warren Henning

Not a fan!

by [troydoodle7](#) on 1/05/07 + 21 diggs

This will kill off bitorrent far better than the riaa ever could. I wonder how long it will take before the other clients start blocking this apps from downloading off them?


A rumor...

by [korimickster](#) on 1/05/07 + 18 diggs

"This will kill off bitorrent far better than the riaa ever could."


Absolutely... and much, much faster.

Wanna try? ☺
dca.ethz.ch/projects/bitthief/



Economic Aspects of Peer-to-Peer Computing (1)

- Participants in the distributed computations can be considered **rational / selfish**, e.g.:
- E.g., the users:
 - Users are selfish, i.e., they exploit music industry by downloading copyrighted material **for free**, or even exploit the p2p system **by not contributing anything themselves!**
 - This may be as simple as **changing the parameters** (or remove files from folder), but also the entire client can be modified (e.g., BitThief); not a big deal, only one person has to do it!
- E.g., the **content distributors** save money by using the users' upload bandwidth or other resources (and **at the cost of ISPs?**)
- Etc.!




Economic Aspects of Peer-to-Peer Computing (2)

- Goal of P2P system designer: How to achieve **incentive-compatibility**?
 - How to make peers act according to the protocol?
 - How to make peers contribute resources?
- **Difficult task!** Today, hardly any system achieves this goal!
 - But still: The systems seem to work! Why? Future?
- How to reason about / tackle these problems?



Field of Game Theory and Algorithmic Mechanism Design!




Economics of Peer-to-Peer Computing

- **Game theory** gives insights into how rational players act in distributed systems
 - Answers the question: **Robust to selfishness?**
- If a game theoretic analysis indicates that the presence of selfish players renders the system inefficient compared to a optimal solution consisting of obedient players only (large **price of anarchy**), appropriate **mechanisms have to be designed**.



VS



Talk Outline




- In this talk, I will present a sample game-theoretic analysis of a **P2P network creation game**.
 - i.e., inefficiency & stability of networks with selfish peers
- Generally, game theory reveals whether a given system is robust to a set of **selfish players**. In practice, there may also be **malicious or irrational players**, e.g., the RIAA, who try to minimize the system performance.
- Second part of talk: Sample analysis of a game with both **selfish and malicious players!**



VS




VS



Talk Outline




- Network Creation Game
- Malicious Players in a Virus Inoculation Game



Talk Outline




- **Network Creation Game**
- Malicious Players in a Virus Inoculation Game



Selfishness in P2P networks


- Collaboration of peers is essential in P2P networks!
 - Each peer should contribute some resources
 - **Selfishness** can cause problems!
- Nothing at all: The Free-Riding Problem
 - **downloading** without uploading
 - Using storage without providing disk-space...
- Sample game: **selfish neighbor selection** in **unstructured P2P systems**
- Goals of selfish peer:
 - It wants to have small latencies, quick look-ups
 - It wants to have small neighbor maintenance overhead



What is the impact on the P2P topologies?



Efficiency

Stability



Model – The “Locality Game”

- Model inspired by **network creation game** [Fabrikant et al, PODC'03]
 - Sparked much future research, e.g., study of **bilateral links** (both players pay for link) rather than unilateral by Corbo & Parkes at PODC'05
- n peers $\{\pi_0, \dots, \pi_{n-1}\}$ distributed in a **metric space**
 - defines distances (\rightarrow 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


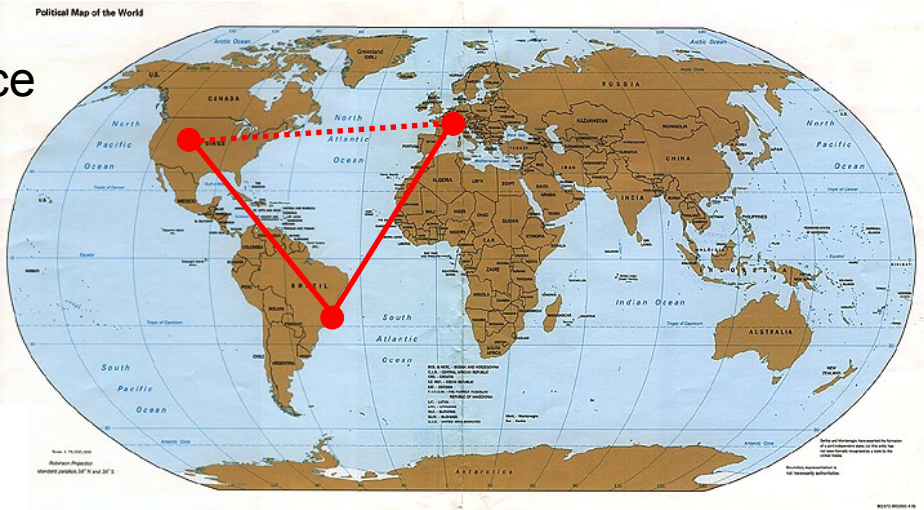
- (1) Maintain a small number of neighbors only (**out-degree**)
- (2) Small **stretches** to all other peers in the system

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!**

- Systems with many small, fast lookups \rightarrow small α
- Storage systems with large files \rightarrow large α



Model – Social Cost




- **Social Cost** is the sum of costs of individual peers
- System designer wants small social costs
- **Social Optimum (OPT)**
 - Topology with minimal social cost of a given problem instance
 - “**topology formed by collaborating peers**”!



- What topologies do *selfish peers* form?

→ Concepts of **Nash equilibrium** and **Price of Anarchy**



Model – Price of Anarchy



- **Nash equilibrium**

- “Result” of selfish behavior → “**topology formed by selfish peers**”
- Network where no peer can reduce its costs by changing its neighbor set

- **Price of Anarchy**


- Captures the impact of selfish behavior by comparison with optimal solution: ratio of social costs



What is the Price of Anarchy of our “Locality Game”?

Is there actually a Nash equilibrium...?


$$PoA := \max_I \frac{NASH(I)}{OPT(I)}$$





Related Work



- The “*Locality Game*” is inspired by the “*Network Creation Game*”
- Differences:
 - In the Locality Game, nodes are located in a **metric space**
 - Definition of stretch is based on metric-distance, not on hops!
 - The Locality Game considers **directed links**
 - Yields new optimization function



Overview




Analysis: Lower Bound for Social Optimum?




Your turn! 😊

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




- Compute upper bound for PoA => need lower bound for social opt
- OPT > ?
 - Sum of all the peers' individual costs must be at least?
 - Total link costs > ? (Hint: directed connectivity)
 - Total stretch costs > ?



Analysis: Social Optimum

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



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


- Now: Upper Bound for NE? In any Nash equilibrium, **no stretch exceeds $\alpha+1$**
 \rightarrow otherwise it's worth connecting to the corresponding peer
- A peer can have at most $n-1$ outgoing links!

$$\text{NASH} \in \mathcal{O}(\alpha n^2)$$

Really...?

Can be bad for large α

$$\text{Price of Anarchy} \in \mathcal{O}(\min\{\alpha, n\})$$





Analysis: Price of Anarchy (Lower Bound)




- Price of anarchy is **tight**, i.e., it also holds that

The Price of Anarchy is $PoA \in \Omega(\min\{\alpha, n\})$

- This is already true in a **1-dimensional Euclidean space**:



Analysis: 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:	$1/2$	α	$1/2 \alpha^2$	α^3	$1/2 \alpha^4$...	$1/2 \alpha^{i-2}$	α^{i-1}	$1/2 \alpha^i$...	$1/2 \alpha^{n-1}$


To prove:

- (1) “is a selfish topology” = instance forms a **Nash equilibrium**
- (2) “has large costs compared to OPT”
= the **social cost** of this instance is $\Theta(\alpha n^2)$


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



Analysis: 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**
 - No link to the right can 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




Analysis: 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 $> \alpha/2$



Analysis: Price of Anarchy (Lower Bound)



- Price of anarchy is **tight**, i.e., it holds that

The Price of Anarchy is $\text{PoA} \in \Theta(\min\{\alpha, n\})$


- This is already true in a **1-dimensional Euclidean space**
- Discussion:

Need no incentive mechanism

→ For small α , the Price of Anarchy is small!

Need an incentive mechanism

→ For large α , the Price of Anarchy grows with n !



What about stability...?




- We have seen:

Unstructured p2p topologies may deteriorate due to selfishness!



- What about other effects of selfishness...?
- ... selfishness can cause even more harm...!



Even in the absence of churn, mobility or other sources of dynamism, the system may never stabilize (i.e., P2P system may never reach a Nash equilibrium)!




Overview




What about stability...?

- Consider the following simple toy-example
- Let $\alpha=0.6$ (for illustration only!)
- 5 peers in **Euclidean** plane...
- ... 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:




What about stability...?




- Example sequence:

Again initial situation

→ Changes **repeat forever!**




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




Stability for general α ?


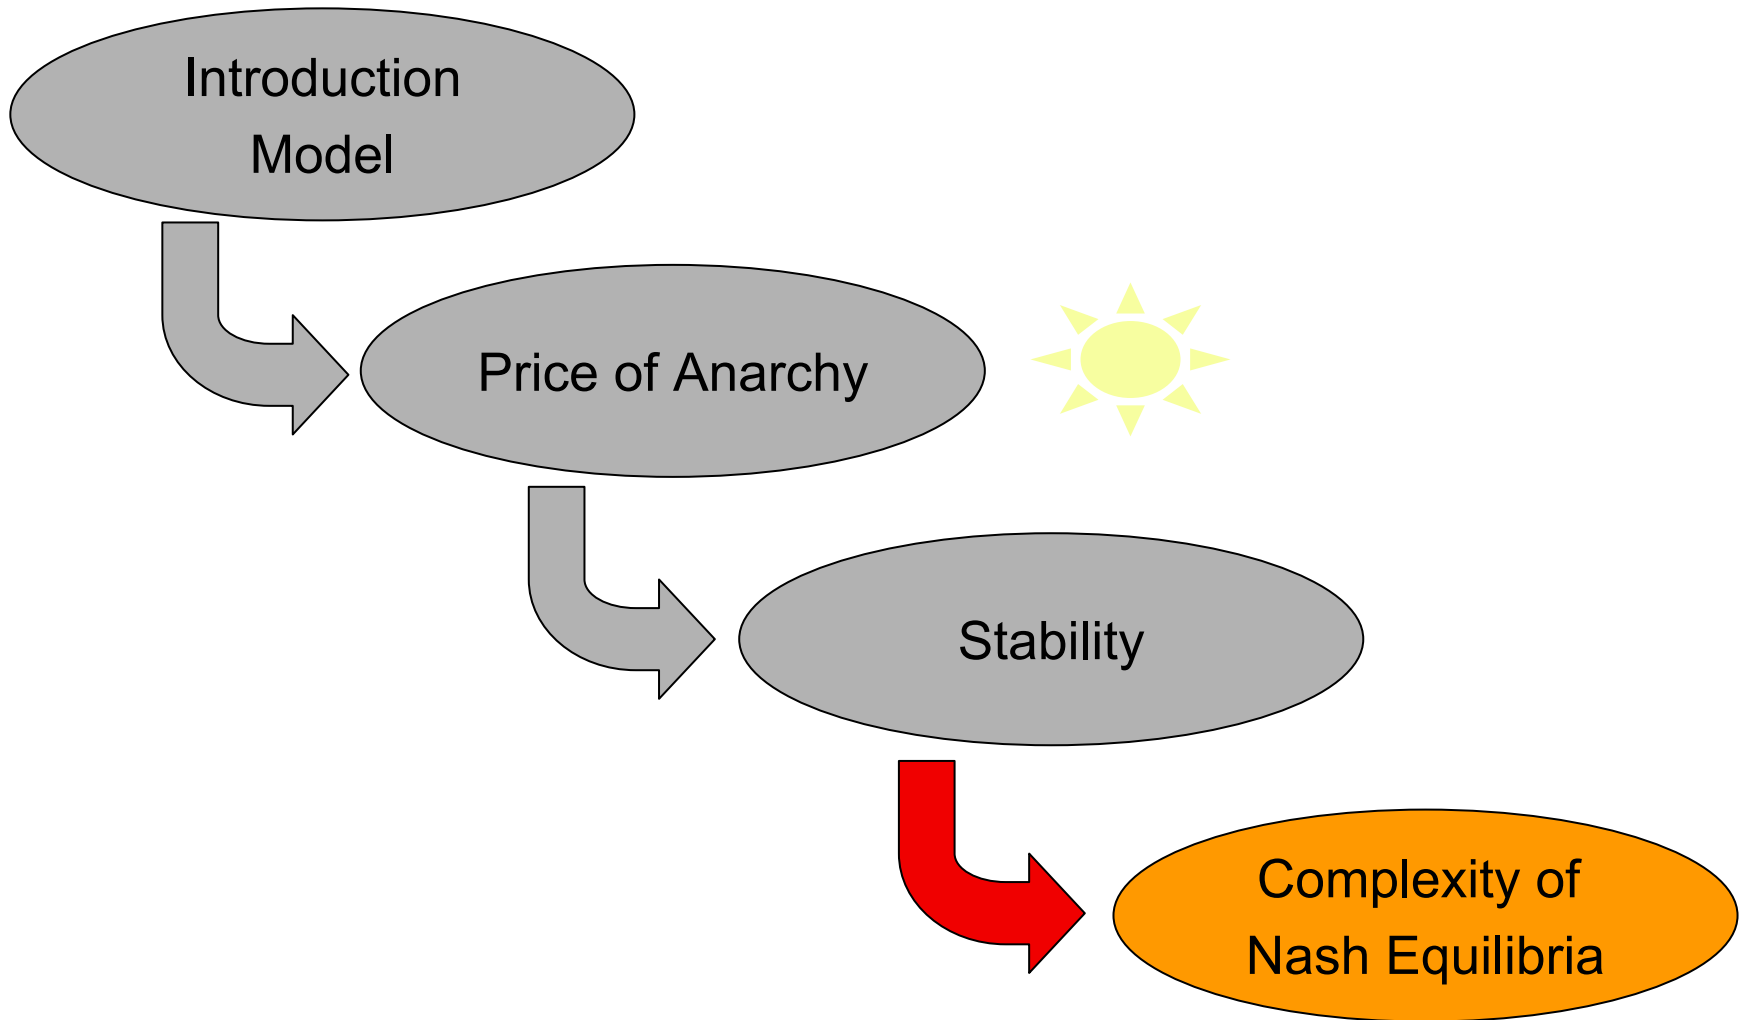
- So far, only a result for $\alpha=0.6$
- With a trick, we can generalize it to **all magnitudes of α**
- Idea, replace one peer by a **cluster of peers**
- Each **cluster has k peers** \rightarrow The network is instable for $\alpha=0.6k$
- Trick: between clusters, at most **one link** is formed (**larger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



$\delta \dots$ arbitrary small number



Overview



Complexity issues...




- **Selfishness can cause instability!**
(even in the absence of churn, mobility, dynamism....)
- Can we (at least) **determine** whether a given **P2P network is stable?**
(assuming that there is no churn, etc...)

→ What is the **complexity of stability**...???




**Determining whether a
P2P network has a (pure)
Nash equilibrium is NP-hard!**




Complexity of Nash Equilibrium

- Idea: **Reduction from 3-SAT** in CNF form (each clause has 3 literals)
 - Proof idea: Polynomial time reduction: SAT formula \rightarrow distribution of nodes in metric space
 - If **each clause** is satisfiable \rightarrow there exists a Nash equilibrium
 - Otherwise, it does not.
 - As reduction is fast, determining the complexity **must also be NP-hard**, like 3-SAT!
 - (Remark: We need that each variable in at most 3 clauses, still NP hard.)
- Arrange nodes as below
 - For each clause, **our old instable network!** (cliques \rightarrow 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 clusters appear in the clause!





Stefan Schmid @ Ha





Complexity of Nash Equilibrium

- Main idea: The **literal clusters help to stabilize!**
 - **short distance from Π^c (by construction), and maybe from Π_z**
- The clue: Π_z can only connect to one literal per variable!
- If a clause has **only unsatisfied literals**, the paths become too large and the corresponding clause becomes instable!
 - Otherwise the network is stable, i.e., there exists a Nash equilibrium.





Complexity of Nash Equilibrium



Complexity of Nash Equilibrium



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



Complexity of Nash Equilibrium





- 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 $(\Pi_y, \Pi_z, \text{Clause})$ does not converge by itself...





Complexity of Nash Equilibrium

- Each node-set Π^c is **connected** to those literals that are in the clause (not to other!)
- if Π_z has link to $\text{not}(x_1)$, there is a “**short-cut**” to such clause-nodes, and C_2 is **stable**
- But not to other clauses (e.g., $C_1 = x_1 \vee x_2 \vee \text{not}(x_3) \dots$)




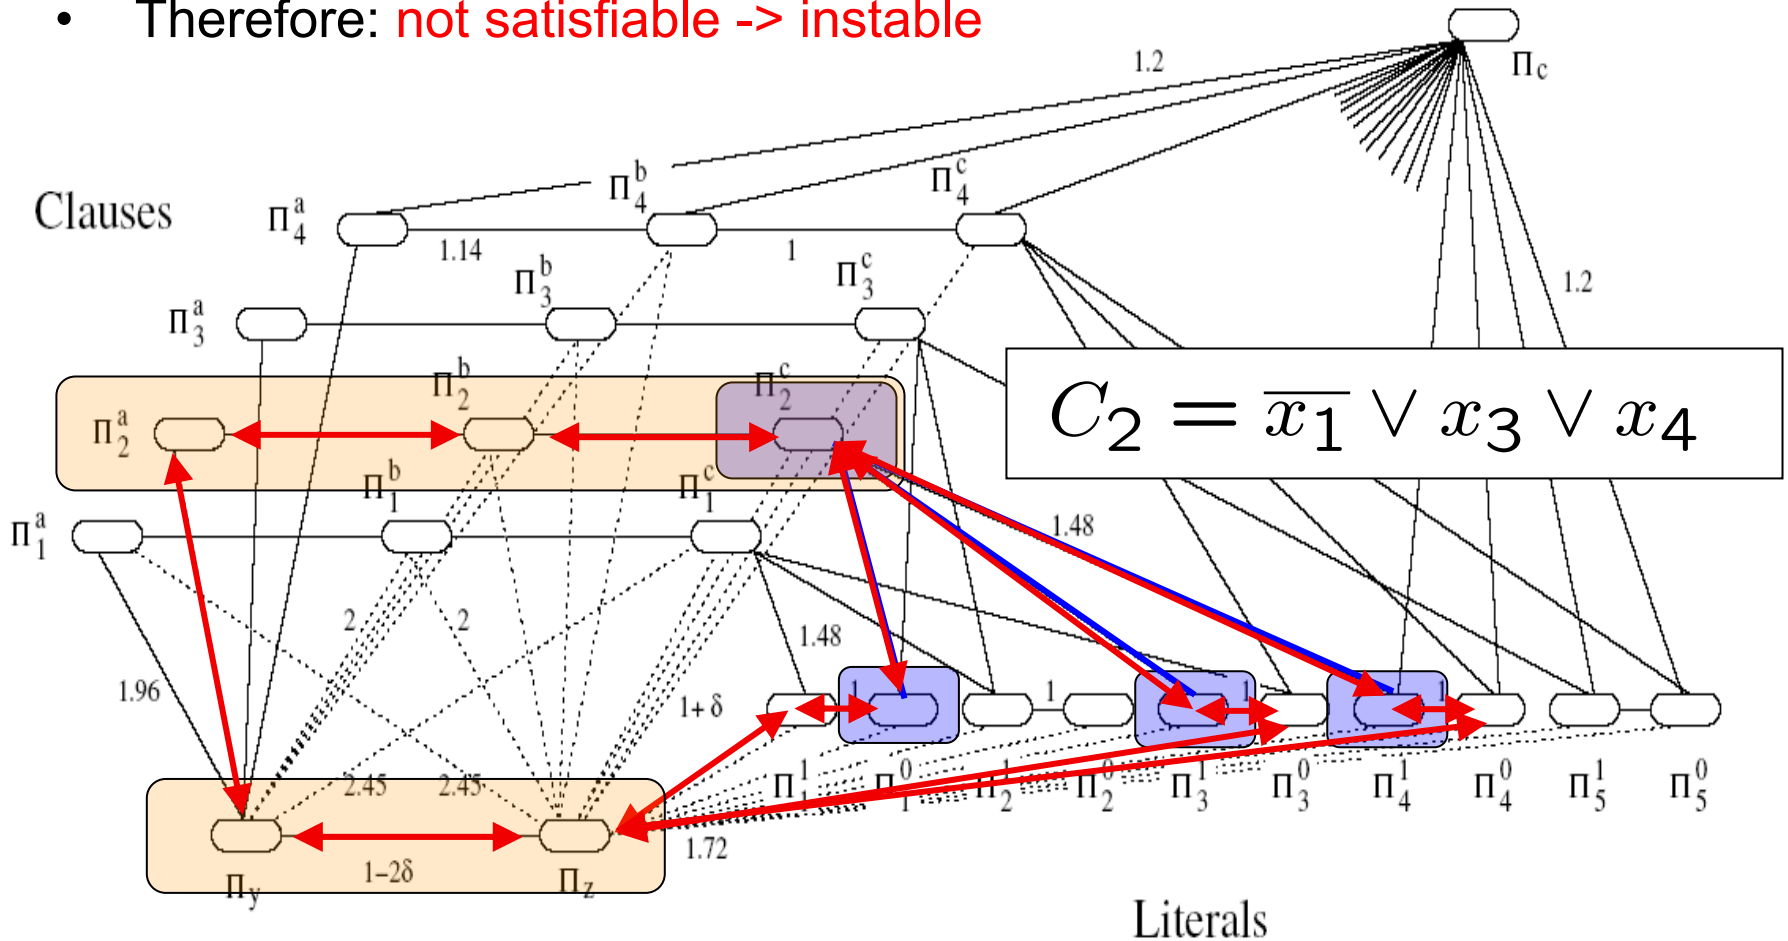
Complexity of Nash Equilibrium

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




Complexity of Nash Equilibrium

- If there is **no** such “short-cut” to a clause, the clause remains **instable!**
- Therefore: **not satisfiable -> instable**



Complexity of Nash Equilibrium

- 3-SAT instance is satisfiable \rightarrow every clause is stable



The Topologies formed by Selfish Peers

- Selfish **neighbor selection** in **unstructured P2P systems**
- Goals of selfish peer:
 - (1) Maintain links only to a few neighbors (small **out-degree**)
 - (2) Small **latencies** to all other peers in the system (fast lookups)




What is the impact on the P2P topologies?

Price of Anarchy $\in \Theta(\min\{\alpha, n\})$


Determining whether a P2P network has a (pure) Nash equilibrium is NP-hard!

Even in the absence of churn, mobility or other sources of dynamism, the system may never stabilize



Future Directions – Open Problems


- Nash equilibrium **assumes full knowledge** about topology!
 - this is of course unrealistic
 - incorporate aspects of **local knowledge** into model
- Current model does not consider **routing** or **congestion** aspects!
 - also, why should every node be connected to every other node?
(i.e., **infinite costs if not**? Not appropriate in Gnutella or so!)
- **Mechanism design**: How to guarantee stability/efficiency..?
- More **practical**: what is the parameter α in real P2P networks?
- Lots more:
 - Algorithms to compute **social opt** of locality game?
 - Quality of **mixed** Nash equilibria?
 - Is it also hard to determine complexity for **Euclidean metrics**?
 - Computation of other equilibria
 - Comparisons to unilateral and bilateral games, and explanations?



Talk Outline




- Network Creation Game
- Malicious Players in a Virus Inoculation Game



Talk Outline




- Network Creation Game
- **Malicious Players in a Virus Inoculation Game**



Motivation




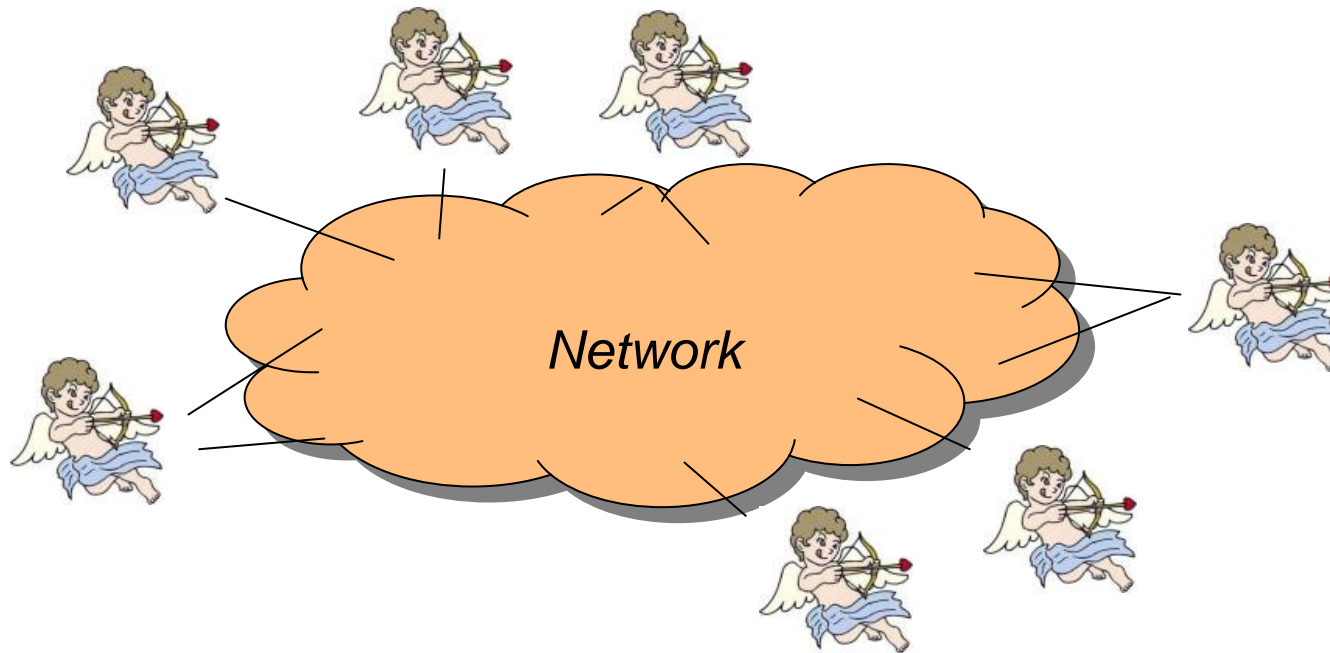
- So far: Selfishness is a threat for peer-to-peer systems!
 - **Game theory** and **algorithmic mechanism design** are tools that help to understand and solve the problem.
- But: Users may also be **malicious** rather than selfish!
 - E.g., the RIAA would like to harm p2p file sharing systems! (e.g., minimize ist performance)
- We have proposed to study the impact of malicious players on the performance of a distributed system.



Modeling Distributed Systems


- One possibility to model a distributed system:

all participants are benevolent!




Selfishness in Networks

- Alternative: Model **all participants as selfish**
→ e.g. our p2p network creation game




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**



What is the impact of **malicious players on selfish systems**...?



“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:


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



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**).



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-loving...
Rational...

Nothing...,
Number of malicious players...
Distribution of malicious players...
Strategy of malicious players...



“Byzantine Game Theory”


- **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...?




- 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 (for simplicity)
- Each peer or node can choose whether to **install anti-virus software**
- Nodes who install the software are **secure** (costs 1) ●
- Virus spreads from a 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 want to minimize its expected cost!



Related Work:

- The VIG was first studied by Aspnes et al. [SODA'05]
- Approximation algorithm
 - General Graphs
 - No malicious players




Virus Inoculation Game

- 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 i th 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:

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




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!
- 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!
- They assume everyone is selfish and rational
- How much can the social cost deteriorate...?

- Simple **upper bounds**:


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

- 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: all components **without** any

malicious node $\rightarrow Cost_{inf}^{Byz} \in O(s)$


(similar to analysis of PoA!)



- On the other hand: a component i with $b_i > 0$

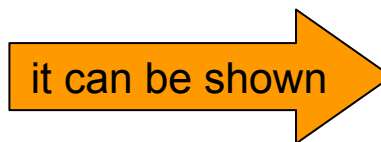
malicious nodes: $\sum_i b_i = b$

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




$$k_i \leq (b_i + 1) \cdot \frac{n}{L} + b_i$$

$$l_i \leq (b_i + 1) \cdot \frac{n}{L}$$



$$Cost_{inf}^{Byz} \in O\left(\frac{b^2 n}{L}\right)$$



Price of Malice – Oblivious case



- Social cost is upper bounded by $O\left(s + \frac{b^2 n}{L}\right)$


- The **Price of Byzantine Anarchy** is at most for $b < L/2$

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

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

- The **Price of Malice** is at most

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



Oblivious Case Lower Bound




- In fact, these bounds are **tight!**

→ bad example: components with **large surface**

(Many inoculated nodes for given component size

=> bad NE! All malicious players together,

=> **one large attack component** => large BNE)



→ this scenario is a Byz Nash Eq.

in the oblivious case.

→ With prob. $((b+1)n/L+b)/n$,


infection starts at an insecure or a malicious node of attack

component of size $(b+1)n/L$

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

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

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



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:

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

This was Price of Anarchy...


- Price of Malice is:

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

- 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!**

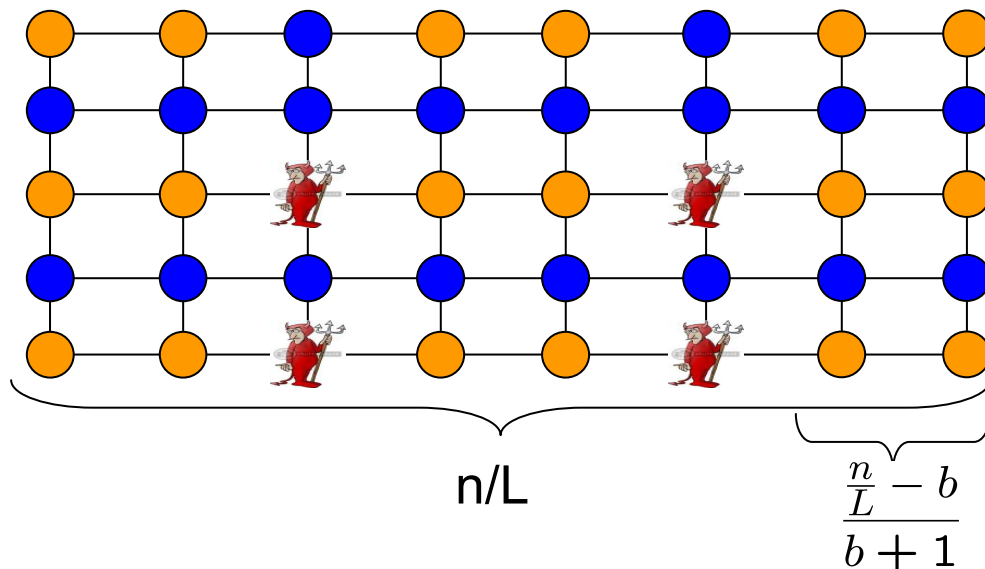


Price of Malice – Non-oblivious case




- Selfish nodes **know** the **number of malicious agents b**
- 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 agents can be removed while attack component size is at most n/L !



Conclusion



- Insights from Byzantine game theory??
- Game-theoretic analysis
 - Large price of anarchy -> **need incentive mechanism**
- Byzantine game theory
 - Large price of malice -> need to do something! But what?
 - E.g., **keep malicious players off** from the beginning!



Future Work




- Plenty of open questions and future work!
- Virus Inoculation Game



- 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**...?


- The **Price of Malice** in other scenarios and games
- Routing, caching, etc...
- Can we **use Fear-Factor** to improve networking...?



Conclusion



- **Selfishness / Non-cooperation** are important challenges in P2P computing!
- In order to build successful systems in practice, it is crucial to understand the **incentives of the different participants!**
- There are **other challenges in P2P computing** which I am interested in.
- For example, P2P systems consist of desktop machines which join the network for a short time only! -> The system must be fully functional in spite of high dynamics (**churn**)!



Overview of Peer-to-Peer Projects (1)



Locality-Awareness, Security, ...

(P2P'06, SRDS'06)





Algorithmic Challenges

Dynamics, Churn
(IPTPS'05, IWQoS'06,
Wicon'06, HiPC'06)



Non-Cooperation
(IPTPS'06, NetEcon'06,
PODC'06, HotNets'06)




Overview of Peer-to-Peer Projects (2)



BitThief (Free-riding client for BitTorrent)

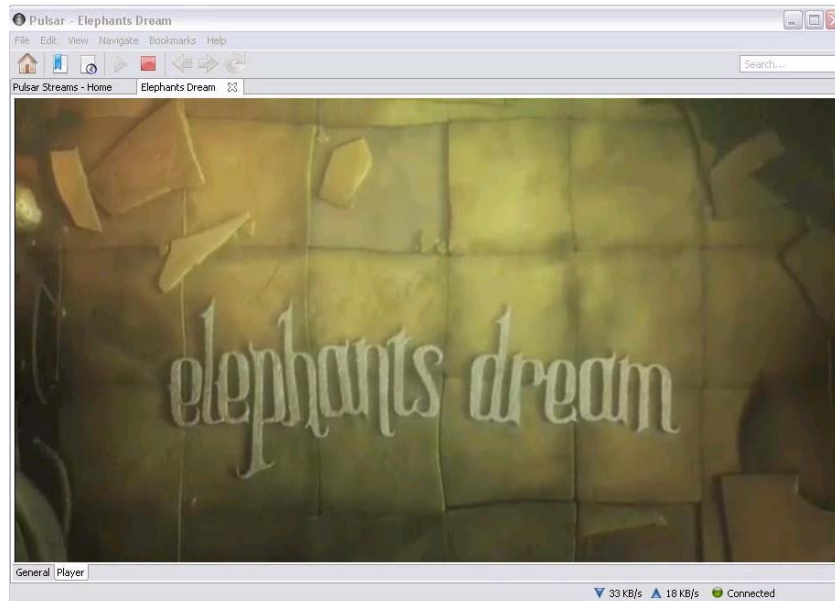


Wuala
(Distributed
P2P Storage and
Social Networking)




*Pulsar**
(Live P2P
Streaming)

* Wanna try? Today and tomorrow, we are streaming the IPTPS workshop on P2P systems (Bellevue, WA)!



07



The Last Slide

THANK YOU !

- Some of our work at DCG, in particular
 - BitThief: A free-riding BitTorrent Client
 - A Network Creation Game
 - Malicious Players in a Virus Inoculation Game
- Questions and Feedback?
- Your work? Discussion?

Literature:

- [1] L. Locher, P. Moor, S. Schmid, R. Wattenhofer: *Free Riding in BitTorrent is Cheap*, HotNets 2006.
- [2] T. Moscibroda, S. Schmid, R. Wattenhofer: *On the Topologies Formed by Selfish Peers*, PODC 2006.
- [3] T. Moscibroda, S. Schmid, R. Wattenhofer: *When Selfish Meets Evil – Byzantine Players in a Virus Inoculation Game*, PODC 2006.

