# Mechanism Design Without Payments for Throughput Maximization

Thomas Moscibroda (MSR)
Stefan Schmid (TUM)

TECHNISCHE UNIVERSITÄT MÜNCHEN

*INFOCOM 2009*

# Distributed Systems 2008/9

Wireless: Many mobile phones today have WLAN

P2P: Olympic games live-broadcast over peer-to-peer networks

„Social“ networks: *Facebook*, *Xing*, *Twitter*...

E.g. US elections 2008: Obama makes extensive use of Internet technologies



# Different Stakeholders

- Distributed systems often consist of different participants and stakeholders
  - e.g., BitTorrent swarm similar to a small society
  - e.g., different strategies define the set of ISP routing rules

- Many distributed systems rely on collaboration and resource contributions

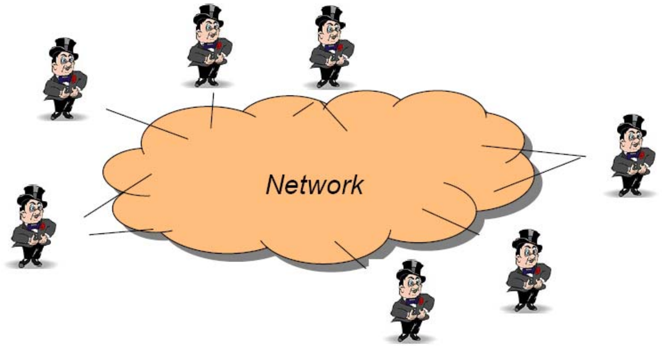


**Incentives for „efficient" collaboration in selfish milieus?**

# The Socio-economic Complexity (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.“**

Triggered much research!

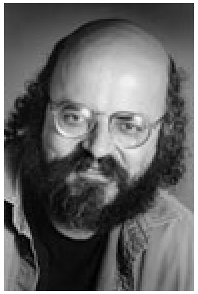

# Monetary Mechanism Design?

- Many economic solutions rely on monetary payments

- This may not be realistic in open distributed systems
  - Open: e.g., peer-to-peer system, wireless network, etc.
  - Anonymous participants
  - No currency available
  - ...

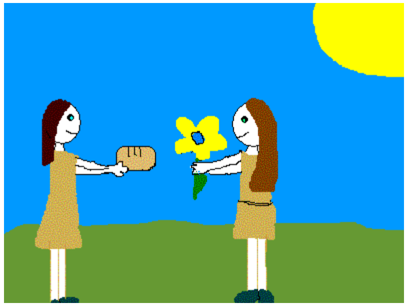

- A *„hot topic“*
  - Distributed mechanism design
  - Solutions with punishments?
  - Money burning mechanisms? (pay with e.g., bandwidth?)
  - Tit-for-tat solutions (e.g., BitTorrent)...

[Interesting work by Papadimitriou, Parkes, Roughgarden, Shenker, etc.]

# Our Research

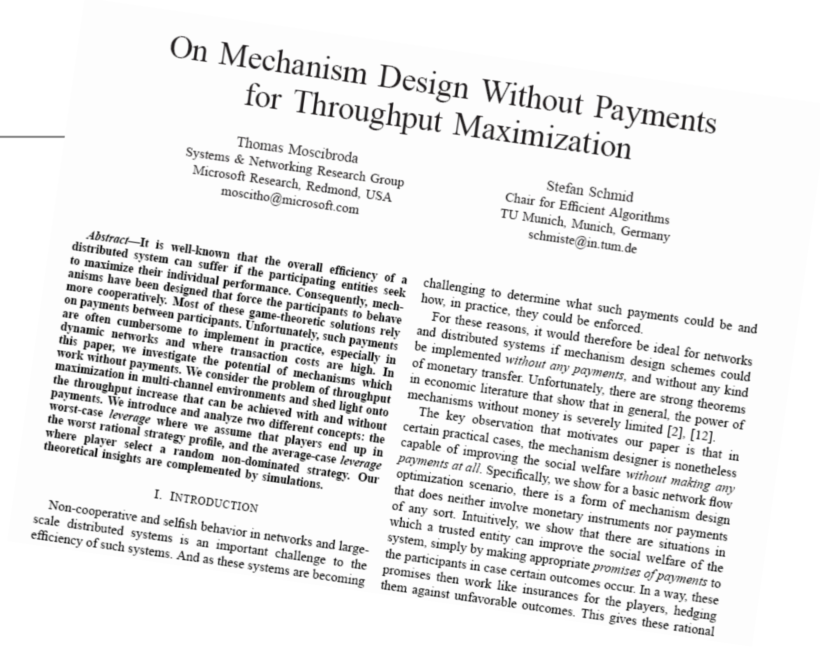

## On Mechanism Design Without Payments for Throughput Maximization

Thomas Moscibroda
Systems & Networking Research Group
Microsoft Research, Redmond, USA
moscitho@microsoft.com

Stefan Schmid
Chair for Efficient Algorithms
TU Munich, Munich, Germany
schmiste@in.tum.de

***Abstract*—It is well-known that the overall efficiency of a distributed system can suffer if the participating entities seek to maximize their individual performance. Consequently, mechanisms have been designed that force the participants to behave more cooperatively. Most of these game-theoretic solutions rely on payments between participants. Unfortunately, such payments are often cumbersome to implement in practice, especially in dynamic networks and where transaction costs are high. In this paper, we investigate the potential of mechanisms which work without payments. We consider the problem of throughput maximization in multi-channel environments and shed light onto the throughput increase that can be achieved with and without payments. We introduce and analyze two different concepts: the worst-case *leverage* where we assume that players end up in the worst rational strategy profile, and the average-case *leverage* where player select a random non-dominated strategy. Our theoretical insights are complemented by simulations.**

### I. Introduction

Non-cooperative and selfish behavior in networks and large-scale distributed systems is an important challenge to the efficiency of such systems. And as these systems are becoming challenging to determine what such payments could be and how, in practice, they could be enforced.

For these reasons, it would therefore be ideal for networks and distributed systems if mechanism design schemes could be implemented *without any payments*, and without any kind of monetary transfer. Unfortunately, there are strong theorems in economic literature that show that in general, the power of mechanisms without money is severely limited [2], [12].

The key observation that motivates our paper is that in certain practical cases, the mechanism designer is nonetheless capable of improving the social welfare *without making any payments at all*. Specifically, we show for a basic network flow optimization scenario, there is a form of mechanism design that does neither involve monetary instruments nor payments of any sort. Intuitively, we show that there are situations in which a trusted entity can improve the social welfare of the system, simply by making appropriate *promises of payments* to the participants in case certain outcomes occur. In a way, these promises then work like insurances for the players, hedging them against unfavorable outcomes. This gives these rational

Observation: Cooperation can be fostered through a trusted entity who makes monetary *promises* to individual players for certain outcomes.

Sometimes, cooperation / outcomes can be improved significantly with only a small amount of money, sometimes even for free (i.e., with mere „creditability“).

# Our Research

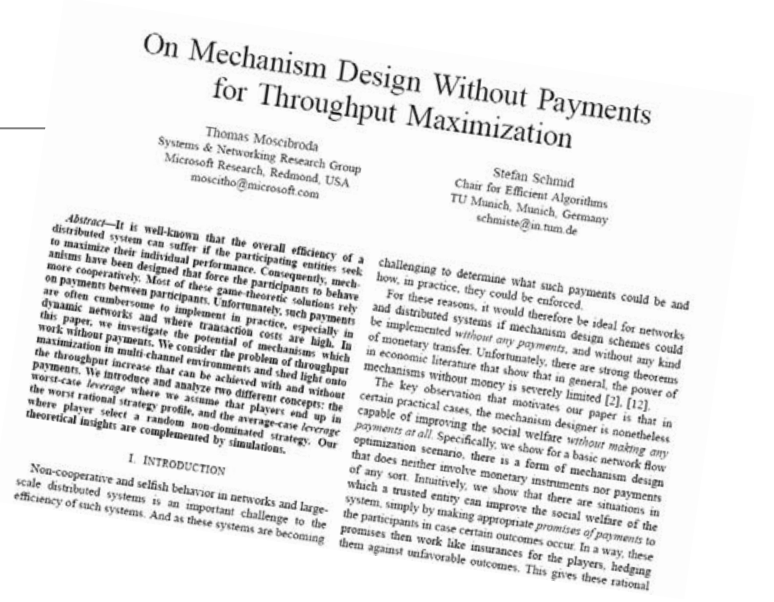

On Mechanism Design Without Payments for Throughput Maximization

Thomas Moscibroda
Systems & Networking Research Group
Microsoft Research, Redmond, USA
moscitho@microsoft.com

Stefan Schmid
Chair for Efficient Algorithms
TU Munich, Munich, Germany
schmiste@in.tum.de

***Abstract*—It is well-known that the overall efficiency of a distributed system can suffer if the participating entities seek to maximize their individual performance. Consequently, mechanisms have been designed that force the participants to behave more cooperatively. Most of these game-theoretic solutions rely on payments between participants. Unfortunately, such payments are often cumbersome to implement in practice, especially in dynamic networks and where transaction costs are high. In this paper, we investigate the potential of mechanisms which work without payments. We consider the problem of throughput maximization in multi-channel environments and shed light onto the throughput increase that can be achieved with and without payments. We introduce and analyze two different concepts: the worst-case *leverage* where we assume that players end up in the worst rational strategy profile, and the average-case *leverage* where player select a random non-dominated strategy. Our theoretical insights are complemented by simulations.**

I. INTRODUCTION

Non-cooperative and selfish behavior in networks and large-scale distributed systems is an important challenge to the efficiency of such systems. And as these systems are becoming challenging to determine what such payments could be and how, in practice, they could be enforced.

For these reasons, it would therefore be ideal for networks and distributed systems if mechanism design schemes could be implemented *without any payments*, and without any kind of monetary transfer. Unfortunately, there are strong theorems in economic literature that show that in general, the power of mechanisms without money is severely limited [2], [12].

The key observation that motivates our paper is that in certain practical cases, the mechanism designer is nonetheless capable of improving the social welfare *without making any payments at all*. Specifically, we show for a basic network flow optimization scenario, there is a form of mechanism design that does neither involve monetary instruments nor payments of any sort. Intuitively, we show that there are situations in which a trusted entity can improve the social welfare of the system, simply by making appropriate *promises of payments* to the participants in case certain outcomes occur. In a way, these promises then work like insurances for the players, hedging them against unfavorable outcomes. This gives these rational

Observation: Cooperation can be fostered through a trusted entity who makes monetary promises to individual players for certain outcomes.

**Can be regarded as an insurance.**

Sometimes, coope ... oved significantly with only a small amount of money, sometimes even for free (i.e., with mere „creditability“).

# Mechanism Design by „Creditability"

Model by Monderer and Tennenholtz

Reliable third party

- *Promises* players payments for certain outcomes
- Selfish players take this into account
- Sometimes, little or no money is sufficient to influence much!

Solution Concept

- Minimal rationality assumption: non-dominated strategies

| $U(x)$ | l | t |
|---|---|---|
| l |  | 0/4 |
| t | 4/0 | 1/1 |

Goal: Implement* **best free** profile, implement **social optimum** at minimal cost, ...

* Implement = make players choose this strategy profile

# „Creditability" in a Prisoner's Dilemma

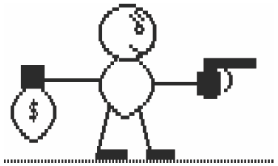

- Two robbers, can *say truth* (t) or *lie* (l)
- Matrix shows *utilities* (e.g., number of saved years in prison)

Social optimum

| $U(x)$ | l | t |
|---|---|---|
| l | 3/3 | 0/4 |
| t | 4/0 | 1/1 |

Unique Nash equilibrium
(dominant strategy profile)

# „Creditability“ in a Prisoner‘s Dilemma

- Third party makes the following promises:

| $U(x)$ | l | t |
|---|---|---|
| l | 4 / 4 | 2 / 4 |
| t | 4 / 2 | 1/1 |

| $Q(x)$ | l | t |
|---|---|---|
| l | 1/1 | 2/0 |
| t | 0/2 | 0/0 |

Now, lying is the only Nash equilibrium.
The social welfare is 6 (instead of 2),
while the third party only pays 2!
**Sometimes, better outcomes without any payments at all!**

# The Throughput Game

- A simple network: *n* players want to communicate from site A to site B over *m* parallel channels
  - Each player must choose one channel; *utility* = *throughput or demand*
  - Throughput is shared among players on same channel

Channel capacities:

$c_1,\ldots,c_m$

Player demands:

$d_1,\ldots,d_n$

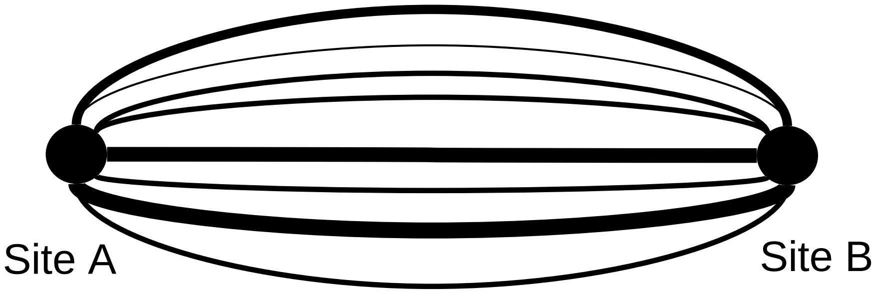


# The Throughput Game

**The goal of the mechanism designer is to distribute the selfish players „well“ among the channels!**

Channel capacities:

$c_1,\ldots,c_m$

Player demands:

$d_1,\ldots,d_n$

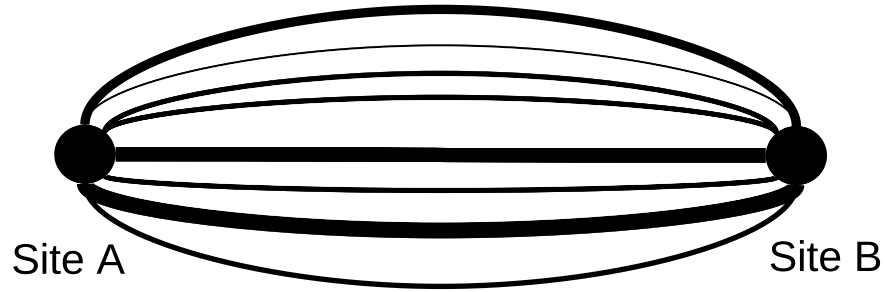


# Leverage: Benefit of Creditability

Worst-case leverage: Players choose worst non-dominated profile!

- Non-dominated: no strategy is always better, independently of other players
- Leverage = how much can be achieved with creditability relative to the social optimum (unbounded in absolute numbers)

Worst non-dominated profile

Payments

*Definition 4.1 (Worst-Case Leverage $\Phi_{wc}$):* The worst-case leverage $\Phi_{wc}(x)$ of implementing a strategy profile $x$ in a game is defined as

$$\Phi_{wc}(x) = \frac{U(x) - U(UDom_{wc}) - Q(x)}{U(OPT)}.$$

$\leq 1$ as $U(x) \leq U(OPT)$

The worst-case leverage $\Phi_{wc}$ of a game is the maximal worst-case leverage over all outcomes $x$: $\Phi_{wc} = \max_{x \in X} \Phi_{wc}(x)$.

# The Throughput Game

Worst-case leverage: Players choose worst non-dominated profile!

- Non-dominated: no strategy is always better, independently of other players („weakest
- Leverage captures
  relative to the social o                                                                                    ers)

**Take into account costs of implementation.**

*Definition 4.1 (Worst-Case Leverage $\Phi_{wc}$):* The worst-case leverage $\Phi_{wc}(x)$ of implementing a strategy profile $x$ in a game is defined as

$$\Phi_{wc}(x) = \frac{U(x) - U(UDom_{wc}) - Q(x)}{U(OPT)}.$$

The worst-case leverage $\Phi_{wc}$ of a game is the maximal worst-case leverage over all outcomes $x$: $\Phi_{wc} = \max_{x \in X} \Phi_{wc}(x)$.

# Formal Results (1)

Appetizer: Can we always implement a unique profile for free?

- Initially, non-dominated profile may not be unique!

But yes: **every Nash equilibrium can be implemented for free!**

- The throughput game always has a (pure) Nash equilibrium.

- Thus, we can e.g., third party can implement
  the best Nash equilibrium!

# Formal Results (2)

- Which profile is optimal, which profile is the best that can be 0-implemented, ...?

- What do I need to pay to implement a certain outcome?
  - Simple and fast computations, e.g., for social opt implementation.

- Bounds for the leverage

- See paper...

*Theorem 6.2:* For $D \rightarrow \infty$, the worst-case 0-leverage $\Phi^0_{wc}$ can be bounded as follows

$$\frac{\left(\sum_{i=1}^{\theta_{low}} c_i\right) - c_\gamma}{\sum_{i=1}^{\min\{n,m\}} c_i} \leq \Phi^0_{wc} \leq \frac{\left(\sum_{i=1}^{\theta_{high}} c_i\right) - c_\gamma}{\sum_{i=1}^{\min\{n,m\}} c_i},$$

where $c_\gamma$, $\theta_{low}$, and $\theta_{high}$ are as defined in Lemmas 5.3 and 6.1, respectively. For general $D$, it holds

$$\frac{\sum_{i=1}^{\theta_{low}} c_i - U(UDom_{wc})}{U(Opt)} \leq \Phi^0_{wc} \leq \frac{\sum_{i=1}^{\theta_{high}} c_i - U(UDom_{wc})}{U(Opt)},$$

with $U(UDom_{wc})$ and $U(Opt)$ as derived in Lemmas 5.1 and 5.3.

# Simulation: Impact of Number of Players (Constant)

- Study of two objectives: free profiles and social optimum
- Assume: 16 and 32 channels with capacity 100, player demand 100
- **Worst-case leverage** is almost 1 for both 0- and OPT-implementations
  - Reason: bad non-dominated profiles
  - For few players, little contention
- **Average-case leverage** (random non-dominated profile) is lower, peak approx. at n=m
  - Reason: random non-dominated profile quite good (third party cannot do much)

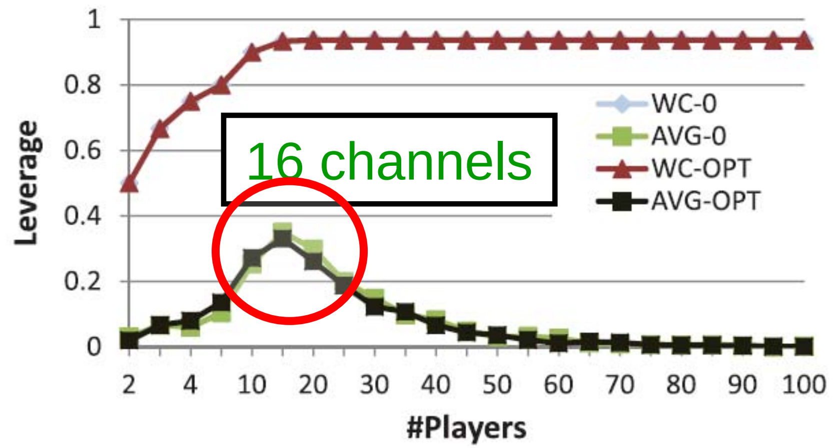


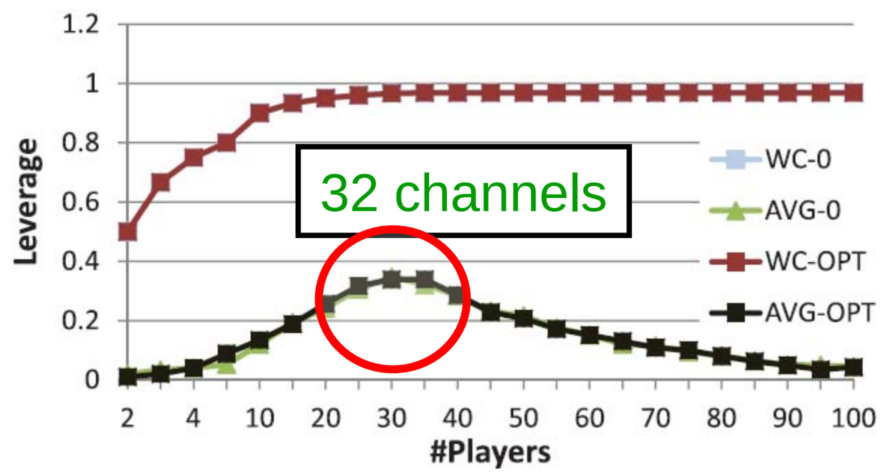

# Simulation: Impact of Number of Players (Constant)

- Study of two objectives: free profiles and social optimum
- Assume: 16 and 32 channels with capacity 100, player demand 100
- **Worst-case leverage** is almost 1 for both 0- and OPT-implementations
  - Reason: bad non-dominated profiles
  - For few players, little contention
- **Average-case leverage** (random non-dominated profile) is lower, peak approx. at n=m
  - Reason: random non-dominated profile quite good (third party cannot do much)

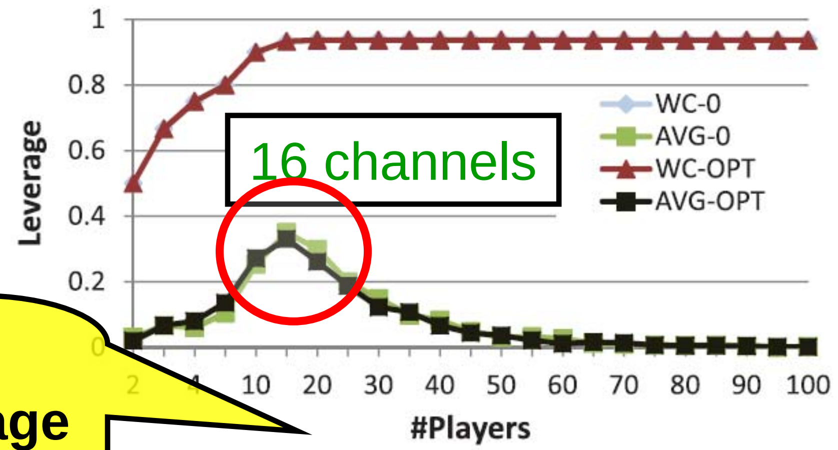


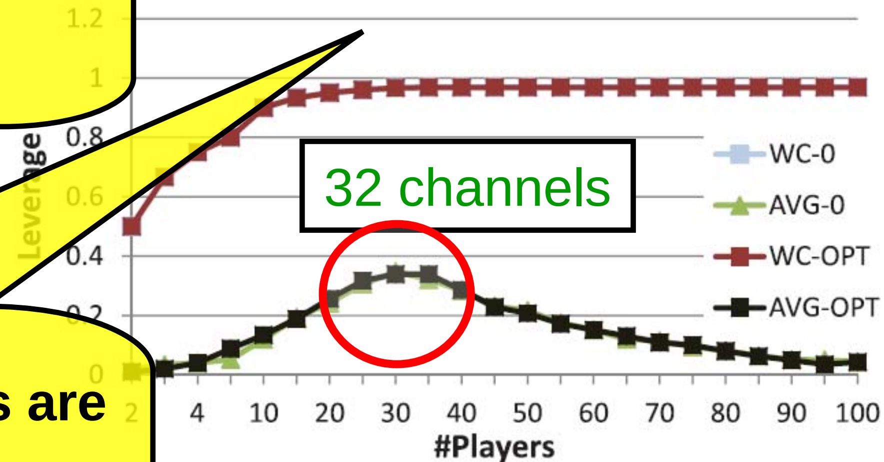


**0-leverage and OPT-leverage similar magnitude.**

**AVG: if n<<m and n>>m, players are likely to use channels well. If n=m, mechanism can help to prevent unnecessary multiple channel use.**

# Simulation: Impact of Number of Players (Uniform)

16 channels

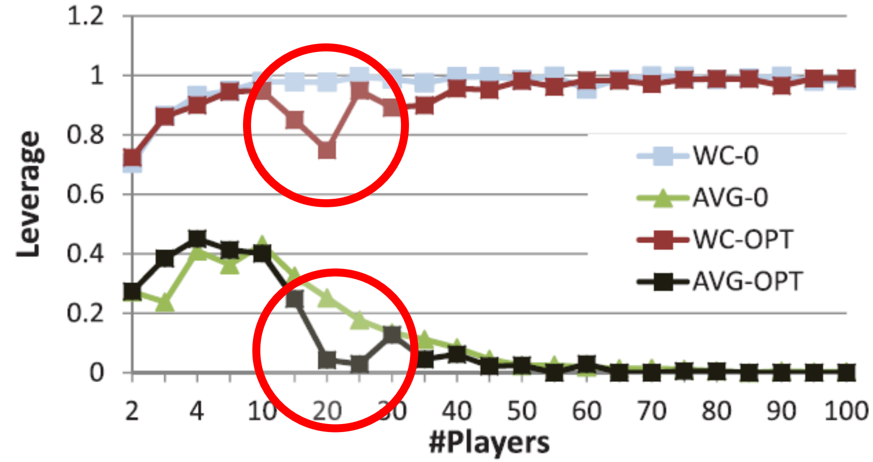


- As before, but channel capacity uniform at random from [1...100] rather than constant
  - variance: there are bad channels

- Interesting phenomenon: similar but at around n=m, there is a drop for OPT-implementations
  - Reason: Need to implement bad channels (costly and little benefit)

- It‘s generally better to implement free profiles rather than the social optimum! (taking implementation cost into account)

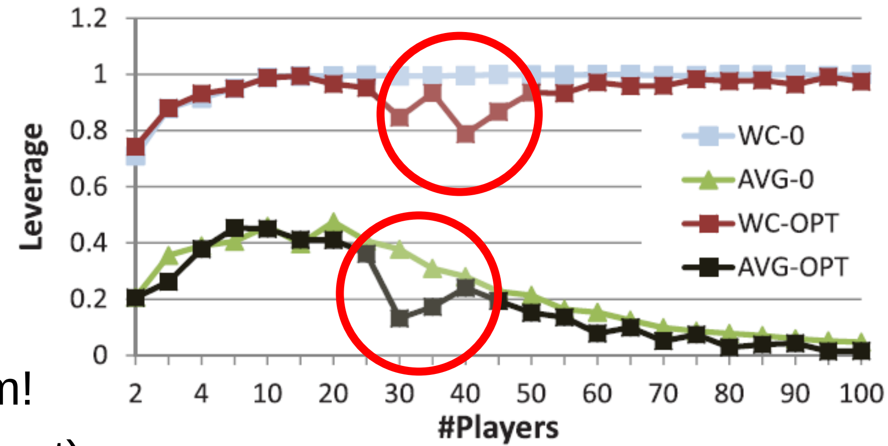


32 channels

# Simulation: Impact of Number of Players (Uniform)

**Reason: In the *social optimum*, in order to make players use the remaining weak channels, large payments are required!**

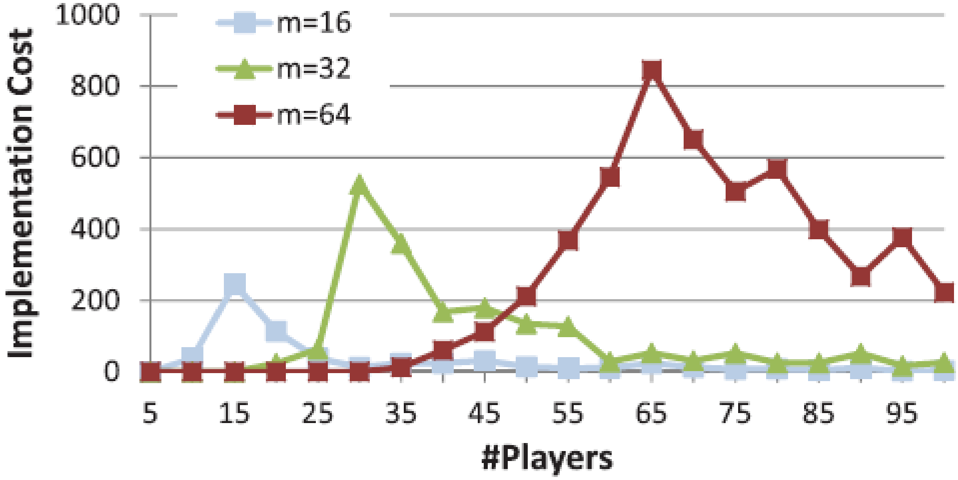


- For n<<m, it‘s cheap to put players on good channels
- For n>>m, it‘s cheap again, as channels are shared by many

# Additional Simulation Results in the Paper

**What is the impact of the players' demands?**

Max leverage for W=1, namely 1/e.

**What is the „impact of scale" W, i.e., if $n = W \cdot m$?**

**What is the leverage at saturation throughput? (Fixed load, but different number of players / demands.)**

Many phenomena and bounds on the leverage can also be explained formally (see paper for derivations).

# Conclusion

- Alternative solutions to mechanism design with payments?

- This paper: Creditability (= insurance?) can help „for free“.
  - Sometimes better than implementing social optimum

- New insights may have implications for other money-less markets
  - E.g., due to ethical or institutional considerations
    (e.g., market for *organ donations*)

- Open questions:
  - In which settings is such a insurance realistic?
    (Not only money problems!)
  - dynamic games, risk-averse behavior, etc.
  - general implementations? (beyond zero and optimal)

# Muito obrigado!

# Nash Equilibria can be implemented for free!

- Two player example:

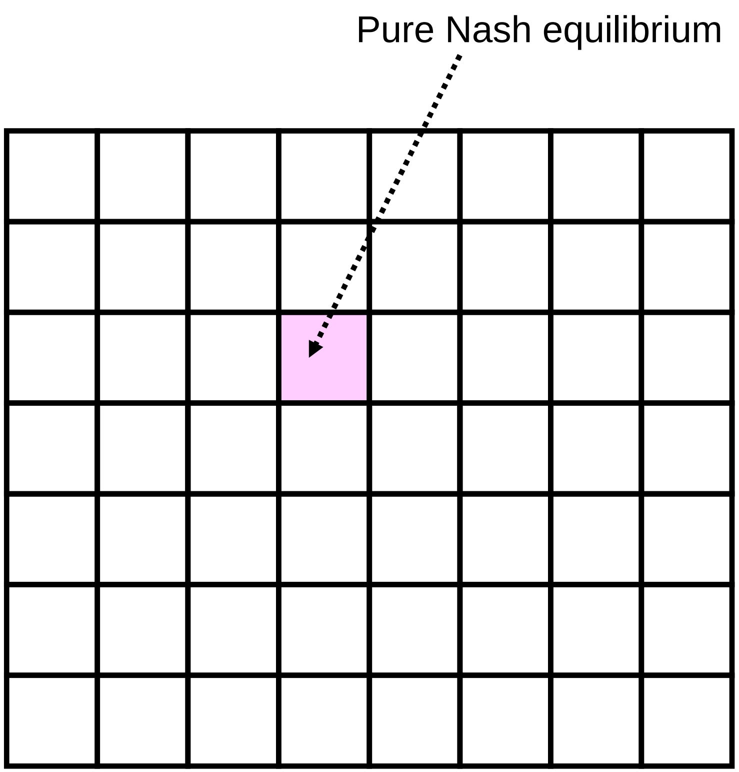

# Nash Equilibria can be implemented for free!

- Two player example:

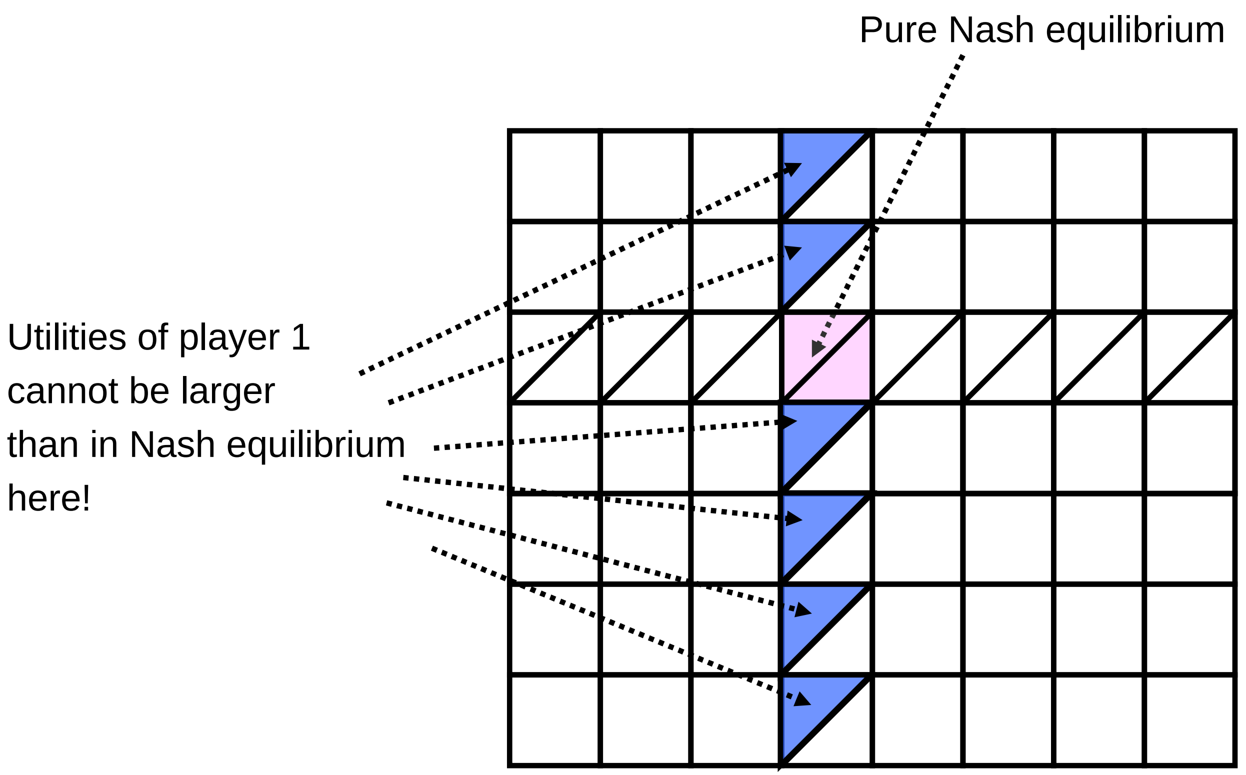

# Nash Equilibria can be implemented for free!

- Two player example:

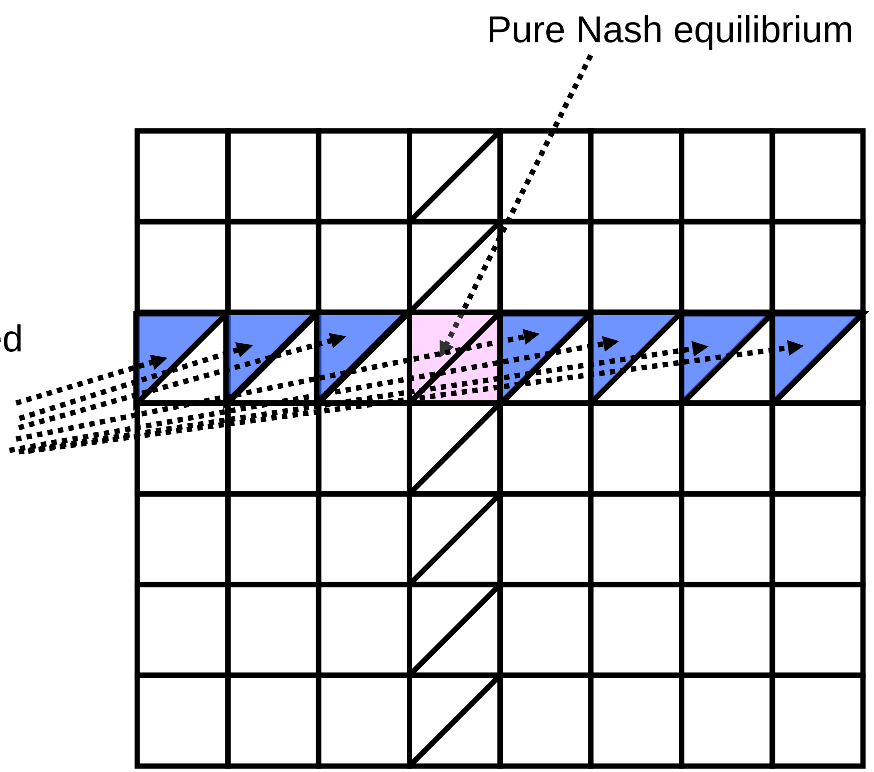


In order to make this the only non-dominated strategy, promise much money here – it's free!

# Nash Equilibria can be implemented for free!

- Two player example:

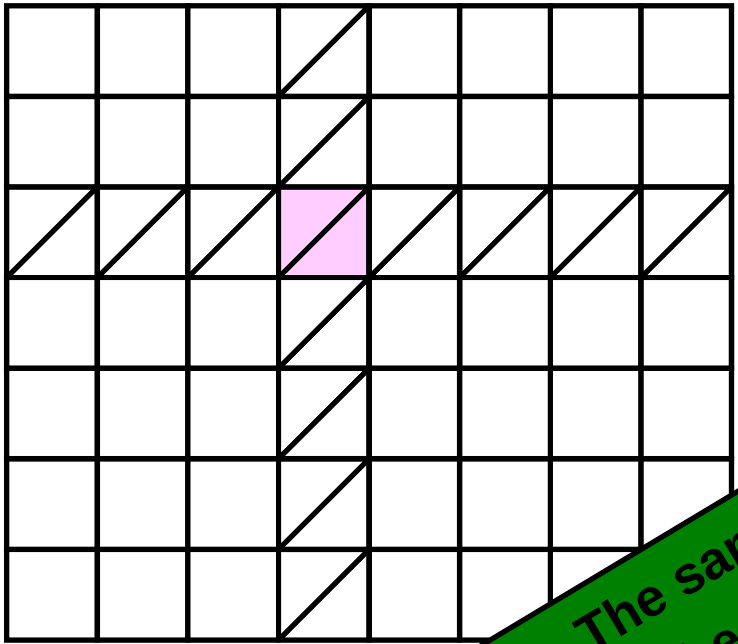

**The same can be done with the second player, and hence the claim follows.**

# Mechanism Design without Money

- Arrow's theorem: Limited impact only

- There are situations where it is still useful

- Approaches today:

  - Punishments for non-cooperative behavior (e.g., jamming)
  - Tit-for-tat solutions (e.g., BitTorrent)
  - Money burning mechanisms: pay by computations or bandwidth

- Further reading (our approach):

  „k-Implementation" by Monderer and Tennenholtz

  - related to correlated equilibria (but with money)
  - indeed: all correlated equilibria are 0-implementable