

# Strategic Properties of Peer-to-Peer Data Structures and Behaviour of Human Peers in Economic Experiments

Stephan Schosser  
schosser@ipd.uka.de  
Universität Karlsruhe (TH), Germany

Rainer Schmidt  
rainer.schmidt@uni-magdeburg.de  
Universität Magdeburg, Germany

Klemens Böhm  
boehm@ipd.uka.de  
Universität Karlsruhe (TH), Germany

Bodo Vogt  
bodo.vogt@ww.uni-magdeburg.de  
Universität Magdeburg, Germany

## Abstract

*Structured peer-to-peer networks manage large data sets. Each peer administers a part of the data. To answer a query, several peers must collaborate. However, experience with peer-to-peer file sharing systems shows that peers tend to abandon collaboration. Such behaviour is also expected in structured peer-to-peer systems. In the recent past, protocols for such systems to counter this kind of behaviour have been proposed. Our objective is to investigate how the behaviour of peers reacts to such protocols. We use idealized game theoretic models to analyze the strategic properties of structured peer-to-peer networks. The models predict cut-off strategies. I.e., if more than a certain percentage of their queries are answered, peers also answer queries in a stationary equilibrium. Another theoretical result is that free riding may occur in equilibrium. To check whether our theoretical predictions are in line with reality, we have implemented our model in an experimental laboratory and have analyzed the behaviour of human peers. These experiments support our theoretical results.*

## 1. Introduction

Peer-to-peer systems are coordinator-free distributed systems. Each node (aka. peer) that is part of a peer-to-peer system is expected to participate in the ongoing work. At the same time, it may also make use of the system as a whole. The peer-to-peer approach has fundamental advantages over systems with a coordinator node, notably reliability and scalability.

In structured peer-to-peer networks, as in other peer-to-peer systems, there typically is more than one peer involved in the evaluation of a query. If one of these peers defects, it is difficult to impossible to identify it. A peer cannot track

its queries. It has to trust in the cooperativeness of the other peers. Only when they decide to handle its query, it will receive a query result.

A natural phenomenon in peer-to-peer systems is free riding, i.e., peers trying to make use of the system without contributing to it. Identifying such uncooperative peers in Content Addressable Networks (CAN) is difficult. CAN are the variant of structured peer-to-peer networks we investigate here. The problem is difficult because each peer only has information about its neighbours. It does not know other peers and cannot observe their reliability. The free riding problem in structured peer-to-peer systems has not received much attention so far. An exception is the variant of Content Addressable Networks described in [2]. With the approach proposed there, peers remember which interactions with other peers have satisfied them. In [2] it is shown that the protocol achieves the discrimination sought between cooperative and uncooperative peers. However, these results only hold under certain model assumptions that are relatively rigid. Some of the assumptions are as follows: The rate of uncooperative peers in the system is known to all peers. Peers either follow the protocol or free ride, but do not resort to other strategies. Elementary questions remain open: Which factors determine cooperation in structured peer-to-peer systems? Given these factors, can we improve the design of the system? Are there any situations that cause the system to collapse when relaxing the model assumptions proposed?

To answer these questions, and to deepen our understanding, we analyse the behaviour of peers in structured peer-to-peer networks from a strategic perspective. We represent the complex peer-to-peer structure with idealized game theoretic models and explore the strategic properties. The equilibria of the models can be regarded as the stable situation of the complex system. We use these models to derive two hypotheses: First, we expect that free riding of

one peer will not cause the system to break down. Second, we expect that the probability of participation of peers in the system depends on the fraction of positive interactions with other peers in the past.

Since it is well known from research in experimental economics that theoretic predictions do not always coincide with actual behaviour, we implement a CAN variant that we can assess with human peers in the laboratory. We are not aware of any other evaluation of protocols for peer-to-peer networks with 'human peers', even though this approach bears a significant advantage: Any assumptions about behaviour are not needed since they are observed! Our motivation is to learn which strategies humans use, and how they affect the robustness, efficiency and amount of traffic of the peer-to-peer structure. Such an approach is superior to solely relying on the intuition of the protocol designers which assumptions are natural, and which ones are not. It turns out that our experiments confirm the hypotheses mentioned above. The paper is organized as follows. Section 2 reviews related work. Section 3 gives a short introduction to Content Addressable Networks. Section 4 describes the idealized economic situations to formalize structured peer-to-peer networks and derives our hypotheses for the experiments. In Section 5 we introduce our experimental design. Section 6 presents and discusses the experimental results. Section 7 concludes.

## 2. Related Work

Empirical studies of peer-to-peer systems, based on usage data of Napster and Gnutella [1] [15], show that many users prefer not to share any resources. Instead, they make heavy use of the resources provided by others. This behaviour, called free riding, has severe negative impact on peer-to-peer systems and can even let the system collapse [11]. Hence it is important to thoroughly evaluate strategies of peers in peer-to-peer systems before such systems become operational.

Ge et al. [8] propose a mathematical model to analyse the performance of peer-to-peer systems. They use their model to compare structured peer-to-peer systems to other peer-to-peer architectures. They show that the amount of data transferred in structured peer-to-peer systems is higher than in other peer-to-peer systems. This approach assumes that the participants are cooperative. It does not investigate the impact of uncooperativeness on the amount of data transferred.

To analyze peer strategies, the cost of participation in peer-to-peer systems has to be taken into account. Fabrikant et al. [6] and Chun et al. [5] provide cost models for internet-like network structures. These approaches model the participant as a non-cooperative player, who has an incentive to participate in the system, and who wants to min-

imize the price for doing so. They regard only the distance between nodes and the degree of connectivity. Christin and Chuang [4] extend this work by considering the load of each peer. They calculate social optima and Nash equilibria of such systems. The results are then used to analyze existing overlay topologies.

Golle et al. [9] give a simple game-theoretic model to analyze agent behaviour in peer-to-peer systems with a central coordinator. They propose several payment mechanisms to encourage sharing. Finally, they use results of simulations to substantiate their theoretical results. Ranganathan et al. [12] use a similar approach. They model peers as uncooperative players and compare reputation-based incentive mechanisms for peer-to-peer systems. They show that these mechanisms can be used to effectively counter free riding. Feldman et al. [7] extend this work by more detailed analyses. They use different reputation-based incentive mechanisms and simulations with different attack strategies to prove the usability of their system.

Buroghain et al. [3] use a differential service based on incentive mechanism to eliminate free riders. They show that the strategy of a peer wishing to join the system only depends on the benefit it can derive from the system. If the benefit is larger than a certain threshold, a peer will join the system.

While these approaches make extensive use of simulations, none of them gives a detailed game-theoretic model and verifies it using economic experiments.

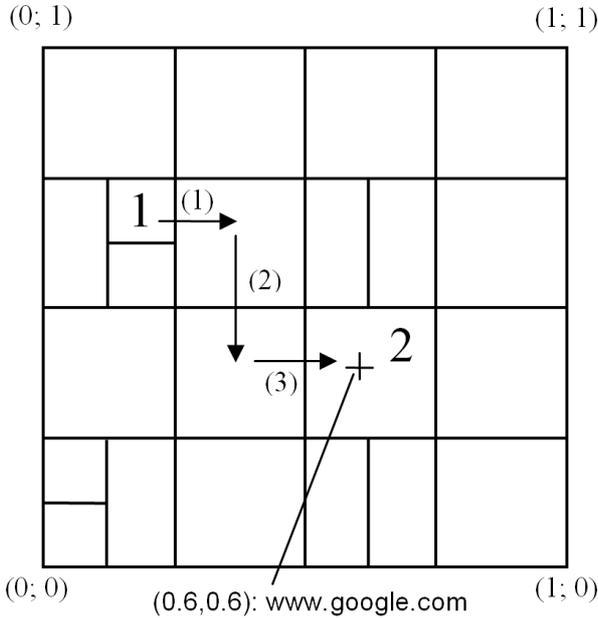
## 3. Content Addressable Networks

Several variants of structured peer-to-peer networks have been proposed recently. CHORD [17], Tapestry [18] and Pastry [14] organize the peers using a one-dimensional hash function. Each peer knows a subset of the peers in the system. Content Addressable Networks [13] in turn use an n-dimensional hash function.

Content Addressable Networks deterministically assign data objects to nodes according to their keys. This allows to guarantee that the data object with a given key is found within a certain number of steps, under several model assumptions [2]. Structured peer-to-peer networks require a hash function that maps keys to coordinates of an n-dimensional space. Each peer is responsible for a partition of the coordinate space, i.e., it stores all data objects whose mapped keys fall into its partition. A peer also knows its neighbours in the key space and their partition.

**Example.** Suppose that we want to store meta information on web sites. E.g., the meta information that would come with `www.google.com` might be "search", the one of `www.cnn.com` might be "news", and the one of `www.mozilla.org` might be "browser". Further, suppose that the hash function maps the URL "www.google.com" to the

coordinates (0.6, 0.6). The node responsible for (0.6, 0.6) now stores the pair (www.google.com, "search"). Figure 1 shows the segmentation of the coordinate space among the nodes. Each rectangle represents the partition a peer is responsible for. ■



**Figure 1. Content Addressable Network**

Each peer needs additional data owned by other peers. Hence, peers trade data by issuing queries. From an economic perspective, the good exchanged is mutual information. To obtain positive utility a node obtains the data object belonging to a given key. When searching data objects within the system, a simple variant of greedy forward routing is used. The peer sends its query, the searched key, to a neighbour whose partition is closer to the coordinates of the key. The neighbour checks whether it can answer the query. If a peer can not answer the query itself, this step is repeated until the query arrives at the peer that can answer the query. This peer then returns the query result.

**Example.** Figure 1 illustrates this algorithm for a query issued by Peer 1. The query asks for meta data on "www.google.com". The query is mapped to the coordinates (0.6, 0.6). Peer 1 now identifies its neighbour with the smallest Euclidean distance to (0.6, 0.6) in the coordinate space. This is the peer right of it. It forwards the query to this peer (Arrow (1)). Since this peer cannot answer the query itself either, the procedure recurs until the query reaches Peer 2 (Arrows (2) and (3)). Peer 2 answers the query and returns the result to Peer 1. ■

To simplify matters, the key space forms a torus. This allows to forward messages without ever reaching an edge.

The standard implementation of CAN, which is de-

scribed here, assumes that each peer forwards or replies to all incoming queries.

## 4 Models and Hypotheses

Before we derive our hypotheses starting from problems in computer science and considering game theoretic models, we first motivate the payoff structure in structured peer-to-peer networks from an economic point of view. In economic modelling the utility of an actor determines his incentives. Utility can be obtained from money as well as from other factors. Obviously costs result in negative utility and positive payoffs in positive utility. Doing work also results in negative utility.

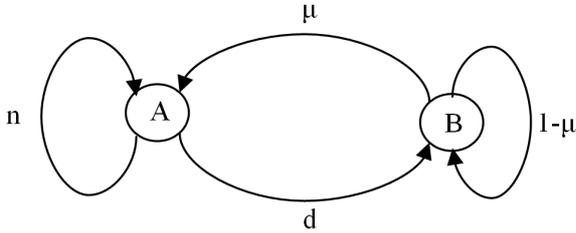
In modelling the peer-to-peer structure we assigned utility to every action of a peer as follows: A peer obtains positive utility if one of its queries has been answered. It obtains negative utility from answering or forwarding a query and from sending one. For reasons of simplicity we identify money with utility in the discussion. This is because the relation between these two is monotone and increasing, even though the relation might be non-linear. In other words, the utility of a participant increases if his payoff increases.

### 4.1. Underlying Model from Economics

We compare structured peer-to-peer networks to the model of Hens and Vogt [10]. In this model, agents can be in two states: In State A they can deliver a service. In State B they can consume one. Delivering a service corresponds to answering or forwarding a query in structured peer-to-peer systems. An agent which delivers a service loses some utility by doing so. To participate in the system, a peer has to invest money for the computing infrastructure. In our design these costs are part of the costs of sending a query. Hence sending a query corresponds to delivering a service. We expect peers which do not send queries to not answer any incoming messages. Hence, their infrastructure costs are 0. This behaviour corresponds to leaving the system. But infrastructure costs exist when issuing queries. Consuming a service corresponds to obtaining a query result in structured peer-to-peer systems. Agents gain positive utility from the service consumed. The amount of this utility exceeds the one of delivering a service.

The model is round-based. In each round an agent in State A may deliver a service (d) or not (n). The agents in State B do not have an action choice. After each round an agent in State A offering a service is matched with an agent in State B. In principle, every agent in State B can consume a service if enough agents in State A offer it. Otherwise, only a fraction  $\mu$  of agents randomly chosen receives the service. After each round the agents that have been matched switch to State A. In consequence an agent does not offer a

service in two subsequent rounds nor consume services one right after the other. Figure 2 serves as a visualization. For an agent the number of rounds in the game is infinite. As in many economic problems, the end of an interaction is not determined.



**Figure 2. Strategic options**

Discounting takes place between two subsequent rounds: The utility drawn from a service consumed later is lower. Without discounting, it would be rational to always cooperate. This is in order to be able to consume the service again as early as possible. With discounting in turn, the picture is less clear: If an agent remains in State B and the discounting is high, the utility drawn from consuming the service will not level up to the one for delivering the service earlier. One can, however, compute when exactly cooperation is rational using economic methods [10]: In equilibrium, cooperation is rational if it pays in the future. For games with finite horizon, the standard way of arguing is backward induction. It would tell us that non-cooperative behaviour is rational. However, backward induction is not applicable since the horizon is infinite. For such a setting, the notion of stationary equilibrium is used for the computation. The main results from [10] are as follows: the key parameter indicating the cooperation of a player is the rate of his service requests that have been successful. If it is above a certain cut-off value he cooperates.

## 4.2. Hypotheses

We perceive this model as a simplification of the interactions in a structured peer-to-peer system. We compare the equilibrium strategies in the model to the strategies used in such peer-to-peer systems. This gives way to the following predictions regarding structured peer-to-peer systems: Peers will use cut-off strategies, and free riding can exist in equilibrium together with cooperation. In other words, based on those theoretical results, we formulate the following two hypotheses, to be evaluated in the experiments that follow.

*Hypothesis S (Cut-off strategies):* If the past success frequency is above a critical value (cut-off value) players will participate in the peer-to-peer system. If the frequency is below a cut-off value the players will not answer, forward

or send queries. The past success frequency is the rate of queries answered in the past by other participants. Hypothesis S reflects the theoretical result that players resort to cut-off-strategies.

*Hypothesis FR (Free Riding without break-down):* Even for small  $k$ , free riding of  $1/k$  of all peers will not cause a break-down of the system, where  $k$  is the number of participants in the system. In our experiments we will show that this holds for  $k$  equals six. In Hypothesis FR we test whether free riding exists in one system without a break down of the system. This is one equilibrium prediction based on [10].

## 5. Experiments

For our experiments we used a simulation environment for structured peer-to-peer systems. The simulation environment controls the structure of the peer-to-peer network, i.e., the neighbours of each peer. The important feature of the simulation environment is that human players can mimic peers. We reduced the number of peers in the system to the number of human players. According to Selten, four are few, and six are many [16]. This quote describes the observation that a small number of persons interacting shows the same behaviour as a large group. We therefore chose six participants for each experiment. Although the participants were told to play against both human and computer-controlled peers, our peer-to-peer system consists of human participants only. The participants did not know the size of the peer-to-peer system. As the strategies of the participants did not depend on the number of peers in the system, we expect the results of our experiments to hold for larger peer-to-peer structures as well.

The theory we have referred to is based on utility considerations. Since utility is not directly observable, payoff of the participants has been money.

Over nine months, we have conducted experiments with 60 students from various disciplines from Magdeburg University. They have been recruited by announcements in the university and have been promised monetary reward contingent on performance in a group decision-making experiment. The participants have received points as payoffs. 100 points were 2 Euros. The average payoff of a participant was 11.05 Euro.

The experimental software was programmed in Java. The experiments were conducted in the computer laboratories of Magdeburg University. The participants were randomly assigned to the seats in the laboratory. The experiment lasted approximately 120 minutes. The first 20 minutes consisted of orientation and instructions. The instructions were provided in written form. The experiment continued with a learning phase during which the participants played for several rounds. The game was played in

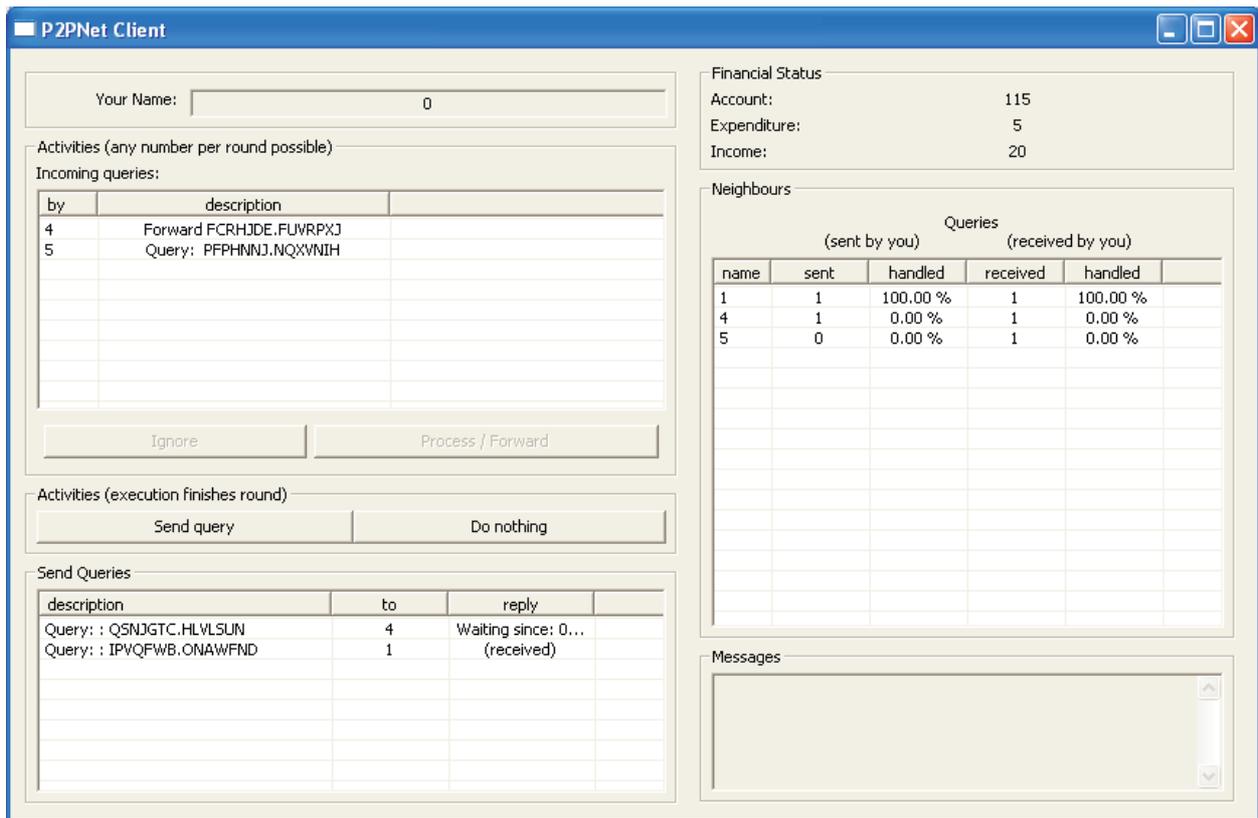


Figure 3. Experiment client

ten groups of six participants via computer terminals. The terminals were well separated from one another. The rationale has been to prevent communication between the participants.

In the beginning, 20 periods were played without discounting. Afterwards we introduced a discounting rate. In the laboratory a discounting rate is implemented as the probability that the game continues after every period. From a theoretical point of view the results do not depend on the interpretation of as a discounting rate or a continuation probability. The chance move whether the game ends or not was performed by rolling a 10-sided dice which was observable by the participants. Hence, in the experiment the participants played the game with a discounting rate of 0.1.

In the beginning of the game every player received a list of 200 strings (numbers and letters) representing keys in a structured peer-to-peer network. For every string on the list of one player there has been a complementary string on the list of another player. The complementary strings represent data objects saved within a structured peer-to-peer network. The distribution of these strings was according to the logic of a Content Addressable Network, as described before. The strings representing keys were randomly distributed among the players. Each player knew the complemen-

tary strings of 200 keys. Every player could send a query for a complementary string. If he received the complementary string he obtained a payoff of 20 points. Every round a player was allowed to ask for just one complementary string for one of his numbers. This should reflect the intuition of a structured peer-to-peer network. Here peers regularly submit queries to other participants. When sending a query the simulation environment selected a string and offered 1-3 peers. These peers were chosen from the neighbours of the players, and they were ordered by their distance to the query result. The player could select one of the peers offered, and the simulation environment then forwarded it the query. The environment offered more than one peer for the following reason: For sending a query, both the probability of having the information as well as the past behaviour are important. For each query received, a player could choose each round whether to process it or not. In the first case, the simulation environment informed the player, whether he could answer the query himself, i.e., whether the string requested was in the list of his own 200 strings. If the player could not answer, he was asked to forward the query. If forwarding was chosen, the next player received the query. He then had the same action choices as his predecessor.

Players were endowed with 100 points in the beginning.

The cost of answering a query was 5 points, the cost of sending a query was 2 points, and the cost of forwarding was 1 point. The payoff for receiving a data object requested was 20 points. See Table 1.

**Table 1. Payoff depending on actions**

Action	Payoff
Initial balance	100
Answering a query	-5
Sending a query	-2
Forwarding a query	-1
Receiving a query result	20

The information available to the participants is shown in the screenshot in Figure 3. The screen shows information about a peer with Identifier 0. Player 0 has sent a query to Player 1 and later a query to Player 4. This can be seen in the bottom left area of the screen. The system (not necessarily Player 1) has already answered the second query. Player 0 only sees that the query sent to Player 1 was answered. Similarly to structured peer-to-peer systems, the simulation environment does not make information on forwarders public. Hence, column 'reply' contains "[received]". On the other hand, Player 0 is still waiting for the result of the query sent to Player 4. Furthermore, Player 0 has received two requests from other players, which he has not yet processed. This is shown in the upper left area of the screen. Player 5 has asked for the result of a query which Player 0 can deliver itself. Player 0 cannot answer the query obtained from Player 4 and can only forward it to other peers. The center right area of the screen tells us that Player 0 has sent one query to Player 1 and one to Player 4. Player 1 has answered the query. This is because the ratio of the messages that have been processed is "100.00 %". Similarly, this area indicates that Player 0 received one query from Player 1. The value "100.00 %" in the column labelled 'handled' indicates that he has already answered this query. The value "Account" in the upper right is the payoff of a player. After every period the participants were informed about the result of the chance move (whether the game continues), their payoff in the last period and their total payoff. They were not informed about any characteristic of the other participants like their scores or their choices. Moreover, participants could not identify other participants. All information was given on the computer screen.

We conducted two games where the participants played against each other. We will call these games treatments. In Treatment 1, the game was played with six human players. In this treatment we tested for the strategies of human beings and wanted to find out whether cut-off strategies were played (Hypothesis S). In Treatment 2 one peer was free riding. Participants were told that they are interacting with human peers and programmed peers. Hence they could ex-

pect any strange behaviour. We used these two treatments to test whether free riding may lead to a collapse of the system (Hypothesis FR).

After the treatments (Treatment 1 and 2), a so-called strategy game was played. In a strategy game, all players are asked to describe their strategies in written form. A strategy game takes place after players have played some treatments. This allows gaining information on the strategies of experienced players: While players learn during the treatments, they have a thorough understanding of the game after the treatments. All participants selected their strategies for all games. The strategies could depend on the parameters of a game and also on the whole history of a play of the game. The participants were free to describe their strategies and its dependencies on certain parameters without any restrictions.

After the experiment was completed each participant was separately paid in cash contingent on her performance.

## 6. Results

In this section we summarize the results of our experiments with human participants and relate them to our theoretical predictions.

### 6.1. Hypothesis S (cut-off strategies)

For a test of Hypothesis S we analyze the results of the strategy game, because we are interested in the behaviour of experienced players and want to minimize the side effects of learning. We also wanted to observe behaviour conditional on the actions of the other players. We obtain this conditional behaviour from the strategies of the players. Another reason to use a strategy game is that we cannot observe a full strategy of a player only by analyzing the treatments. This is because the chance move will stop the game after several periods, and some parts of the strategy will not be observed.

The strategies between Treatment 1 and Treatment 2 did not differ. Most strategies had a start phase and a main phase. The behaviour of the participants in the start phase, which ended after a few periods, was very diverse. It might not depend on the past success frequency in the beginning, as only few interactions were observed to that moment. Since we can only analyze behaviour of peers in an equilibrium using economic theory, we concentrate on the strategies after the start phase. Table 2 shows the result of our analysis.

We classified the strategies as follows: whether an action is chosen is

- a) only depending on the past success frequency of own queries (with different cut-offs),

- b) depending on the past success frequency of own queries, but also other factors,
- c) not in line with Hypothesis S.

Three participants chose strategies that depend on the absolute number of unanswered queries and not on the success frequency. Five used different types of strategies. All other strategies depend on the past success frequency. 41 strategies were according to Hypothesis S. Three participants tried to anticipate the end of the game, and three introduced the condition of doing less than the others (a kind of free riding). This can be regarded as a slight modification of a strategy according to Hypothesis S. The remaining five strategies show stronger differences to a cut-off strategy.

**Table 2. Types of strategies in strategy game**

Strategy	Category	# persons
Cut-off strategy only depending on past success frequency	a)	41
Cut-off strategy depending on past success frequency plus end phase	b)	3
Cut-off strategy depending on past success frequency plus the condition to answer less queries than own queries are answered	b)	3
Cut-off strategy depending on past success frequency for own queries, answering randomly	b)	4
Cut-off strategy depending on past success frequency, answering queries after several non answered queries	b)	1
Cut-off strategy depending on the absolute number of own queries not answered	c)	3
Different types of strategies	c)	5

We test the Null-hypothesis that the participants select any other than a cut-off strategy in a binomial test. We only count the 41 cut-off strategies (out of 60) as contradicting the Null hypothesis. Then we can reject the Null hypothesis on a significance level of 1%. This test result confirms Hypothesis S.

This shows that cut-off strategies are accepted by the majority of players in structured peer-to-peer systems. Hence implementing mechanisms that allow the exchange of data depending on the percentage or amount of queries answered in the past is an accepted behaviour. A small number realized the chance of stopping cooperation at the end of the game. This implies that countermeasures against defecting in the ending phase have to be found.

## 6.2. Hypothesis Fr (Free Riding without Break down)

In Treatment 2 one peer was instructed to send queries and not to process queries of others. This player is called destructive player.

**Table 3. Average payoff per round**

Group	Treatment 1		Treatment 2		
	all peers	destructive peer in Treatment 2	all peers	destructive peer	other peers
1	5.85	7.71	4.11	3.45	4.25
2	2.63	5.90	0.89	1.27	0.82
3	4.40	6.52	1.85	2.35	1.75
4	6.79	10.20	2.81	8.00	1.78
5	7.78	7.36	4.47	3.83	4.59
6	2.31	5.95	0.99	3.67	0.46
7	4.31	-0.50	1.47	2.17	1.33
8	5.30	0.45	0.93	0.50	1.02
9	5.01	7.95	4.61	5.20	4.50
10	2.10	0.09	1.27	3.56	0.81
Mean	4.65	5.16	2.34	3.40	2.13

We compare the results of Treatment 1 and Treatment 2 in Table 3. Comparing the average payoff per round between the two treatments yields the following results. In 10 out of 10 groups the payoff is smaller in Treatment 2 with a free-rider. In a binomial test this result is significant on the 1% level. The average payoff per round can be regarded as a function of the efficiency. We therefore conclude that the efficiency decreases. On average the payoff is reduced to about half of its value in Treatment 1, but the payoff is still higher than zero. It is significantly higher (0.1% in a binomial test) than half of the one in Treatment 1. A general result from game theory tells us that in the end of the treatment the peer-to-peer system might break down. But even if we skip the first 10 rounds from our analysis during which cooperation tends to be the strongest we obtain the same test results (see Table 4). This result confirms Hypothesis FR.

Comparing the payoff of the destructive player to the average payoff of the remaining peers in Treatment 2 we see that the destructive player has a slightly higher payoff on average. In 8 out of 10 groups he performs better. This shows that the destructive player tends to perform better than the average remaining player. However, if we compare the average payoff of the destructive player between both treatments he receives a lower payoff in Treatment 2. In 7 of 10 groups he performs better in Treatment 1. This shows that free riding does not pay for the participants.

Players might not always be motivated to obtain a maximal payoff for themselves, but by obtaining more than oth-

**Table 4. Average payoff after round 10**

Group	Treatment 1		Treatment 2		
	all peers	destructive peer in Treatment 2	all peers	destructive peer	other peers
1	5.77	8.29	5.07	4.67	5.15
2	3.58	7.80	0.83	0.09	0.97
3	4.61	8.00	0.97	-0.46	1.26
4	7.43	9.73	2.74	5.14	2.26
5	8.40	7.67	4.74	1.86	5.31
6	2.50	1.90	1.07	3.14	0.66
7	4.29	0.21	0.99	-0.57	1.30
8	6.01	2.92	0.68	-2.00	1.21
9	5.71	7.50	4.98	2.00	5.57
10	2.31	-1.23	0.95	4.94	0.15
Mean	5.06	5.28	2.30	1.88	2.39

ers. One may perceive this as a special case of competition. We measure this effect by calculating relative payoffs. Relative payoffs are the quotients of one's own payoff and the average payoff of the remaining players. If a player is motivated by relative payoff, free riding is attractive. This is because he earns more than the other players. However, if one is motivated by his own payoff in absolute terms, free riding does not pay. This is because the player also receives a payoff lower than in the games in which he has cooperated. In most cases such as in our setting, relative payoffs are not known. Hence, free riding does not seem to be attractive. In our experiment nobody always defects (unless he has been instructed to do so). This shows that the incentive from relative payoffs is not available to the peers and/or not high enough. Our design protects from free riding.

## 7. Conclusions

In our paper we have combined economic and computer science tools to analyze strategic aspects of peer-to-peer structures. Complex peer-to-peer structures are used in a variety of applications. For these applications the following problems have to be solved: stability of the system, cooperation of peers, free riding of peers. We used a game theoretic model to explore the strategic opportunities of the participants. The model predicts that peers use cut-off strategies. It is rational to participate in the system, and even a working of the system with free riding is rational and can be expected.

We confronted our predictions with an (economic) experimental test with human peers. The experimental data confirms our predictions. We obtain cut-off strategies as predicted by theory and a working system with free-riders. If the information of the payoffs of others is not known,

free riding is not attractive in our setting. Experiments with human peers allow to test which kind of behaviour does not occur, and which kind of behaviour the designers of peer-to-peer systems must reckon with. To the best of our knowledge, this is the first investigation where humans assume the role of peers in peer-to-peer structures. The advantage is that we do not need to make any assumptions about the behaviour of peers, compared to simulations.

In the future we will extend our analysis to more complex peer-to-peer systems. We will introduce feedback in the experiments. We also plan to implement computer-controlled peers using the strategies found here and let them play against human participants.

## 8. Acknowledgments

We thank the anonymous reviewers for their comments. They have helped a lot to improve the paper.

## References

- [1] E. Adar and B. Huberman. Free riding on gnutella. Technical report, Xerox PARC, 2000.
- [2] K. Böhm and E. Buchmann. Free riding-aware forwarding in content-addressable networks. *VLDB Journal*, to appear:–, 2005.
- [3] C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in p2p systems. In *Proceedings of the 3rd International Conference on Peer-to-Peer Computing*, page 48, 2003.
- [4] N. Christin and J. Chuang. A cost-based analysis of overlay routing geometries. In *Proceedings of IEEE INFOCOM05*, 2005.
- [5] B.-G. Chun, R. Fonseca, I. Stoica, and J. Kubiawicz. Characterizing selfishly constructed overlay routing networks. In *Proceedings of the 23rd IEEE International Conference on Computer Communications*, 2004.
- [6] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proceedings of ACM PODC*, 2003.
- [7] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 102–111, 2004.
- [8] Z. Ge, D. R. Figueiredo, S. Jaiswal, J. Kurose, and D. Towsley. Modeling peer-peer file sharing systems. In *Proceedings of the INFOCOM*, 2003.
- [9] P. Golle, K. Leyton-Brown, and I. Mironov. Incentives for sharing in peer-to-peer networks. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 264–267, 2001.
- [10] T. Hens and B. Vogt. Money and reciprocity. Iew - working papers, Institute for Empirical Research in Economics - IEW, 2003.
- [11] L. Ramasamy and L. Liu. Free riding: A new challenge to peer-to-peer file sharing systems. In *Proceedings of the 36th International Conference on System Sciences*, 2003.

- [12] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster. To share or not to share: An analysis of incentives to contribute in collaborative file sharing environments. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems*, 2003.
- [13] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 161–172, 2001.
- [14] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Proceedings of the International Conference on Distributed Systems Platforms*, 2001.
- [15] S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In *Proceedings of Multimedia Computing and Networking*, 2002.
- [16] R. Selten. A simple model of imperfect competition where four are few and six are many. *International Journal of Game Theory*, 2:141–201, 1973.
- [17] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Transactions on Networking*, 11:17–32, 2003.
- [18] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical report, U. C. Berkeley, 2001.