


Universität Karlsruhe - Fakultät für Informatik - Bibliothek - Postfach 6980 - 76128 Karlsruhe

**Competition and Cooperation in
Heterogeneous Structured P2P Systems
- Are They Mutually Exclusive?**

Stephan Schosser, Klemens Böhm und Bodo Vogt

Interner Bericht 2007-17



Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

ISSN 1432-7864



Fakultät für **Informatik**

Competition and Cooperation in Heterogeneous Structured P2P Systems – Are They Mutually Exclusive?

Stephan Schosser Klemens Böhm
Universität Karlsruhe (TH), Germany
{schosser|boehm}@ira.uka.de

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

Abstract

We analyze how to ensure efficiency (in the economic sense of the word) in structured P2P systems, with a focus on heterogeneity. In general, two factors influence efficiency, the degree of cooperation and the degree of competition. In competitive situations, the degree of cooperation tends to be very low. Structured P2P systems in turn, where peers have reason to behave in a competitive way as well as to cooperate, have turned out to be surprisingly efficient - at least as long as peers are homogeneous [20]. The reason is that there is a rather rigid arrangement of peers, i.e., structure, giving way to what is called indirect partner interaction. However, the homogeneity assumption is restrictive, as peers face different costs and benefits due to bandwidth limitations etc. Economic experiments in general (not with a focus on structured P2P systems) suggest that heterogeneity is in the way of efficiency, i.e., reduces the sum of payoffs received. By means of behavioral experiments, we show that most peers in structured P2P systems cooperate, even if they earn less. This result is an important step towards establishing structured P2P systems that are operational.

1. Introduction

Structured P2P systems administer large amounts of data. Peers in such systems are run and maintained by different individuals/organizations, and the infrastructure costs are distributed among the peers. In this paper, we analyze how to ensure efficiency (in the economic sense of the word) in such systems.

1.1 Competition and Cooperation

In general, from an economic perspective, two motives influence the behavior of interaction partners: cooperation and competition. In competitive situations, i.e., several individuals compete, the degree of cooperation tends to be very low. The effect of this antagonism is

that ensuring efficiency is difficult in many situations. For instance, think of unstructured P2P systems for file sharing which cannot guarantee efficiency and therefore high payoffs. (The payoff is the overall benefit minus the overall costs of participating.) Peers have an incentive to free ride, and most of them actually do [1]. There is a broad range of settings giving way to both competitive and cooperative behavior: Think of two colleagues working in the same department. Under normal circumstances, they cooperate, e.g., by sharing knowledge. As soon as they both apply for the same position within the company, they are likely to stop cooperating and show selfish behavior. This is in the way of efficiency, i.e., the department will not have a high output. In competitive situations, selfish behavior of participants, i.e., they only care for their own payoff, leads to a breakdown of cooperation. Nevertheless, competitive markets exist where selfish behavior of individuals maximizes the overall payoff [25]. Again, think of the two colleagues, from the perspective of the employer: When they compete with another department also consisting of two individuals for a bonus, all employees will work harder and maximize their output. In the end, the colleagues who do not get the bonus will accept this if the competition has been fair.

Thus, the combination of cooperation and competition is promising as it can increase (economic) efficiency. However, a system designer has to be careful: Subtle design issues can have a significant effect on the outcome and can even decrease efficiency.

1.2 Heterogeneity

Behavioral experiments on interactions with two participants show that the situation is even worse if the participants are heterogeneous, i.e., they gain different benefits and face different costs from their actions. Humans try to ensure fairness among participants [5][10], i.e., ensure that the payoffs of participants are equal. This usually leads to a further decrease in efficiency [5]. However, while heterogeneity is a problem of two-person interaction, it is not a problem in com-

petitive markets. If persons behave in a competitive way, the outcome of a market can be very heterogeneous: the better trader gets more since he is a better trader. In this sense, competition yields stability.

1.3 Structured P2P Systems

In structured P2P systems, each peer can issue queries and should contribute to the ongoing work. A peer benefits by obtaining query results, while it encounters costs for processing queries. Typically, different peers are under control of different individuals/organizations, who are selfish and want to maximize their payoff. A key aspect of the design of P2P systems is efficiency, i.e., how much do peers benefit from the system on average. Hence, it is important to ensure high levels of cooperation. The more peers contribute to the ongoing work, the higher is efficiency.

In P2P systems, competition alone cannot help to reach efficiency. The peers face a dilemma of cooperation [8]: They can increase their payoffs by not contributing. On the other hand, systems with an explicit assignment of data and services to participants – like structured P2P systems – are promising. Here, indirect partner interaction takes place [21]. This is a novel economic paradigm for situations where participants are organized according to a certain structure. The structure helps to ensure competition among the peers: On the one hand, to receive query results, each peer needs to be at least as cooperative as most other peers in the eyes of others. On the other hand, the system does not suffer from uncooperative peers: As several message routes exist to forward queries, cooperative peers can replace uncooperative ones. In other words, a peer competes with other peers. The structure results in cooperation among participants, even if they do not interact repeatedly. Thus, this particular combination of cooperation and competition yields high levels of efficiency – if all peers face equal costs and benefits.

1.4 Contributions

In almost any P2P system, peers are heterogeneous. They have different costs for storing data, due to different hard disk sizes, for issuing messages, due to different bandwidth limitations, etc. As mentioned before, economic theory insinuates that heterogeneity has a negative effect on cooperation. This calls for an analysis of the impact of heterogeneity on the efficiency of structured P2P systems. We use behavioral experiments, i.e., let human participants assume the role of peers in structured P2P systems. Since any peer in structured P2P systems is under control of a human

after all (i.e., a human has implemented its behavior), such experiments are instructive with regard to (economic) characteristics of such systems.

While fairness leads to a decrease in efficiency in other systems, we show that it does not play much of a role in structured P2P systems: Here, competition among peers fosters cooperation. High levels of efficiency are reached. Individual peers are motivated to contribute if cooperativeness leads to higher average payoffs.

Paper outline: Section 2 reviews related work. We introduce the basic concepts of structured P2P systems in Section 3. In Section 4 we derive several hypotheses concerning our experiments, which we evaluate in Section 6, using the methodology described in Section 5. We conclude in Section 7.

2. Related Work

Measurement studies of existing P2P systems have shown that peers in such networks tend to behave uncooperatively [1][19]. They download content offered by others without contributing to the system. [17] shows that such behavior, called free-riding, can lead to the break-down of the system.

To counter such behavior, reputation mechanisms have been proposed [13][16]. Here, the peers observe each other and behave reciprocally. Different measures to quantify peer behavior are available. Kamvar et al. [13] propose to use the Eigentrust reputation measure [14]. Kung and Wu [15] use an Eigenvector-based reputation system. Both publications analyze the usability of different reputation measures to enforce cooperative behavior between different peers by using simulations. Neither one analyzes the impact of different strategies of the participating peers on their mechanisms.

Buroghain et al. [4] use game theory to analyze the strategic aspects when contributing in unstructured P2P systems. They show that a mechanism with differential service provision will eventually operate at a Nash Equilibrium. They use simulations to confirm that this equilibrium is even reached when peers are heterogeneous. [7] shows that the peers have no incentive to free-ride in homogeneous systems. Only if peers are heterogeneous they observe, by models and simulations, that peers tend to be uncooperative. [8] extends this work: It compares how well different mechanisms, such as contribution observations and feedback, ensure cooperation. The authors show that the contribution level sustains cooperation in smaller systems, but feedback is needed in larger systems to cope with white-washing. Blanc et al. take a similar approach [2] and show that the reputation system used is robust in the presence of malicious peers and noise. [20] analyzes

the behavior of human participants who control the behavior of different peers. [21] extends this work by introducing the concept of indirect partner interaction and showing that the usage of feedback does not improve the contribution level of the participants in structured P2P systems. Several properties of peers in such systems are identified. First, peers behave reciprocally, i.e., they differentiate between different contacts and forward or answer queries only for peers that have forwarded or answered queries for them before. Second, peers use cut-off strategies. I.e., a peer observes the behavior of another peer. Only if it is better than a certain threshold, it cooperates with the other peer.

While some of the publications mentioned analyze the impact of heterogeneity on system performance, they do not consider its impact on the strategies used. In particular, they do not investigate whether efficiency is sustained in the presence of heterogeneity.

3. Structured Peer-to-Peer Systems

In this section, we give a short overview over structured P2P systems and Content-Addressable Networks (CAN) in particular. Although our experiments were based on CAN, we expect our findings to be applicable to any system allowing indirect partner interaction, i.e., interaction via peers who interact repeatedly.

Structured P2P systems such as CAN [18], Chord [28] or Pastry [27] maintain large sets of (key, value)-pairs. Each key is mapped to a certain location in the key space. The key space is divided into several peer zones. Each peer is responsible for one of these zones. Next to the zone, each peer administers a list of peers with adjacent zones and their zones. After issuing a query (asking for a certain key), the key is transformed into the corresponding location in the key space, the query point. If the query point lies in its zone, the peer returns the query result, i.e., the value. Otherwise the query is forwarded to the peer closest to the query point. This recurs until the query finally reaches the peer having the query result. This peer then sends the query result to the initial sender.

Example: Figure 1 shows a CAN. The keys are transformed into two-dimensional coordinates using a hash function known to all peers. The (key, value)-pair (“Mars“, “Chocolate Bar”) for example is mapped to the coordinates (0.45, 0.3). Each square in the figure represents a peer zone. I.e. Peer F stores the value corresponding to (“Mars”). If Peer A is interested in information concerning “Mars”, it can forward its query using one of its contacts: Peer B, C, D, or E. Therefore, “Mars” is first transformed into the coordinates (0.45, 0.3) using the public hash function. Then the query is forwarded to the peer closest to (0.45, 0.3), Peer B. Peer B does not ad-

minister the key “Mars”, hence it forwards the query to one of its contacts until the query arrives at Peer F. Peer F finally sends the query result to Peer A. ■




Figure 1: Content-Addressable Network

Structured P2P systems have two important properties: First, all peers interact with their contacts only. Second, several peers have to cooperate to answer a query. Under realistic assumptions it can be shown that if only 5% of the peers are uncooperative, 40% of the queries in such a system remain unanswered [3].

To distinguish between cooperative and uncooperative peers, each peer can monitor the behavior of its contacts. Based on these observations each peer can autonomously decide which contacts it deems cooperative or uncooperative. Further, it can decide whether to process a query received from one of its contacts and which one to forward it to, when necessary. Experiments have shown that peers forward queries for cooperative peers only [20].

In our analysis we assume that each peer benefits when receiving query results, while it faces costs when forwarding, issuing or answering a query. In what follows, we only consider scenarios where the benefit of obtaining a query result exceeds the costs of processing it. In this paper, we analyze the effect of heterogeneity among the peers on the strategies observed. We focus on the influence of different benefits and costs among the peers on peer behavior.

4. Methodology

From an economic perspective, the designers of P2P systems face a market-design problem. The system as a whole is a market of information. Each peer controls a fraction of the information in the system and can issue queries for some of the information in the system. An objective of market design is Pareto-efficiency [29]. A Pareto-efficient allocation is one where no participant can increase his payoff without another participant facing a decreased payoff. As this definition is difficult to observe in complex settings, we define a market to be efficient if the sum of all payoffs is maximal. This goal

is in line with the goal of the designer of a P2P system: From his perspective, systems are optimal if the sum of the payoffs of all peers is maximal. This is the case if all peers fully cooperate, and all messages are handled. To analyze how efficiency is reached under heterogeneity, we will use behavioral experiments, which have one important advantage: Existing analyses of P2P systems rely on assumptions. [7] for instance assumes that heterogeneity does not influence the strategy, but only the payoffs. With behavioral experiments, assumptions regarding the behavior of peers are not necessary – the behavior is simply observed. This lets us observe not only the impact of heterogeneity on the payoff, but also on the behavior of peers. Analyzing the behavior of human participants mimicking peers in structured P2P systems provides valuable insights: Researchers, be they biologists or sociologists, agree that evolution has led to patterns of human behavior (aka. strategies) which are mature and sophisticated [11]. They are a good basis for developing protocols for systems where software agents interact.

5. Hypotheses

Whether efficiency is reached depends on several specific design issues. In this section, we show under which conditions efficiency is reached. Next, we derive predictions on the efficiency of structured P2P systems using economic arguments. Finally, we state why we expect several existing mechanisms in structured P2P systems, such as observations of other peers and cut-off strategies, to be beneficial in this context.

5.1 Cooperation in Two-Person Interaction

To better understand P2P structures which consist of two-person interactions (e.g., a message is forwarded from one peer to another one) and the evolution of cooperation over time, we use the helping game [24] as an example: One participant has the role of a donor, while the other participant is the recipient. The donor can decide to make a donation c . In this case, the recipient receives a benefit b . The cost of donating c is lower than the benefit b of the recipient. If the donor decides not to donate, both participants do not receive any reward. Game theory predicts that nobody will make a donation. But behavioral experiments [24] show that some participants (approximately 22%) actually donate. The motivation for such cooperative behavior is called fairness. But fairness also implies that participants favor same payoffs for all. Literature predicts [6] that this is counterproductive in case of het-

erogeneity, e.g., if the recipient is known to be ‘rich’ already when the game starts: Inequality usually leads to a decrease in cooperation and efficiency due to fairness considerations [6]. Given a scenario where the recipient is already ‘rich’, the donor would not give him any money. In other words, participants have an “inequality aversion”, i.e., prefer allocations where all participants have the same payoff, and choose their actions accordingly¹. This reasoning predicts a decrease in efficiency for structured P2P systems that are heterogeneous.

5.2 Competition

This subsection discusses competition, which has been proposed as a means to ensure efficiency under heterogeneity. We explain where exactly competition occurs in structured P2P systems. We then say why competition alone will not sustain cooperation in such systems. In competitive markets everybody can provide a good to everybody else and obtain a good at any point in time. Payoffs can be significantly different: Competition ensures efficiency.

In structured P2P systems, different peers have different goods (different pieces of information). Goods only exist more than once in case of replication (which we leave aside in our current study). Exchange also occurs over time: Usually the point of time at which a peer requests information is not identical to the point of time when others request information from it.

In systems without competition cooperation tends to decrease with the size of the network: [9] describes an experiment where peers form a circle, and messages are routed along the circle (comparable to a one-dimensional CAN). In the experiment, participants only show conditional reciprocity, i.e. Participant A cooperates on the same level with B as the predecessor of A in the circle C cooperates with A. This lead to a decrease of cooperation with increasing network size. In the presence of only one free-rider, the system collapses: the neighbors of the free-rider stop cooperating, then their neighbors etc. The situation changes if competition is introduced, i.e. if participants can forward their queries along different routing paths. Here, participants play cut-off strategies [20]. These strategies work independent of the size of the network and even in the presence of free-riders: On the one hand, a peer along a routing path is exchangeable and feels competition. On the other hand, each peer decides for which peers it forwards queries. If a peer is not trustworthy in the eyes

¹ Heterogeneity has never been a focal point in the analysis of the helping game.

of others, they do not process its queries any more. Hence, each peer should more or less be at least as cooperative as the average peer. This is because the cut-off value is likely to depend on the average behavior observed. Note that these elements of competition are present in structured P2P systems, but not in unstructured ones.

In structured P2P systems, cooperation and competition lead to reciprocal behavior. The outcome is efficient – if all peers have the same costs and benefits [21]. This is due to the structure of such systems: Although peers forward and answer messages on behalf of different peers, they only interact with a small group of peers that does not change, their contacts. Because of repeated interactions, peers can estimate the degree of cooperativeness of their contacts reliably. This leads to high levels of cooperation. However, these issues have only been investigated for the homogeneous case so far. It is currently unclear if these findings hold for the heterogeneous case as well.

Competitive scenarios lead to efficiency in the homogeneous as well as the heterogeneous case. This gives way to the assumption that structured P2P systems with heterogeneous peers are efficient as well.

5.3 Predictions for Our Experiments

Section 5.1 has presented argued that heterogeneity in structured P2P systems leads to inefficiency. Section 5.2 gives way to another conclusion. The motive ‘cooperation’ makes peers contribute in structured P2P systems. Competition in turn might be the motive for accepting heterogeneous outcomes. Put differently, a peer can be replaced, but a sequence of cooperative peers is necessary to process a request. All this calls for a holistic analysis.

Literature predicts that participants who earn less try to reduce the payoffs of participants who profit more. Think of a participant who benefits five times as much as all other participants. His contacts could enforce equality by only forwarding every fifth of his queries. In our setup however, we expect competition to be a stronger motive: As mentioned before, every participant is interested in having a good standing in the eyes of his contacts. If a participant drops several queries from another one, his standing becomes worse. Consequently, the probability that his queries in turn are processed properly goes down as well. Hence, we formulate the following hypothesis:

Hypothesis No Strategy Change: Strategies do not change under heterogeneity:

- a) Participants do not change their strategies towards a peer which earns more.

- b) Even if the distinguished peers, i.e., the ones whose utilities are different, are known, this does not result in different payoffs.

Hypothesis Payoff/Efficiency: Payoff and efficiency is the same in the homogenous and the heterogeneous case: The payoff of a participant (with fixed benefit) is the same in the homogeneous and in the heterogeneous case. Introducing heterogeneity among peers does not influence the payoff of participants.

Another reaction to heterogeneity might be that participants leave the system because they deem it unfair. I.e., they might stop issuing any queries. We do not expect such behavior. Namely, participation is beneficial when each participant looks at itself in isolation.

Hypothesis No-Leaves: Participants participate at the same level in homogenous and heterogeneous systems. I.e., they issue queries at the same rate in the heterogeneous system as in the homogeneous one.

6. Experiment Design

We have conducted economic experiments with human participants. In an experiment, a participant controls the strategy of one peer. Each participant decides whether the controlled peer forwards, answers or issues a query (for each query individually) and decides which contact to forward the query to².

We conducted the experiments in rounds. Each peer was allowed to issue one query per round. The participant decided whether to issue the query in each round. At the end of each round the participant was shown the current total payoff of his peer during the last round and whether the experiment continued. The properties of other participants were kept secret.

The experiment environment hides some internals of the structured P2P system from the participants: For instance, it manages the list of contacts. When forwarding or issuing a query, it calculates the distance of all contacts towards the query point. It then shows a list of contacts ordered by their distance to the query point. The order represents the probability that the peer has the desired query result.

We conducted the experiments with six participants each. Even if the participants realized that there were six peers only, this would not have influenced their behavior – in qualitative terms: The larger the system, the more different payoffs do peers have. In small systems, a con-

² A website containing screenshots, a detailed description of the game (the written instructions for the players) and the experiment results are downloadable from: <http://www.wipd.ira.uka.de/~schosser/2007-17/>

tact with higher payoffs is identifiable, and participants can readily ‘punish’ it. This is difficult in larger systems where interaction typically is via several peers. Further, conducting experiments with six participants is in line with the observation of Selten [22] that more than five humans show the same behavior as in large groups of any size..

To prevent communication among the participants, their terminals were separated from each other. Participants could not see the assignment of peers to other players.

In the beginning of each experiment, all participants were randomly seated in the laboratory. During the first 20 minutes we introduced the participants to the game: We handed out a description of the game in written form and played several test rounds. Then we played different treatments. A treatment consisted of twenty rounds each without discounting. We then rolled a ten-sided dice. If the dice showed 1, the game ended, otherwise it continued. This simulates a discounting rate of 0.1. This approach is common to diminish end game effect: Participants do not anticipate the end of the game and behave as if it continued.

In all treatments, peers controlled zones of the same size. All peers had the same number of contacts. Hence, they could expect the same number of incoming queries. The experiment environment randomly generated the queries participants could issue.

In all treatments the participants knew which share of their queries forwarded via a certain Peer i the system had answered. In other words, the experiment environment displayed the value α_i for all contacts.

$$\alpha_i = \frac{\text{number of query results sent via } i \text{ received}}{\text{number of own queries issued via } i}$$

However, while a participant knew which fraction of his queries sent via a specific peer had been answered, the environment did not reveal whether Peer i or any subsequent peer had forwarded/answered the query or not.

For each treatment, the participants initially had a balance of 100 points. When receiving a query result, the balance increased by 20 points. Issuing a query cost 2 points, forwarding 1 point and answering a query cost 5 points. This cost structure should reflect the costs and benefits in structured P2P systems: While receiving a query result imposes a big benefit, issuing and forwarding incurs small costs, at least from the perspective of the issuer: Only one message needs to be sent. Answering includes sending the query result to the issuer, hence the costs are higher. If the costs of forwarding, issuing and answering in sum were higher than the benefit of receiving query results, there obviously would not be any incentive to participate in the system.

We played three treatments. The first treatment was equal to the game described above (Standard Treatment).

I.e., the peers had standard payoffs. In the second treatment one peer was randomly chosen and received five times the payoff of the other peers (Treatment Luck). In this treatment the participants did not know the identity of the distinguished peer. (They knew that there was one such peer.) Our last treatment was equal to Treatment Luck, except that every participant knew the identity of the distinguished peer (Treatment Special).

After the treatments we conducted a strategy game [23]. Strategy games are common in economic experiments: The participants are asked to describe their strategy in own words. While they tend to learn and refine their strategies during the treatments, their behavior in the strategy game tends to be sophisticated. The strategies may depend on treatment parameters and on the history of the treatment. The participants are confronted with situations they know from the treatments, but on an abstract level. I.e., participants in the strategy game specify their behavior for all situations they might encounter. From a game-theoretic point of view, a strategy game lets us observe complete strategies.

After the experiment, we paid all participants depending on their success in the treatments. Their payoff corresponded to the points earned in the treatments (€ 2.00 per 100 points). Each participant earned € 18.10 on average.

7. Evaluation

In this section, we evaluate our hypotheses based on the treatments and the strategy games conducted.

7.1 Hypothesis No Strategy Change

To analyze whether participants change their strategies towards peers that earn much more, we look at the results of the strategy game. Here we asked the participants to describe their strategies towards the distinguished peer and towards the other ones.

Strategy	Category	# Players
Same strategy towards all peers	a)	38
Queries sent to distinguished peer preferably	b)	10
Queries never sent to distinguished peer	c)	2

Table 1: Behavior Changes When Issuing

Tables 1 and 2 list the results. We used three different groups of strategies to classify the behavior observed:

- a) Same strategy towards all peers independent of their payoffs.
- b) More cooperative strategy towards the distinguished peer.

- c) Less cooperative strategy towards the distinguished peer.

In most cases the strategies towards the distinguished peer are the same as towards all other peers (Category a)). While several participants prefer to issue their queries via the distinguished peer (Category b)), only few chose to never send any queries through it (Category c)). This changes when it comes to forwarding and answering. A majority of peers still does not make the distinction. Among the other peers, less prefer to cooperate with the distinguished peer over other peers compared to issuing (Category b)). Furthermore, more participants show uncooperative behavior towards the distinguished peers (Category c)). I.e., there is a small tendency among the participants to punish the peer earning more. But this tendency is behind any statistical significance. On average, the strategies do not change. Even in case of answering, where more participants change their behavior towards the distinguished peer compared to issuing and forwarding, a binomial test confirms this on significance level of 1%.

Strategy	Cat.	# Players	
		Answer	Forward
Same strategy towards all peers	a)	34	37
Queries of distinguished peer are handled preferred to others	b)	4	5
Queries of distinguished peer are not handled	c)	12	8

Table 2: Behavior Changes When Answering/Forwarding

In addition, we analyze whether the payoff of the distinguished peer changes between Treatment Luck, where this peer is not known, and Treatment Special. Table 3 shows the results. The payoff decreases in half of the games, and it increases in the other half of the games. I.e., we cannot confirm at an acceptable significance level that the payoff increases or decreases.

7.2 Hypothesis Payoff/Efficiency

To learn whether introducing heterogeneity influences the efficiency of the system, we analyzed the payoffs of all participants with standard payoffs in Treatment Standard and Treatment Luck (see Table 3).

While the payoff of all peers with standard payoffs increases in half of the treatments and decreases within the other half, we cannot confirm that the payoff increases or decreases at a certain significance level. I.e., although the participants in the luck treatment know that another peer earns a lot more than they do, they keep playing as in the homogeneous treatment.

	Standard Peers		Distinguished Peer	
	Treatment Standard	Treatment Luck	Treatment Luck	Treatment Special
1	9,16	8,72	53,87	53,33
2	7,40	10,57	87,33	22,80
3	7,01	7,65	88,40	66,53
4	9,21	5,81	42,27	42,67
5	5,56	5,67	80,73	53,87
6	9,50	8,71	75,27	81,73
7	8,63	10,25	46,13	67,60
8	8,49	10,36	73,87	73,60
9	6,01	3,49	12,60	61,00
10	6,61	1,92	23,40	30,47

Table 3: Payoffs

Treatment Special confirms these results as well. See the web page with the experiment results for details.

7.3 Hypothesis No Leaves

To analyze whether participants tend to leave the system when heterogeneity is introduced we compare the number of queries issued per standard peer in the different treatments. See Table 4.

	Treatment Standard	Treatment Luck	Treatment Special
Group 1	14,17	15,40	17,60
Group 2	14,50	18,60	19,60
Group 3	16,33	17,20	18,20
Group 4	16,50	17,80	18,00
Group 5	14,00	17,20	19,00
Group 6	16,50	16,00	18,60
Group 7	16,83	19,40	19,60
Group 8	15,83	19,00	19,20
Group 9	14,67	16,00	15,80
Group 10	14,83	11,60	15,60

Table 4: Number of Queries Issued

The number of queries issued in the standard treatment is higher than in the other treatments for one group only. I.e., despite the fact that one peer earns more than the others, peers keep participating in the system.

8. Conclusions

P2P systems are heterogeneous, i.e., peers have different benefits and cost structures. We have studied the effect of heterogeneity on efficiency in structured P2P

systems. Efficiency is an important design objective. It gives peers an incentive to participate.

A key aspect of our study is how structured P2P systems actually reconcile the two different behavioral motives cooperation and competition. We use behavioral experiments to investigate this. Behavioral experiments are the means of choice – standard economic models assuming only selfish actors do not always predict actual behavior well if cooperation is concerned. Our results are as follows: Efficiency in structured P2P systems does not decrease, the strategies of the peers do not change, no additional leaves occur, and cooperation persists if we compare the homogenous to the heterogeneous case. This shows that structured P2P systems are systems that can actually reconcile the two behavioral motives cooperation and competition. This result is important: It means that existing protocols already help to reach efficiency under heterogeneity.

9. References

- [1] Adar, E. & Huberman, B.: Free Riding on Gnutella *Xerox PARC*, 2000.
- [2] Blanc, A. et al.: Designing Incentives for Peer-to-Peer Routing, *Proc. of the IEEE Infocom*, 2005
- [3] Buchmann, E. & Böhm, K.: FairNet - How to Counter Free Riding in Peer-To-Peer Data Structures, *Proc. of the International Conference on Cooperative Information Systems*, 2004.
- [4] Buragohain, C. et al.: A Game Theoretic Framework for Incentives in P2P systems, *Proc. of the International Conference on Peer-to-Peer Computing*, 2003.
- [5] Fehr, E. & Falk, A.: Wage Rigidity in a Competitive Incomplete Contract Market, *Journal of Political Economy*, 107, 106-134, 1999.
- [6] Fehr, E. & Schmidt, K. M.: A Theory of Fairness, Competition, and Cooperation, *The Quarterly Journal of Economics*, 114 (3), 817-868, 1999.
- [7] Feldman, M. et al.: Quantifying Disincentives in Peer-to-Peer Networks, *Workshop on Economics of Peer-to-Peer Systems*, 2003
- [8] Feldman, M. et al.: Robust Incentive Techniques for Peer-to-Peer Networks, *Proc. of the ACM Conference on Electronic Commerce*, 102-111, 2004.
- [9] Greiner, B. & Levati, V.: Indirect reciprocity in cyclical networks – An experimental study, *Journal of Economic Psychology*, 26 (5), 711-731, 2005.
- [10] Grosskopf, B.: Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition, *Experimental Economics*, 6 (2), 141-158, 2003.
- [11] Hagen, E. & Hammerstein, P.: Game theory and human evolution: A critique of some recent interpretations of experimental games. *Theoretical Population Biology*. 69, 339-348, 2006.
- [12] Isaac, R. & Plott, C.: Price Control and the Behavior of Auction Markets: An Experimental Examination, *American Economic Review*, 71, 448-459, 1981.
- [13] Kamvar, S. et al.: Incentives for Combatting Freeriding on P2P Networks, *Proc. of the International Euro-Par Conference*, 1273-1279, 2003.
- [14] Kamvar, S. et al.: The EigenTrust Algorithm for Reputation Management in P2P Networks, *Proc. of the WWW*, 2003.
- [15] Kung, H. & Wu, C.: Differentiated Admission for Peer-to-Peer Systems: Incentivizing Peers to Contribute their Resources, *Proc. of the Workshop on Economics of Peer-to-Peer Systems*, 2003.
- [16] Ngan, T. et al.: Enforcing Fair Sharing of Peer-to-Peer Resources, *Proc. of the International Peer to Peer Symposium*, 2003.
- [17] Ramasamy, L. & Liu, L.: Free Riding: A New Challenge to Peer-to-Peer File Sharing Systems, *Proc. of the International Conference on System Sciences*, 2003.
- [18] Ratnasamy, S. et al.: A Scalable Content-Addressable Network, *Proc. of the Conference on Applications, Technologies, Architectures and Protocols for Computer Communications*, 161-172, 2001.
- [19] Saroiu, S. et al.: A Measurement Study of Peer-to-Peer File Sharing Systems, *Proc. of Multimedia Computing and Networking*, 2002.
- [20] Schosser, S. et al.: Incentives Engineering for Structured P2P systems - a Feasibility Demonstration Using Economic Experiments, *Proc. of the ACM Conference on Electronic Commerce*, 280-289, 2006.
- [21] Schosser, S. et al.: Indirect Partner Interaction in P2P Networks - Stimulating Cooperation by Means of Structure, *Proc. of the ACM Conference on Electronic Commerce*, 2007.
- [22] Selten, R.: A simple model of imperfect competition where four are few and six are many, *International Journal of Game Theory* 2, 141-201, 1973.
- [23] Selten, R. Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen einer Oligopolexperimente, *Beiträge zur experimentellen Wirtschaftsforschung*. Heinz Saueremann (ed.), 136-168, 1967.
- [24] Seinen, I. & Schram, A. Social Status and Group Norms: Indirect Reciprocity in a Helping Experiment. *European Economic Review*, 50, 581-602, 2001.
- [25] Smith, V.: Microeconomic Systems as an Experimental Science, *The American Economic Review*, 72 (5), 923-955, 1982.
- [26] Smith, V. & Williams, A.: The boundaries of competitive price theory: convergence, expectations, and transaction costs, *Bargaining and Market Behavior: Essays in Experimental Economics*, 286-319, 2000.
- [27] Rowstron, A. & Druschel, P.: Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems, *Proc. of the IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, 329-350, 2001.
- [28] Stoica, I. et al.: Chord: a scalable peer-to-peer lookup protocol for internet applications, *IEEE/ACM Transactions on Networking*, 11, 17-32, 2003.
- [29] Varian, H.: *Intermediate Microeconomics*, W. W. Norton and Company, 5. Edition, 1999.