Competition vs. Fairness – Analyzing Structured Networks by Means of User Experiments

Stephan Schosser
Universität Karlsruhe (TH)
schosser@gmx.de

Klemens Böhm
Universität Karlsruhe (TH)
boehm@ira.uka.de

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

Abstract

We investigate how to ensure efficiency (in the economic sense of the word) in structured networks, with a focus on heterogeneity. A network is structured if the network designer has predefined some relationships between individuals (aka. nodes). Structured networks have turned out to be surprisingly efficient – at least as long as nodes face the same costs and benefits, i.e., are homogeneous [25]. However, the homogeneity assumption is unnatural and restrictive. Economic experiments in general (not with a focus on structured networks) suggest that heterogeneity is in the way of efficiency, i.e., reduces the sum of all payoffs. This is because individuals favor outcomes where everybody earns the same. This paper describes behavioral experiments that investigate this issue, i.e., the influence of heterogeneity on efficiency in structured networks. We show that most nodes in structured networks cooperate even if they earn less than others. Our explanation is that – with our design – competition enhances cooperation. This effect is rarely observed with other networks as well as in other, less specific settings where competition is in the way of cooperation. This result is an important step towards establishing networks that yield more tangible payoffs for its nodes.

1. Introduction

With the proliferation of the internet, the question how to ensure efficiency (in the economic sense of the word) in networks is fundamental. In networks, nodes are connected with other nodes they can interact with. A node may be a human or a system node whose behavior is controlled by a human. So the question how humans behave in networks is important. *The network is efficient* if all nodes behave cooperatively in any interaction. Given that nodes maximize their own income, i.e., behave competitively, networks often do not reach high levels of cooperation and therefore efficiency. If nodes behave competitively, economic theory predicts that no exchange of services will take place. Thus, how to design networks where nodes are motivated to contribute and continue to do so?

Behavioral experiments [26] show that structure can increase cooperativeness in networks and help to overcome the negative effects of competition on efficiency. 'Structure' means that nodes cannot freely choose their interaction partners; instead, some authority has assigned them, at least to some degree. In online communities, a system designer obviously is in the position to specify the structure (even though discussing concrete applications is beyond the scope of this paper).

1.1 Benefits of Structure

A core contribution of this paper is to offer a plausible explanation for this important finding (the one that structure incurs cooperation – previous work has only observed the phenomenon [26], but has not offered an explanation) and to substantiate it by means of behavioral experiments. To obtain such a deeper understanding, this paper will show that an investigation of the heterogeneous case is helpful. I.e., how does a node behave if it interacts with another node whose benefit from a service executed is significantly higher? Two motives might play a role: fairness and competition. (1) Fairness means that individuals favor outcomes where everybody earns the same. Ensuring fairness is problematic in the presence of heterogeneity. Here, the payoff for all individuals is equal only if the system processes less service requests issued by 'privileged' individuals. But on the other hand, this reduces efficiency. (2) While heterogeneity is a problem of two person interaction, it is not a problem in competitive markets. If persons behave competitively, the outcome of the market can be very heterogeneous: The better trader gets more since he is a better trader. In this sense competition is a stabilizing factor. Both motives, fairness and competition between nodes, might influence node behavior in networks.

Example 1: Think of a network where nodes offer services, and other nodes want to consume these services. A node is both client and server – it offers one or several services and at the same time wants to consume services offered by others. We assume that services offered are evenly distributed over the network, i.e., it is not the case that some nodes offer ‚better' or more attractive services than others. We distinguish between *issuers*, *intermediate nodes* (aka. *forwarders*), and *responders*: An issuer issues a service request and sends it to another node it is connected to. This recipient typically is an intermediate node, i.e., it may forward the request, be it to another intermediate node, be it to the responder. A responder may then process the service request. We assume that there always are several potential forwarders, i.e., nodes $N_1, \ldots, N_l$ where a node, subsequently referred to as N, can forward the service request to. Our assumption is that in structured networks – in contrast to networks without structure – there always is a minimal number k of potential forwarders, i.e., $l \geq k$, and $k \geq 2$. Node N_i, the forwarder chosen by N, has to take a decision: Should it forward the service request or ignore it? In this context, several aspects play a role: *(a)* Ignoring a message reduces the benefit of another node. In this way, N_i can ensure fairness. Suppose that N earns 5 times the amount N_i earns with a service execution, and N_i knows this. To reach equal, i.e., fair payoffs, it only processes every fifth request obtained from N. *(b)* Suppose that N_i forwards significantly less service requests obtained from N than N_j ($i \neq j$). N will observe this eventually. A natural reaction of N will be that it likewise does not forward any/forwards fewer requests obtained from N_i. This will either harm N_i directly (if N_i has issued the request), or, if N_i has been a forwarder, bring down N_i's reputation with its other contacts. Thus, to avoid such behavior by N, N_i should be about as reliable as N_j – N_i is in competition. Since, by definition, the structure of structured networks is given, at least to some degree, N_i cannot defy this competition offhand, i.e., without leaving the system at least temporarily. ▪

The research question of this paper is: Which behavioral motive – fairness or competition – is stronger? Suppose that N profits much more from the processing of a request it has issued, compared to others. Given that many individuals value fairness highly [9], will N_i 'punish' N? Or will N_i cooperate with N, in the spirit of competition? – In homogeneous systems where fairness is more or less given, this conflict (the one between fairness and competition) does not occur. Hence, structured networks that are heterogeneous allow to study how strong (in qualitative terms) competition is as a behavioral motive.

1.2 Method

We now briefly discuss which method is appropriate to address the questions raised above. The following methods have been applied in the past to analyze networks: (a) game theory, together with a cost model, i.e., formal analysis, (b) simulations, (c) behavioral experiments, and (d) observation of node/participant behavior in existing real-world systems.

(a) is not feasible here because it assumes rational behavior. However, existing game-theoretic models cannot explain fairness, i.e., the assumption is too strong. The rationality assumption does not necessarily hold in networks. We can observe this in existing settings, e.g., P2P filesharing: Game-theoretic models predict that nodes free-ride [11][16]. Some cooperate nevertheless, according to measurement studies [1].

Regarding (b), which behavior exactly is it which we could simulate? If we knew which behavioral motive is stronger, we could of course explore efficiency and stability of the network by means of simulations, but this question needs to be explored first. (d) is not an option either: We are not aware of any real-world network that is structured, and nodes are autonomous. From a slightly different perspective, our objective is complementary to the one of measurement studies on existing systems, as we try to foresee system characteristics at development time. Further, (d) does not let us control all factors which could influence node behavior. A certain network may have specific characteristics, e.g., it is new, and participation is fashionable, or there may be other (informal) networks which exist in addition, beyond our control. We need to rule out that such effects influence results.

The approach of our choice is (c). We derive hypotheses from known patterns of human behavior and evaluate them by means of experiments with humans. This method (behavioral experiments) is widely used by experimental economists, see e.g. [3][8][29], when existing economic models do not explain human behavior. A slight disadvantage might be that one frequently can only obtain qualitative results, at least when system complexity goes beyond a certain point [19].

1.3 Results

Our main findings are as follows: For structured systems with heterogeneous nodes the desire for fairness does not play much of a role. Competition among nodes is the dominant behavioral motive and fosters cooperation. In our own previous work [25][26], we have shown that structure fosters cooperation – in the homogeneous case. There, 91.7% of the individuals play reciprocal cut-off strategies. I.e., they cooperate with someone if he has been cooperative in the past. Here, we show that the degree of cooperation in the heterogeneous case is roughly the same as in the homogeneous one. In this way, we show that structure leads to cooperation in the heterogeneous case as well.

Paper outline: Section 2 reviews related work. We describe the basic concepts of structured P2P systems in Section 3. In Section 4 we derive hypotheses concerning our experiments, which we evaluate in Section 6 using the experiment design introduced in Section 5. Section 7 features a discussion. Section 8 concludes.

2. Related Work

Measurement studies of existing P2P systems have shown that nodes in such networks tend to behave uncooperatively [1][24]. They download content offered by others without contributing to the system. [22] shows that such behavior, called free-riding, can lead to the break-down of the system. Hence, ensuring cooperation is essential in such systems.

From an economic perspective, competition influences the behavior of interaction partners. In competitive situations, the degree of cooperation tends to be low in general. The effect of this antagonism is that ensuring efficiency is difficult in many situations. For instance, think of unstructured P2P systems for file sharing: Nodes compete for bandwidth of others to download files [10]. Such systems cannot guarantee efficiency and therefore high payoffs. One exception are feedback systems. For them, Bolton et al. [3] recently showed that adding competition can increase efficiency.

Smith observes that competitive markets exist where selfish behavior of individuals maximizes the overall payoff (i.e., efficiency) [30]. However, subtle design issues can have a significant effect on the outcome and can even decrease efficiency: 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 or face different costs from their (same) actions[1]. Humans try to ensure fairness [7][14], i.e., ensure that the payoffs are equal. This usually leads to a further decrease in efficiency [7]. However, while heterogeneity is a problem of two-person interaction, it is not a problem in markets that face only competition. Here, competition ousts fair behavior [12], i.e., even 'fair-minded' people behave selfishly under competition.

3. Structured Peer-to-Peer Systems

This section is a short overview of structured P2P systems and Content-Addressable Networks (CAN) in particular. In our analysis, we will focus on CAN as an instance of structured networks. Here, the service requests mentioned are simply queries. Even though our experiments are based on the grid structure, we expect our findings to be applicable to other structured networks as well, as we will explain at the end of this section.

Structured P2P systems such as CAN [23], Chord [32] or Pastry [31] maintain large sets of (key, value)-pairs. A hash function that is commonly known maps each key to a certain location in the key space. The key space is divided into peer zones. Each node is responsible for one zone, i.e., administers all pairs whose key falls into its zone. Next to the zone, each node administers a list of nodes with adjacent zones and their zones. When a node issues a query (asks for a certain key), it hashes the key to the corresponding location in the key space, the query point. If the query point lies in its zone, the node itself already has the query result, namely the value. Otherwise, it forwards the query to the neighbor node closest to the query point. This recurs until the query finally reaches the node having the query result. This node then sends the query result to the initial sender.

Example 2: Figure 1 shows a CAN. The keys are transformed to two-dimensional coordinates using a hash function known to all nodes. The (key, value)-pair ("Mars", "Chocolate Bar") is mapped to the coordinates (0.45, 0.3). Each square in the figure represents a peer zone. I.e., Node N_x stores the value corresponding to ("Mars"). If Node N' is interested in information concerning "Mars", it can forward its query using one of its contacts: Node N_1, N_2 or N_3. Therefore, "Mars" is first transformed into the coordinates (0.45, 0.3) using the public hash function. The query is then forwarded to the node closest to (0.45, 0.3), Node N_2. Node N_2 does not administer the key "Mars", hence it forwards the query to one of its contacts until the query arrives at Node N_x. Node N_x finally sends the query result to Node N'. ▪

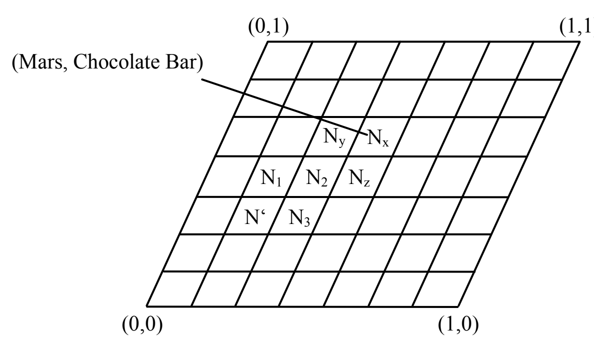

Figure 1: Content-Addressable Network

Structured P2P systems have two important properties: First, all nodes interact with their contacts only. Second, several nodes have to cooperate to answer a query. Under realistic assumptions, if only 5% of the nodes are uncooperative, 40% of the queries in such a system remain unanswered [4].

[1] We use the term 'payoff' when referring to the total income of a participant over the entire lifetime, whereas benefits and costs are characteristics of single actions.

To distinguish between cooperative and uncooperative nodes, each node can monitor the behavior of its contacts. More specifically, it observes if the messages it has sent to a neighbor have been answered. (If not, it does not know whether its neighbor or a subsequent node have dropped the message.) Based on its observations, each node can decide for itself which contacts it deems cooperative. Further, it can decide whether to process a request received from a contact and which contact to forward it to, when necessary. Experiments in the homogeneous case have shown that nodes forward queries only for nodes they deem cooperative [25].

We assume that each node benefits when receiving answers for queries, while it faces costs when forwarding, issuing or answering one. In what follows, we only consider scenarios where the benefit of obtaining an answer for a query exceeds the sum of the costs of processing it.

Our findings should be applicable to any system allowing for indirect partner interaction and competition. *Indirect partner interaction* means that (a) an individual provides services on behalf of strangers, and (b) direct interaction only occurs between nodes which interact frequently, i.e., partners [26]. In Chord for instance, if nodes are autonomous, indirect partner interaction takes place: Each node can access the information stored by any other node (stranger) by only interacting with the nodes in its finger table (partners). Competition takes place as well: If a node deems another node untrustworthy, it can forward its queries via a third node in the finger table in most cases.

4. Hypotheses

From an economic perspective, the designers of structured networks face a market-design problem. The system is a market of services. Each node controls a fraction of the services in the system and can request some services. An objective of market design is Pareto-efficiency [33]. An allocation is Pareto-efficient if nobody can increase his payoff without someone else 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 structured network: From his perspective, systems are optimal if the sum of the payoffs of all nodes is maximal. This is the case if all nodes fully cooperate, and all requests are processed properly.

Whether efficiency is reached depends on several design issues. In Section 4.1 we discuss the impact of fairness on efficiency, before we reason about the impact of competition in Section 4.2. Using these discussions we derive our hypotheses in Section 4.3.

4.1 Cooperation in Two-Person Interaction

To better understand structured networks which consist of two-person interactions (e.g., a message is forwarded from one node to another one) and the evolvement of cooperation over time, we use the helping game [29] as an example: One participant has the role of a donor, while the other one 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 donate. But behavioral experiments [29] show that some participants (approximately 22%) actually do. The motivation for such cooperative behavior is called fairness. But fairness also implies that participants favor the same payoffs for all. Literature predicts [9] that this is counterproductive in case of heterogeneity, 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 [9]. 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 accordingly[2]. This reasoning predicts a decrease in efficiency for structured networks that are heterogeneous.

4.2 Competition

In systems without competition cooperation tends to decrease with the size of the network: [13] describes an experiment where nodes form a circle, and service requests 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 C of A in the circle cooperates with A. This leads 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 nodes can forward requests along different routing paths.

Example 3: Node N' wants to issue a query (see Figure 1). It can forward it to one of its neighbors (N_1, N_2 or N_3). If N' chooses to forward the query to N_1, N_1 has to decide whether to forward it or not. The reasoning from Example 1 is applicable here as well. N_1 is in competition with N_2 and N_3. In a structured P2P system without

[2] Heterogeneity has never been a focal point in the analysis of the helping game.

replication only one node exists which can answer a given query (N_x in the example). Nevertheless, potential responders face competition as well. N_y can observe the fraction of queries answered which it had forwarded to N_2, N_x and N_z. If this fraction for N_x lies significantly below the one of the other neighbors of N_y, N_y is likely to be less cooperative with N_x. ▪

In presence of competition, nodes play reciprocal cut-off strategies [25]. (When Node A plays a *cut-off strategy*, it cooperates with Node B if B's degree of cooperation in the eyes of A exceeds a certain cut-off value.) These strategies work independently of the network size and even in the presence of free-riders. Reciprocal cut-off strategies are beneficial in structured networks where partners interact, but not in unstructured ones where interaction primarily is between strangers.

4.3 Predictions for Our Experiments

Section 4.1 has argued that heterogeneity in structured networks leads to inefficiency. Section 4.2 gives way to another conclusion. The motive 'cooperation' makes nodes contribute in structured networks. Competition in turn might be the motive for accepting heterogeneous outcomes. Put differently, a node can be replaced, but a sequence of cooperative nodes is necessary to process a request. All this calls for a holistic analysis.

In our setup, we expect competition to be the stronger motive: As mentioned, an individual is interested in having a good standing in the eyes of his contacts. If a node drops several queries from another one, his standing becomes worse. Consequently, the probability that its queries in turn are processed properly goes down as well. Hence, we formulate the following hypotheses:

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

a) Nodes do not change their strategies towards a node which earns more.
b) Even if the distinguished nodes, 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 node (with fixed benefit) is the same in the homogeneous and in the heterogeneous case. Introducing heterogeneity among nodes does not influence the payoff of nodes.

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

Hypothesis No-Leaves: Nodes 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.

5. Experiment Design

To analyze how efficiency is reached under heterogeneity, we use behavioral experiments. In these experiments, a participant controls the strategy of one node in a CAN. Each participant decides whether the controlled node forwards, answers or issues a service request/query (for each request individually) and decides which contact to forward the request to.

5.1 Experiment Environment

We conducted the experiments in rounds.[3] Each node was allowed to issue one service request per round. The participant decided whether to do so in each round. For each incoming query, a participant could choose whether to process the query (i.e., forward or answer it) or not. If he decided to process it, he was shown a list of all nodes he could forward the query to.

In all treatments, nodes controlled zones of the same size. All nodes had the same number of contacts. Hence, they could expect the same number of incoming requests. The experiment environment randomly generated the service requests for the participants.

In all treatments the participants knew which share of their requests forwarded via a certain Node *i* the system as a whole had answered. 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 the requests sent by his node via Node *i* had been answered, the environment did not reveal whether *i* or any subsequent node had forwarded/answered the request or not.

At the end of each round the participant was shown the current total payoff of his node during the last round and whether the experiment continued. The properties of other nodes were kept secret.

The experiment environment hides some internals of the network from the participants: For instance, it manages the list of contacts. When forwarding or issuing a query, it calculates the distance of all contacts to the node which could answer the query. It then shows a list of contacts ordered by their distance to this contact, together with some statistics on their past behavior (see the screenshots on our web page for details).

[3] See http://www.ipd.uni-karlsruhe.de/~schosser/wi08/ for screenshots, a detailed description of the game (the written instructions for the players) and detailed results.

We conducted the experiments with six participants each and repeated them with ten groups, i.e., 60 individuals have participated in our experiments. To prevent communication among the participants, their terminals were separated from each other. They could not see the assignment of nodes to other players.

5.2 Experiment Procedure

In the beginning of each experiment, participants were randomly seated in the laboratory. During the first 20 minutes we introduced them 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. The game ended if it showed 1, otherwise it continued. This simulates a discounting rate of 0.1. This approach is common to diminish end game effects: Participants do not anticipate the end of the game and behave as if it continued.

For each treatment, the nodes initially had a balance of 100 points. When receiving an answer for a request, the balance increased by 20 points. Issuing a request cost 2 points, forwarding 1 point and answering a request cost 5 points. This cost structure should reflect the costs and benefits in structured networks: While receiving an answer for a service request result imposes a big benefit, issuing and forwarding incurs small costs, at least from the perspective of the issuer: Only one request message needs to be sent. Answering includes conducting the service and sending the 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 answers for service requests, there obviously would not be any incentive to participate in the system.

We played three treatments. The first one was equal to the game described above (Baseline Treatment). I.e., the nodes had standard benefits. In the second treatment one node was randomly chosen and received five times the benefit of the other ones (Hidden Information Treatment). In this treatment the participants did not know the identity of the distinguished node. (They knew that there was one such node.) Our last treatment was equal to the Hidden Information Treatment, except that every participant knew the identity of the distinguished node (Information Treatment).

After the treatments we conducted a strategy game [28]. Such games are common in economic experiments: The participants are asked to anonymously 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 mature. The strategies may depend on treatment parameters and on the history of the treatment. 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. In sum, we paid the participants € 1086.00 for participating in the experiments.

5.3 Discussion of Experimental Setup

The exogenous parameters of our setup are (a) costs of issuing, forwarding, processing a service request; benefit of a service request processed, (b) degree of heterogeneity, (c) structure of the network, i.e., number of neighbors and size of the network. Note that the level of competition is an endogenous parameter and depends on node behavior. In the following, we say why an investigation with the values mentioned above should give way to insights, and why we do *not* deem an investigation of other values promising.

The cost structure (parameters of Category (a) and 'size of the network') is not important as long as the benefit of a service exceeds the total costs (i.e., costs incurred for all nodes) significantly. (A parameter setting where the costs exceed the benefit is not interesting.) [25][26] give way to this conclusion – the decision which one of the few possible strategies a node plays only depends on this difference and not on any other characteristics of the exogenous parameters. [25][26] features experiments with different parameter values. Regarding the number of neighbors, participants play cutoff strategies in the homogeneous case, with each neighbor separately, according to [25][26].

Both the group size and the overall number of participants are characteristics of behavioral economics experiments in general, e.g., [2][5][6]. The number of participants mentioned is sufficient to reach statistical significance in our context, as we will show. Further, six participants per experiment are sufficient, for various reasons: First, even if the participants realized that there were six nodes only, this would not have influenced their behavior – in qualitative terms: Our experiments consider an extreme case where exactly one participant has a significantly larger benefit than all others. Almost all participants have him as neighbor and can punish him. In one treatment everybody even knows him explicitly. Therefore, trying to be fair in the presence of heterogeneity is more difficult in larger systems. Put differently, if we can show for our extreme case that heterogeneity has no negative effect on efficiency, this will also be the case in larger systems. Further, conducting experiments with six participants is in line with the re-

sult by Selten [27] that more than five humans show the same behavior as in larger groups.
Why does our evaluation rely on the strategy game? Observing the user behavior in the explicit games does not yield complete strategies. Further, our own experience shows that it simply contains too much noise. Finally, we told the participants that compensation would depend on the strategies from the strategy game – we therefore can expect truthful strategy descriptions in this game. (We ended up paying the participants contingent on their performance in the explicit games, but participants have not been aware of this.)

6. Evaluation

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

6.1 Hypothesis No Strategy Change

To analyze whether participants change their strategies towards nodes that earn much more, we look at the results of the strategy game of the Information Treatment. Here we asked the participants to describe their strategies towards the distinguished node and towards the other ones. We only analyze 50 strategy games: We ignore the answers of participants who controlled distinguished nodes.
Tables 1 and 2 list the results. We use three different groups of strategies to classify the behavior observed:

a) Same strategy towards all nodes independent of their payoffs.
b) More cooperative strategy towards the distinguished node.
c) Less cooperative strategy towards the distinguished node

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

In most cases the strategies towards the distinguished node are the same as towards all other nodes (Category a)). While several participants prefer to issue their queries via the distinguished node (Category b)), only few chose to never send any requests through it (Category c)). This changes when it comes to forwarding and answering. A majority of nodes still does not make the distinction. Among the other nodes, fewer nodes prefer to cooperate with the distinguished node over other nodes, compared to issuing (Category b)). Furthermore, more participants show uncooperative behavior towards the distinguished nodes (Category c)). I.e., there is a small tendency among the participants to punish the node 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 node compared to issuing and forwarding, a binomial test confirms this on a 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 node changes between the Hidden Information Treatment, where this node is not known, and the Information Treatment. 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.

6.2 Hypothesis Payoff/Efficiency

To learn whether heterogeneity influences the efficiency of the system, we analyzed the payoffs of all participants with standard payoffs in the Baseline Treatment and the Hidden Information Treatment (see Table 3).
While the payoff of all nodes with standard benefits increases in half of the treatments and decreases in the other half, we cannot confirm that the payoff increases or decreases at a certain significance level. I.e., although the participants in the Hidden Information Treatment know that another node earns a lot more than they do, they keep playing as in the homogeneous case. Thus, heterogeneity does not decrease efficiency. The result is even more obvious when analyzing the Information Treatment. Here, the payoffs of the peers increase in seven out of ten groups.
These results also indicate that structured networks are more efficient than unstructured ones: Measurement studies of unstructured systems [1] show that 75% of the participants resort to free riding. Studies of unstructured systems with countermeasures against free riding show [34] that 75% of the participants still behave in a way similar to free riding: They only contribute the absolute minimum enforced by the system. If the participants were free-riding to this degree, we could not have observed the high payoffs in our treat-

ments (compare the values in Column Info. Treatm. in Table 3 to the theoretical maximum in a system with cooperative nodes only and without forwarding: 20 (benefit of result) – 5 (cost of answering) – 2 (cost of issuing) = 13). Table 4 shows the fraction of queries processed. From these numbers, we conclude that structure increases cooperation (as we have shown for the homogeneous case [25][26]). In the latter ones, only 12 out of 50 participants (24%) showed a tendency towards free-riding against the distinguished peer (see Table 2). We could not analyze whether the participants show free-riding behavior during the treatments, due to noise and learning effects. Finally, we did not observe any free-riding towards nodes with the same payoff in the strategy games. This is in line with our previous results [25][26],

	Standard Peers			**Distinguished Peer**	
	Baseline Treatm.	**No Info Treatm.**	**Info Treatm.**	**No Info Treatm.**	**Info Treatm.**
1	9.16	8.72	10.21	53.87	53.33
2	7.40	10.57	11.87	87.33	22.80
3	7.01	7.65	9.01	88.40	66.53
4	9.21	5.81	9.17	42.27	42.67
5	5.56	5.67	7.89	80.73	53.87
6	9.50	8.71	7.43	75.27	81.73
7	8.63	10.25	9.67	46.13	67.60
8	8.49	10.36	10.37	73.87	73.60
9	6.01	3.49	7.95	12.60	61.00
10	6.61	1.92	5.27	23.40	30.47

Table 3: Payoffs Per Round

The Information Treatment confirms these results as well. See the web page with the results for details.

6.3 Hypothesis No Leaves

To analyze whether participants tend to leave the system when heterogeneity is introduced we look at the number of queries issued per node. Here, we do not consider the distinguished node. See Table 4.

	Baseline Treatment		**No Information Treatment**		**Information Treatment**	
	Issued	**Handled**	**Issued**	**Handled**	**Issued**	**Handled**
1	14.17	85.9 %	15.40	74.2 %	17.60	72.2 %
2	14.50	64.4 %	18.60	81.4 %	19.60	73.2 %
3	16.33	62.2 %	17.20	62.9 %	18.20	69.4 %
4	16.50	75.8 %	17.80	54.7 %	18.00	64.8 %
5	14.00	57.1 %	17.20	55.7 %	19.00	57.1 %
6	16.50	76.8 %	16.00	74.0 %	18.60	64.6 %
7	16.83	69.3 %	19.40	69.2 %	19.60	69.5 %
8	15.83	72.6 %	19.00	78.3 %	19.20	75.9 %
9	14.67	60.2 %	16.00	44.3 %	15.80	75.5 %
10	14.83	60.7 %	11.60	33.3 %	15.60	47.8 %

Table 4: # Queries Issued and Fraction Processed

The number of requests issued in the Baseline Treatment is higher than in the other treatments for one group only. I.e., despite the fact that one node earns more than the others, nodes keep participating in the system with the same intensity.

7. Discussion

We now address the relevance of our results for system design from a computer-science perspective. We do so by describing two plausible incentive mechanisms for CAN that differ regarding one specific design issue and explaining that our study helps to choose between them. This design alternative is just a simple example; there are further alternatives where our current work is insightful. Two reputation mechanisms for CAN are conceivable: (a) The cut-off value is predefined, i.e., fairness, and (b) a peer can fix the cut-off value itself, e.g., to the average degree of cooperation of its neighbors (competition). Several real-world file-sharing systems [24][34] have implemented a mechanism similar to (a), i.e., this current discussion is not constructed. (a) does not expose the nodes to competition. A participant is likely to favor fair outcomes, by trying to contribute as much as the others. Efficiency is low, and participants are unlikely to contribute more than they have to. Our experiments suggest that introducing competition into the system, i.e., Alternative (b), could change this. Nodes could quantify the degree of cooperation of a neighbor as its contribution divided by its consumption. A node that wants to consume more services also has to contribute more. With Alternative (b), one would have an incentive to contribute more, leading to competition and to more cooperation.

8. Conclusions

Structured networks are heterogeneous, i.e., nodes have different benefits and cost structures. Efficiency is an important design objective. We have studied the effect of heterogeneity on efficiency in one specific structured network. A key aspect of our study is how competition and fairness affect the efficiency of networks. We use behavioral experiments to investigate this. Such experiments are the means of choice – standard economic models assuming only selfish actors fall short when it comes to predict the degree of cooperation in scenarios similar to the one studied here: For instance, several game-theoretic models predict free-riding of all participants in peer-to-peer filesharing networks. Nevertheless, measurement studies show that a fraction of the participants actually contributes [1].

Our results are as follows: Efficiency in structured networks does not decrease, the strategies of the nodes do not change, no additional leaves occur, and cooperation persists if we compare the homogenous to the heterogeneous case. This shows that structured networks are systems that actually benefit from competition. This result is important: It means that structure can help to reach efficiency under heterogeneity, and that fairness considerations do not influence the behavior of the participants.

9. References

[1] Adar, E. and Huberman, B.: Free Riding on Gnutella, *Xerox PARC*, 2000.

[2] Berninghaus, S. et al.: Conventions and Local Interaction Structures: Experimental Evidence, In: *Games and Economic Behavior*, 39 (2), 177-205, 2002.

[3] Bolton, G. et al.: How Social Reputation Networks Interact with Competition in Anonymous Online Trading: An Experimental Study, *Working Paper*, 2006

[4] Buchmann, E. and Böhm, K.: FairNet - How to Counter Free Riding in Peer-To-Peer Data Structures, In: *Int. Conf. on Cooperative Information Systems,* 2004.

[5] Callander, S. and Plott, C.: Principles of Network Development and Evolution: An Experimental Study, In: *Journal of Public Economics*, 89 (8), 1469-1495, 2005

[6] Charness, G. et al.: Bargaining and Network Structure: An Experiment, In: *Journal of Economic Theory*, 2007

[7] Fehr, E. and Falk, A.: Wage Rigidity in a Competitive Incomplete Contract Market, In: *Journal of Political Economy,* 107, 106-134, 1999.

[8] Fehr, E. and Gächter, S.: Altruistic punishment in humans, In: *Nature*, 415, 137-140, 2002.

[9] Fehr, E. and Schmidt, K.: A Theory of Fairness, Competition, and Cooperation, In: *The Quarterly Journal of Economics*, 114 (3), 817-868, 1999.

[10] Feldman, M. et al.: Quantifying Disincentives in Peer-to-Peer Networks, In: *Workshop on Economics of Peer-to-Peer Systems*, 2003

[11] Feldman, M. et al.: Robust Incentive Techniques for Peer-to-Peer Networks, In: *ACM Conference on Electronic Commerce*, 102-111, 2004.

[12] Fischbacher, U. et al.: *Fairness and the Power of Competition*, Working Paper, University of Zürich, 2003.

[13] Greiner, B. and Levati, V.: Indirect reciprocity in cyclical networks – An experimental study, In: *Journal of Economic Psychology*, 26 (5), 711-731, 2005.

[14] Grosskopf, B.: Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition, In: *Experimental Economics*, 6 (2), 141-158, 2003.

[15] Hagen, E. and Hammerstein, P.: Game theory and human evolution: A critique of some recent interpretations of experimental games, In: *Theoretical Population Biology*. 69, 339–348, 2006.

[16] Hales, D. and Arteconi, S.: SLACER: A Self-Organizing Protocol for Coordination in P2P Networks, In: *IEEE Intelligent Systems*, 21(2), 29-35, 2006.

[17] Kamvar, S. et al.: Incentives for Combatting Freeriding on P2P Networks, In: *Int. Euro-Par Conf.,* 2003.

[18] Keser, C. and van Winden, F.: Conditional Cooperation and Voluntary Contributions to Public Goods. In: *Scandinavian Journal of Economics*, 102, 23-39, 2000.

[19] Levitt, S. and John A.: What do Laboratory Experiments Measuring Social Preferences tell us about the Real World, In: *Journal of Economic Perspectives*, 2007

[20] Mori, J. et al.: Real-World Oriented Information Sharing Using Social Networks, In: *Conf. on Supporting Group Work*, 81-84, 2005.

[21] Nov, O.: What motivates Wikipedians?, In: *Communications of the ACM*, 50(11), 60-64, 2007.

[22] Ramasamy, L. and Liu, L.: Free Riding: A New Challenge to Peer-to-Peer File Sharing Systems, In: *Int. - Conf. on System Sciences*, 2003.

[23] Ratnasamy, S. et al.: A Scalable Content-Addressable Network, In: *Conf. on Applications, Technologies, Architectures and Protocols for Computer Communications*, 161-172, 2001.

[24] Saroiu, S. et al.: A Measurement Study of Peer-to-Peer File Sharing Systems, In: *Multimedia Computing and Networking,* 2002.

[25] Schosser, S. et al.: Incentives Engineering for Structured P2P systems - a Feasibility Demonstration Using Economic Experiments, In: *ACM Conf. on Electronic Commerce*, 280-289, 2006.

[26] Schosser, S. et al.: Indirect Partner Interaction in P2P Networks - Stimulating Cooperation by Means of Structure, In: *ACM Conf. on Electronic Commerce*, 2007.

[27] Selten, R.: A simple model of imperfect competition where four are few and six are many, In: *Int. Journal of Game Theory* 2, 141-201, 1973.

[28] Selten, R. Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eine Oligopolexperiments, In: *Beiträge zur experimentellen Wirtschaftsforschung*. Heinz Sauermann (ed.), 136-168, 1967.

[29] Seinen, I. and Schram, A. Social Status and Group Norms: Indirect Reciprocity in a Helping Experiment, In: *European Economic Review*, 50, 581-602, 2001.

[30] Smith, V.: Microeconomic Systems as an Experimental Science, In: *The American Economic Review*, 72 (5), 923-955, 1982.

[31] Rowstron, A. & Druschel, P.: Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems, In: *IFIP/ACM Int. Conf. on Distributed Systems Platforms (Middleware)*, 329-350, 2001.

[32] Stoica, I. et al.: Chord: a scalable peer-to-peer lookup protocol for internet applications, In: *IEEE/ACM Transactions on Networking*, 11, 17-32, 2003.

[33] Varian, H.: *Intermediate Microeconomics*, W. W. Norton and Company, 5. Edition, 1999.

[34] Zghaibeh, M. and Anagnostakis, K. G.: On the Impact of P2P Incentive Mechanisms on User Behavior, In: *Joint Workshop on The Economics of Networked Systems and Incentive-Based Computing,* 2007