

Do Humans Identify Efficient Strategies in Structured Peer-to-Peer Systems? (Short Paper)

Stephan Schosser

Universität Karlsruhe (TH)
P.O. Box 6980, 76128 Karlsruhe
+49 (721) 608-3968

schosser@ipd.uka.de

Klemens Böhm

Universität Karlsruhe (TH)
P.O. Box 6980, 76128 Karlsruhe
+49 (721) 608-3968

boehm@ipd.uka.de

Bodo Vogt

Otto-von-Guericke-University Magdeburg
P.O. Box 4120, 39106 Magdeburg
+49 (391) 67-18426

bodo.vogt@ww.uni-magdeburg.de

ABSTRACT

In the last years, distributed coordinator-free systems, e.g., peer-to-peer systems (P2P systems), have attracted much interest among researchers and practitioners. In these systems it is difficult to motivate participants to cooperate. To this end, researchers have proposed various incentive mechanisms. In this paper we are interested in the following question: Do human beings indeed use the strategies that are rational in presence of the incentive mechanism? As humans control the agents in distributed coordinator-free systems, e.g., the peers in peer-to-peer systems, answering this question is essential. We conduct human experiments in the context of structured P2P systems to answer it. This paper shows that humans tend to find it difficult to resort to the strategies expected by the system designer.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems – *human factors, human information processing*; H.3.4 [Information Storage and Retrieval]: Systems and Software – *distributed systems, information networks*

General Terms

Design, Economics, Reliability, Experimentation, Human Factors, Theory, Verification.

Keywords

Economic Experiments, Game Theory, Social Exchange, Mechanism Design, Peer-to-Peer Networks.

1. INTRODUCTION

In the last years, distributed coordinator-free systems have attracted a lot of interest, be it by economists investigating social networks, be it by computer scientists interested in peer-to-peer systems (P2P systems). Such systems rely on the cooperation of their participants. Researchers have designed incentive mechanisms to eliminate uncooperative behavior, i.e., free riding. The designer of such mechanisms sees participants as rational utility maximizers. Using game theory, he identifies the strategies a par-

ticipant can use and calculates the utility of every strategy, given the strategies of the others. Using the utility of the agents in distributed coordinator-free systems, e.g., the peers in peer-to-peer systems, the researcher derives the equilibria. The objective of mechanism design is slightly different: Here, the system designer modifies the utilities of all participants, so that it is rational for them to resort to the strategy favored by him. With real-world applications, participants need to identify and actually use these strategies. Recently Zghaibeh and Anagnostakis [16] have shown that existing incentive mechanisms for P2P file-sharing applications do not lead to the degree of cooperation expected. In the systems investigated, the extent of free riding has been comparable to the one in systems without any mechanisms. Hence, system designers need to keep one question in mind: After having designed an incentive mechanism, do participants indeed resort to the strategies expected? According to behavioral economics, while humans easily find good strategies in situations they are familiar with, they have difficulties to identify strategies which are optimal under rationality [3]. In open distributed systems different individuals and organizations (called participants in the remainder) control the agents. By programming or modifying existing code, they specify the behavior of their agents. Hence, it is important to investigate how human participants perceive distributed coordinator-free systems, and under which circumstances they choose rationally optimal strategies. In this paper, we take a first look at this problem.

For our analysis, we use Content-Addressable Networks (CAN) [10], a prominent variant of structured P2P systems [2], as the basis. In CAN, strategies exist which lead to efficient outcomes (cf. Section 3). Using behavioral experiments, we investigate whether humans identify and use these strategies when controlling their agents. We analyze two different approaches to design strategies: (a) Humans who were never confronted with the problem domain find their strategies by playing (hot strategy design). (b) Humans get used to CAN by implementing different agents. After having gained a thorough understanding of the system, we ask them to realize a strategy successfully cooperating with the strategies of others (cold strategy design). Both hot and cold strategy design is expensive and time-consuming. Further, to make the results comparable, we need to guarantee the same conditions, e.g. comparable monetary rewards, for all participants.

Our results show that both groups, namely cold and hot strategy designers, resort to strategies which drive out free riding, namely

Cite as: Title (Short Paper), Author(s), *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

reciprocity and cut-off strategies [7]. Here, reciprocity means that participants cooperate with others who have cooperated with them in the past. Participants playing cut-off strategies distinguish between cooperative and uncooperative agents by observing their ratio of cooperative actions divided by the number of all actions. If this ratio is above a threshold value, the cut-off value, they are cooperative; otherwise they are not.

In Section 2 we review related work, before we introduce the basics of structured P2P systems in Section 3. We discuss the design of our experiment in Section 4, before we evaluate its results in Section 5. We conclude in Section 6.

2. RELATED WORK

When designing multi-agent systems, two tasks exist [5]: First, the system designer must specify the protocol covering the strategy space of the participants and define how the strategies influence the outcome of the system. Second, the participants need to specify the strategies for all agents. Dash et al. [5] propose computational mechanism design to analyze such systems. Here, a central instance calculates the optimal strategies based on the preferences of the participants and chooses the corresponding strategy profile. Due to the use of a central instance, this approach is problematic in open, distributed systems. Conitzer and Sandholm [4] suggest that the participants define the strategy they deem optimal. A central authority then decides which allocation to choose given the strategies specified by the participants. In this context Papadimitriou [9] proposes the *price of anarchy* as the efficiency difference between a scenario in which the central authority chooses the system-wide optimum and a scenario in which the participants choose the best strategy.

Empirical studies show the impact of the price of anarchy in existing systems. Such analyses of unstructured P2P file sharing systems show that the majority of users free rides in systems lacking an incentive mechanism [1][11]. A study of Zghaibeh and Anagnostakis analyzes the impact of incentive mechanisms in such systems [16]. They observe, that although the mechanisms lead to an increase of participation, the contributions are not as high as expected by the system designer. For structured P2P systems, however, empirical studies are not conductible: Up to now, no such systems are operational in real world.

To cope with the lack of real world applications behavioural experiments can be used. Here, humans mimic agents. Using such experiments, we have shown that inexperienced participants play cut-off strategies [12] and do not resort to feedback, if they can gather information themselves [13]. Both publications do not address the question at what strategies humans arrive in structured peer-to-peer systems. In these analyses, we also do not consider the cold approach.

3. STRUCTURED P2P SYSTEMS

Next, we give an overview over structured P2P systems. We then review game-theoretic literature to predict effective strategies in structured P2P systems.

3.1 Content-Addressable Networks

In our analysis, we focus on Content-Addressable-Networks (CAN) [10], a prominent variant of structured P2P systems. In structured P2P systems (key, value)-pairs are managed. A hash function maps each key to a point in the key space, the query

point. All participants in the system know this hash function. The key space is divided into several zones. Each agent within the system controls one of these zones. I.e., each agent is responsible for the (key, value)-pairs whose key is mapped into its zone. Each agent also knows all agents with adjacent zones, its neighbors, and the zones the control. Queries are request for values given the corresponding key. To answer a query, the query point for the key is calculated. Given the query point lies in the zone of the current agent, the (key, value)-pair is returned. Otherwise, the agent forwards the query to the neighbor closest to the query point. The neighbors repeat this step until the query point is finally reached.

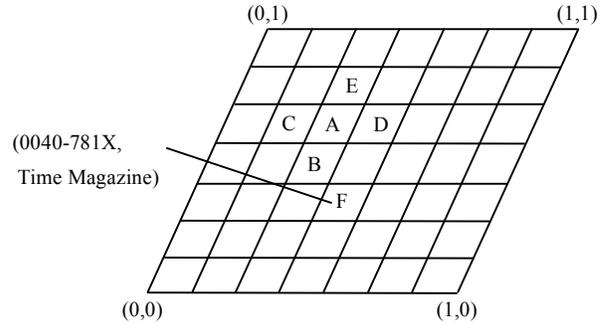


Figure 1: Content-Addressable Network

Example: Figure 1 represents a CAN. A hash function maps the keys to two-dimensional query points. For instance, it maps the key of (0040-781X, “Time Magazine”) to (0.45, 0.3). Each rectangle represents a zone, i.e. Agent F manages the key value pair belonging to “0040-781X”. Be Agent A interested in this (key, value)-pair, it forwards the query to one of its neighbors, Agents B, C, D or E. Therefore, Agent A calculates the query point of “0040-781X”, (0.45, 0.3). It forwards the query to Agent B, the neighbor closest to the query point. Agent B does manage the query result and forwards the query to one of its neighbors. This recurs until the query reaches Agent F. It returns the query result to Agent A. ■

3.2 Game-Theoretic Predictions

Each agent in a CAN plays one strategy. A strategy specifies how the agent forwards, answers and issues queries. If it forwards queries, the strategy also specifies whom to forward the query to. From a game-theoretic perspective agents in structured P2P systems should play reciprocal cut-off strategies: Nowak and Sigmund [8] have shown that reciprocal behavior leads to an equilibrium. Hens and Vogt motivate cut-off strategies [7]. I.e., agents should distinguish between cooperative and uncooperative agents by observing their share of cooperative actions in the past. If it is above a threshold, the other agent is cooperative, otherwise not.

4. EXPERIMENT DESIGN

To analyze under which conditions human participants (who control the behavior of agents) find the strategies predicted by theory and actually use them, we conduct two types of experiments: Hot experiments correspond to hot strategy design. Here, each participant directly controls one agent, which at the same time interacts with the agents of others. Cold experiments correspond to cold strategy design. In cold experiments, each participant implements the strategy of one agent. The implementations of different participants then interact in a test bed.

Each participant controls the strategy of one agent. The experiment environment controls everything else, i.e., it manages the (key, value)-pairs, the neighbors and the query processing (except for the decisions whether and whom to send messages to) for each agent. If a participant decides to forward/issue a query, the experiment environment calculates the distance of all possible recipients to the query point. Further, it assigns the zones to participants and generates all queries. The experiment environment does this randomly and distributes them over the key space using an equal distribution.¹

In the experiments, all zones have the same size, all agents have the same number of contacts and are likely to receive the same number of queries. We conduct all experiments in rounds. Each agent may issue one query per round, and it can answer/forward all messages others have sent to it. If an agent wants to issue/forward a query, the experiment environment calculates the distance of all neighbors to the query point. Based on this distance, it generates a list and shows it to the participant. Thereby, the agent with the lowest distance is on top of the list, the second closest is at the second position, etc. I.e., the position of a neighbor in the list represents the probability that an agent manages the query result. Using the list, an agent can find a trade-off between reputation and distance to the query point of the recipient.

According to the success of its strategy an agent receives points, reflecting the utility of participation in the system. Contributing imposes negative utility, while receiving query results is beneficial: In our experiments, issuing a query costs 2 points, forwarding 1 point and answering 5 points. For receiving a query result an agent receives 20 points. Initially all participants are endowed with 100 points. The intuition behind these points is to reflect the costs and benefits as they would occur in real world applications. Further, the costs of processing one message need to be smaller than the benefits of receiving a query result. Otherwise, participation obviously would not be beneficial. In our experiments, similar to real structured P2P systems, agents only observe for which of their queries they have received query results, and whom they have sent the query initially. They also know how much they have earned in preceding rounds. The experiment environment does not show any other properties of the structured P2P system. E.g., for dropped messages a participant does not know whether the first agent forwarding it has dropped it or another one.

After the experiments, we conducted a strategy game [15]: Here, we asked the participants to describe the strategies they have used in the experiments. More specifically, we gave them abstract descriptions of several system states and asked them how they would have reacted to them. By analyzing several system states in this way, we end up with complete strategies.

We conducted one treatment using both methods (cold and hot experiments) and analyze the results of the experiments.

4.1 Hot Experiment

We conducted the hot experiments with six participants each. According to Selten [15], humans tend to show the same behavior in groups of five participants as they would in groups of more participants. I.e., more participants do not have any *qualitative*

effect on the strategies. With our setup no participant knows the assignment of agents to participants. To prevent communication other than the one within the experiment environment, we physically separated the terminals from each other. In the beginning of the experiment we assigned the participants to random terminals.

After each round after the 20th round, we rolled a six-sided dice. The game continued if the dice showed a number different from one, otherwise it ended. This is an accepted technique to rule out end-game behavior in experimental economics. After the experiments, we paid the participants depending on the points they had received. 100 points have corresponded to € 2.00.

4.2 Cold Experiments

A group of computer-science students participated in the cold experiments. The experiment was a laboratory course lasting one semester. At the beginning of the semester, fourteen students took part in the course. Eleven students remained until the end.

In a first lecture, we gave an introduction to of structured P2P systems. The programming part of the course consisted of 5 iterations. Each iteration lasted two weeks. At the beginning of an iteration, we asked the participants to implement a strategy given the understanding they had gained so far. At the end of an iteration, we showed the participants the success of each implementation and revealed to them the strategies of the other participants. We calculated the ‘success’ of a strategy by running a simulation of a structured P2P system with all implementations developed by the students. The experiment environment randomly assigned the implementations to agents. The success of a strategy was the number of points the corresponding agent had earned. To support the students, we gave them the simulation environment where they could plug their implementations for testing. The participants spent 6.67 hours per week working on their implementations on average, according to a questionnaire we handed out at the end of the semester.

In all iterations, the participants kept refining their strategies. We used the first four iterations to make the participants familiar with the problem domain. The fifth iteration was the basis for our analysis: It is this iteration when we paid the participants depending on the success of their strategy. We limited both the analysis and the payment to the last iteration to keep the costs of the whole experiment down.

Similarly to the hot experiments, the cold experiments lasted for at least 1000 rounds. Afterwards, we rolled a six-sided dice. If it showed one, the experiment ended, otherwise it continued. At the end of the course, we paid the participants depending on the success of their strategies in the fifth iteration. For 1,000 points earned, the received a payoff of € 1.00. This payoff scheme has the characteristic that the participants in hot experiments would earn approximately the same as in cold experiments if all participants used the same strategy.

5. EXPERIMENTAL RESULTS

As described in Subsection 3.2, one could expect reciprocal cut-off strategies to reach the efficient equilibrium, according to game theory. We analyze the results of the strategy game of both experiments to find out whether the participants play cut-off strategies. To this end, we came up with three categories of strategies: cut-off strategies depending on the past success frequency of own

¹ Find the software we use, the description of the experiments together with screenshots, log files and implementations at <http://www.ipd.uni-karlsruhe.de/~schosser/aamas08>

queries (Category A), cut-off strategies with modifications (Category B), strategies which are not cut-off strategies (Category C).

Table 1: Strategies Observed in Hot vs. Cold Experiment

Strategy	Cat.	# pers. (hot)	# pers. (cold)
<i>Cut-off strategy depending on past success frequency (CSPSF) only</i>	A	35	5
<i>CSPSF plus end phase or start phase</i>	B	9	1
<i>CSPSF plus limit for answering queries of others per round</i>	B	0	1
<i>Cut-Off strategy depending on the absolute number of own queries not answered</i>	C	0	1
<i>Unconditional cooperation</i>	C	11	0
<i>Unconditional cooperation plus condition that account is high</i>	C	3	0
<i>Unconditional cooperation plus limit for number of answers</i>	C	1	
<i>Different free-riding strategies</i>	C	0	2
<i>Different types of strategies</i>	C	1	1
Sum		60	11

The results of the strategy game (see Table 1) show that a majority of the participants plays cut-off strategies. A binomial test confirms this on significance level of 1% for the hot experiments and 15% for the cold experiments. Another binomial test confirms on significance level of 1% that participants in both experiments behave reciprocally: 44 of 60 (eight of eleven) participants used such strategies in the hot (cold) experiments. Nevertheless, 16 of 60 (three of eleven) participants played strategies different from the expected reciprocal cut-off strategies. This shows that not all behavior is in line with the behavior expected according to theory.

Next, we analyze how cooperative the participants have been in the experiments. We simulated systems of six implementations students had programmed as part of the cold experiments for twenty rounds. In these simulations, we used every possible combination of implementations, i.e., we created every combination of six strategies given the eleven implementations we had received. We repeated the simulations with each combination of strategies five times by assigning different positions to the different strategies. In this way, we could compare hot and cold experiments. In the hot experiments, the average degree of cooperation was 82% (51% in the cold experiments). Behavioral experiments using other mechanism, such as punishment (64%) [6] or feedback (69%) [14], guarantee similar degrees of cooperation. In systems without any incentive mechanism as in the cold experiments, lower degrees of cooperation occur (40% - 50%) [14]. This result shows that participants in cold experiments find it difficult to play the efficient equilibrium.

6. CONCLUSIONS

Incentive mechanisms in distributed coordinator-free systems have received much interest in the recent past. Efficiency in such systems depends on the strategies of the participants. We have analyzed this using Content-Addressable Networks, a prominent variant of structured P2P systems. According to game-theoretic literature, participants in such systems should use reciprocal cut-off behavior. In both hot and cold experiments most participants

recognize the usefulness of reciprocity and cut-off behavior. Nevertheless, a large fraction of participants uses different strategies, such as unconditional cooperation or free riding, in both settings. This results in systems which are not efficient.

7. REFERENCES

- [1] Adar, E. and Huberman, B. 2000. Free Riding on Gnutella. Technical Report. Xerox PARC.
- [2] Androutsellis-Theotokis, S. and Spinellis, D. 2004. A survey of peer-to-peer content distribution technologies. *ACM Comput. Surv.* 36, 4, 335-371.
- [3] Cosmides, L. and Tooby, J. 1994. Better than Rational: Evolutionary Psychology and the Invisible Hand. *Am. Econ. Rev.* 84, 2.
- [4] Conitzer V. and Sandholm T. 2004. Self-interested Automated Mechanism Design and Implications for Optimal Combinatorial Auctions. In *Proceedings of the ACM Conference on Electronic Commerce*.
- [5] Dash, R. et al. 2003. Computational Mechanism Design: A Call to Arms. *IEEE Intell. Syst.*
- [6] Fehr, E. and Gächter, S. 2002. Altruistic punishment in humans. *Nature.* 415, 137-140.
- [7] Hens, T. and Vogt, B. 2002. Money and Reciprocity. Working Paper, University of Zürich, in revision.
- [8] Nowak, M. and Sigmund, K. 1998. Evolution of Indirect Reciprocity by Image Scoring. *Nature.* 393, 573-577.
- [9] Papadimitriou, C. 2001. Algorithms, Games, and the Internet. In *Proceedings of the ACM Symposium on Theory of Computing*.
- [10] Ratnasamy, S. et al. 2001. A Scalable Content-Addressable Network. In *Proceedings of the Conference on Applications, Technologies, Architectures and Protocols for Computer Communications.* 161-172.
- [11] Saroiu, S. et al. 2002. A Measurement Study of Peer-to-Peer File Sharing Systems. In *Proceedings of the Multimedia Computing and Networking*.
- [12] Schosser, S. et al. 2006. Incentives Engineering for Structured P2P Systems - a Feasibility Demonstration Using Economic Experiments. In *Proceedings of the ACM Conference on Electronic Commerce*.
- [13] Schosser, S. et al. 2007. Indirect Partner Interaction in P2P Networks - Stimulating Cooperation by Means of Structure. In *Proceedings of the ACM Conference on Electronic Commerce*.
- [14] Seinen, I. and Schram, A. 2006. Social Status and Group Norms: Indirect Reciprocity in a Helping Experiment. *Eur. Econ. Rev.* 50, 581-602.
- [15] Selten, R. 1967 Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eine Oligopol-experiments. *Beiträge zur experimentellen Wirtschaftsforschung.* Heinz Sauermann (ed.). 136-168.
- [16] Zghaibeh, M. and Anagnostakis, K. 2007. On the Impact of P2P Incentive Mechanisms on User Behavior, In *Proceedings of the Joint Workshop on The Economics of Networked Systems and Incentive-Based Computing*.