Game theory in peer-to-peer networks

Teemu Typpi
Helsinki University of Technology
ttyppi@cs.hut.fi

Abstract

Game theory has been used to model the behaviour of selfish participants in P2P networks in which certain incentives to share resources are present, and also to model the resulting efficiency of the network. This enables the design of incentives which cause the network to operate in an optimal way. This paper surveys the state of research into how game theory may be applied to peer-to-peer (P2P) networks. We give a brief introduction to game theory, describe the main problem in P2P networks, i.e. free-riding, and present the commonly proposed incentives for sharing resources in a P2P network. Finally, we present some relevant studies which utilize game theory to find solutions to the free-riding problem.

KEYWORDS: Game theory, free-riding, peer-to-peer, incentives

1 Introduction

A peer-to-peer (P2P) network is a system where the participants are directly connected to each other and can act as both consumers and providers of resources in the network. This is in contrast to the more traditional model where a single server or multiple servers provide all the resources and the participants of the system are clients which only consume.

The greatest advantages of P2P networks over centralized networks are scalability and robustness. They are scalable because when a new participant joins a P2P network, not only the load on the system increases but at the same time the resources of the system increase. The network is also less vulnerable to attacks or failures because each node is only a small part of the system, whereas in a centralized network an attack against one of the main servers can severely reduce the performance of the whole network.

A major problem for the P2P networks are free-riders, i.e. users who do not contribute anything, or contribute a negligible amount compared to the resources they consume. Studies have shown that in a system where sharing is not rewarded and lack of sharing is not punished in any way, the majority of users choose to free-ride. For example, [15] discovered that 85 percent of users in the Gnutella network share no files with others.

Most applications nowadays which use the P2P paradigm are file sharing applications, such as BitTorrent and Gnutella [2, 1]. Other popular usages for it are grid computing, distributed backup, software updates and instant messaging software such as Skype [3].

	P2 Action 1	P2 Action 2
P1 Action 1	w1,w2	x1,x2
P1 Action 2	y1,y2	z1,z2

Table 1: Example representation of a strategic game with two players and two possible actions for both

Originally, it was thought that the users in P2P networks would contribute resources unselfishly for the common good. Because nowadays this assumption is not realistic, there is a need for methods to encourage or force cooperation. A P2P network can be considered as a system consisting of selfish participants whose goal is to gain as many resources as possible from the network compared to the cost for themselves. Game theory has been used as a way to design incentives for sharing in P2P networks and in this paper we will outline some of the research in this area.

The rest of this paper is organized as follows. Sec. 2 gives a short introduction to game theory. In Sec. 3 we describe the problem of free-riding in more depth. Sec. 4 presents proposed incentives and their issues. In Sec. 5 we introduce and summarize some relevant studies. In Sec. 6 we take a look at some ways to cheat in a BitTorrent network. Finally, in Sec. 7 we conclude the paper.

2 Introduction to Game Theory

This section is a brief introduction to game theory and is based on [20]. The purpose of this section is only to explain lightly what game theory is about, for those new to the subject.

2.1 Strategic Games

A strategic game models interactive decision making between the participants. The components of a strategic game are a set of players, a set of actions for each player, and for each player the set of actions he would prefer the other players to take when he takes a certain action. The outcome of the game is defined after all actions are taken. It is often clearer to consider the payoffs for each player after the actions are taken instead of the preferred actions of others. A representation of a two player strategic game is shown in Table 1.

The rows represent possible actions for player 1, and columns represent the possible actions for player 2. In each box there are the payoffs for both players in the corresponding outcome of the game.

	Don't confess	Confess
Don't confess	3,3	0,4
Confess	4,0	1,1

Table 2: Payoff matrix for the Prisoner's dilemma

2.2 Nash Equilibrium

An important concept in game theory is the concept of Nash equilibrium. A Nash equilibrium in a game is a state in which all players predict the actions of the other players correctly, and none of the players can benefit by switching his action to some other. In this situation the set of chosen actions and payoffs of the outcome constitute the Nash equilibrium. The Nash equilibrium, however, is not necessarily the best outcome collectively, i.e. when counting the payoffs of all players together.

2.3 Prisoner's Dilemma

Prisoner's dilemma is a classical example of a simple strategic game. It models a situation where two suspects are under arrest and are interrogated separately. If neither of them confesses, both get only a short prison sentence of one year. If both confess, both will get a sentence of three years. If only one confesses, he will get free and his testimonial will be used against the other who will get a long sentence of four years. The payoff matrix is presented in Table 2.

Assuming that neither player wants to spend time in prison, the best choice for a player is to confess, since it gives the best payoff for that player, no matter what the other one decides. Even though both players would be better off if neither confessed, the only Nash equilibrium for the game is that both confess.

If the game would be repeated several times with the same players, the outcome in which the players cooperate (do not confess) in each iteration can be stable. The reason for this is that the players may believe that once they defect, the other player will punish them by starting to defect too, which results in a worse payoff in the long run. We will later find out that the interactions between users in P2P systems are better modelled as repeated games, and that this threat of getting punished later on may work as an important incentive for contributing to the system.

3 The problem of Free-riding

A free-rider in a P2P system is a user who only tries to use resources and does not offer them himself. If the contribution cost is high and there are no incentives for sharing, the system is in danger to collapse because many users would choose to free-ride.

In [10] the authors outline several important questions which have to be taken into account when designing incentive systems to alleviate free-riding.

Firstly, it is important to find out which factors determine how widespread free-riding becomes in a system. An example of such factor is the contribution cost. Secondly, it should be examined how free-riding affects the system's performance. It would be of little use if the cost of eliminating free-riding brought the system's performance close to zero.

When designing incentive schemes, the characteristics of the P2P system have to be considered. Such characteristics are for example how dynamically users join and leave, whether it is easy to gain cheap identities, traceability and visibility of actions, and prevalence of collusive behaviour in the system.

In addition, the available mechanisms to induce cooperative behaviour must be thought of. Some of these will be introduced later in this paper.

Another question is also the effect of available cheap identities on user behaviour in the system. It has to be taken into account that users might rejoin the network under a new identity if it helps to improve their reputation cheaply enough.

Finally, it should be considered how rational the users in the P2P system actually are and how possible irrationalities would affect the system.

4 Proposed Incentives

This section will describe incentives for sharing resources as proposed in [10]. On the basic level, there are three reasons why users would share resources in a P2P network. These reasons are (1) altruism, (2) the knowledge that they get something in return and (3) the fear of being punished. Next we take a look at the proposed incentives for sharing.

4.1 Inherent Altruism

Empirical studies have shown that it is not accurate to assume that all users behave completely selfishly [8]. [6] introduces a model of giving based on the assumption that people gain utility, a "warm glow", from the act of giving to other people. This generosity has enabled P2P systems without incentive mechanisms to function, especially when the cost of sharing resources is reasonably low. The authors of [12] create a model of user behaviour in P2P systems which takes into account the inherent generosity of users. We will introduce their study later in this paper.

4.2 Monetary Schemes

The idea in monetary schemes is that participants pay money when they consume resources and get paid when they share resources to others. [13] proposes several payment methods and analyzes their effects on user behaviour in P2P networks by using a formal game theoretic model of the system. The incentives in these systems are obviously strong, but the implementation of such systems in practice is difficult and requires infrastructure for accounting and micropayments.

4.3 Reciprocity-based Schemes

There are two kinds of reciprocity-based schemes, namely direct reciprocity and indirect reciprocity. In direct reciprocity schemes, the decision a user makes about how to serve another user is based on the quality of the service he received from that user previously. In indirect reciprocity schemes, also called reputation based schemes, the overall generosity of the user is taken into account when deciding the quality of service he deserves to obtain.

4.3.1 Direct Reciprocity

In direct reciprocity schemes when a user serves another user well, that user will serve well in return. These kind of schemes work well in P2P systems where session times are long. For example, in BitTorrent networks an incentive mechanism based on a variation of tit-for-tat [9] is used, although its effectiveness in preventing free-riding has been shown to be questionable in [16]. EMule [5] uses a credit system, in which the more user A uploads to user B, the less user A has to queue when requesting files from user B. The credits are not global but are always between two specific users, so this is a direct reciprocity scheme.

Direct reciprotity schemes have also been applied in P2P multicast streaming applications. [19] suggests that the multicast tree is periodically rebuilt to enable the threat of a free-riding peer getting punished by a node it previously refused to serve.

4.3.2 Indirect Reciprocity

In indirect reciprocity schemes the reputation of users is tracked and those with good reputation get good service. The general idea is that the reputation of a node increases when it serves others and decreases when it does not. The difference to direct reciprocity is that a user does not need to serve a certain user to get good service from him, as long as the overall reputation is high. This is an advantage for example if that user is not interested in anything that the other one could offer. Indirect reciprocity schemes allow for greater scalability than direct reciprocity schemes in P2P systems where populations are large and users leave and join frequently, because it is unlikely that two peers will repeatedly interact and thus there will not be opportunities to retaliate against uncooperative peers [11].

5 Game Theoretic Studies on P2P systems

This section summarizes some relevant studies which use game theory to analyze P2P systems.

5.1 Free-Riding and Whitewashing in Peer-to-Peer Systems [12]

In [12] the authors develop a model to examine free-riding in P2P systems. In their model, a user's inherent generosity is taken into account with a separate parameter. A user in the network will share if his generosity is high enough compared to the current contribution cost in the system, which is inversely propotional to the amount of contributors. The paper studies the effects of different mechanisms which punish free-riders. The authors also examine the problem of white-washers, i.e. users who rejoin the network under new identities to avoid getting penalized from bad reputation.

First the authors study how high the generosity level of the society has to be for the system to be stable and what percentage of free-riders will remain.

The authors then study how a penalty imposed on users who get identified as free-riders based on their behaviour affects the prevalence of free-riding and the performance of the system. The penalty can be exclusion from the system with some probability, or a reduction of the service free-riders get. The authors find out that if the penalty is set high enough, or if it is possible to recognize free-riders with a high enough probability, optimal system performance may be reached.

The authors consider the issue of white-washing when the cost of new identities is infinite, and when new identities are free. In the case of free identities, they suggest that to discourage white-washing, a penalty could be imposed on all newcomers. This reduces system performance since this penalty is unavoidable and also new contributors suffer it in addition to white-washers. However, the authors show that the loss in performance is significant only if the user turnover rate in the system is high. If the cost of a new identity is infinite, identities are permanent, and thus white-washing is impossible and new users do not need to be punished at all.

5.2 Game Theory as a Tool to Strategize As Well As Predict Node's Behavior In Peer-to-Peer Networks [14]

In [14] the authors propose a basic reputation based incentive scheme for a P2P system and examine how successfully it prevents freeriding and how well the system performs.

They assume that all users are rational and strategic and want to maximize their utility for minimal cost. They also assume that users do not gain utility by being altruistic. In their solution, when a user requests for service, the probability for obtaining it is directly propotional to the user's reputation. Reputation increases when the user accepts a service request and decreases when it denies one. The authors use different weights for how much past and current actions affect reputation.

The authors modeled the system as an infinitely repeated game. Each game corresponds to a time period during which all users request service once, and either choose to serve or not serve others. The authors then calculate the Nash equilibrium strategies for the users.

One important finding is that the Nash equilibrium strategy for the users is to serve with a probability that is greater than zero in each time period. This means that free-riding will not be a problem, because it simply is a bad strategy. Another finding is that if the utility that can be gained from the network is much greater than contribution costs, optimal strategy can never be to share with a probability higher than 50 percent. It is noted that this leads to a reduced system efficiency, since more than 50 percent of all service requests are denied. The authors also show that the more weight is given to past actions, the better the system performs since users have to regularly provide services to keep their reputation high enough.

5.3 A Game Theoretic Framework for Incentives in P2P Systems [7]

In [7] the authors study the interaction of strategic and rational peers with the help of game theory. They propose an incentive scheme in which the peers receive differential service based on their behaviour to reduce freeriding and improve system's performance. They start by studying the case of a homogenous system, i.e. a system where all users gain equal benefit from everyone else. They show that there exists two Nash equilibria contribution levels for the system and that the one resulting in the better overall performance is realized.

Next the authors study the Nash equilibrium in a heterogenous system through simulation. They show that these systems also converge to the desirable Nash equilibrium and that the result does not depend on the number of peers in the system. However, the bigger the system is, the more robust it is against loss of performance caused by peers leaving the system. They also note that the effect of non-cooperative peers on the system is to bias the equilibrium towards their contribution value, but that the effect is generally insignificant.

Finally, the authors suggest how current P2P networks could be modified to implement the incentive scheme proposed. They suggest that each peer accepts requests from other peers with a probability which is a function of the requesting peer's contribution level. A short history of requests should be kept in memory to enable checking that no peer tries to circumvent the system by sending several service requests quickly in a row. The probability function could be a part of the systemt's architecture and thus same for all users and not modificable by them. The contribution level of a peer could be calculated from the peer's uptime and its shared disk space. The contribution level could then be attached in file request messages as an additional header. To enable new users to start using the system properly right away, they can be given a default initial contribution level for some time. To prevent peers from reporting their contribution levels falsely, the authors suggest implementing a neighbour audit scheme in which peers monitor the disk spaces and uptimes of their neighbour peers.

5.4 Clustering and Sharing Incentives in BitTorrent Systems [17]

In [17] the authors perform an experimental investigation of the peer selection strategy in the BitTorrent protocol. BitTorrent uses a choking algorithm for peer selection, which mainly prefers users with high upload rates, thus giving an incentive to share. In general the algorithm performs well, but some studies suggest that it does not solve the free-riding problem perfectly [16, 18].

The authors ran simulations with a single initial seed and 40 leechers. A seed in a BitTorrent network means a user who has the complete file and is uploading it to others. A leecher is a user who does not have the full file yet, but might be also contributing already by uploading those parts of the file which he has downloaded so far. The simulations were run with leechers separated into two classes by upload limits, then three, and finally with a uniform distribution of upload limits for the leechers. Separate runs were also made with the original seeder being well-provisioned (high upload speed compared to the leechers) or underprovisioned (lower upload speed).

The results considering the existence of incentives for sharing were that when the initial seed was well-provisioned, there were effective sharing incentives. Those leechers who uploaded most also finished the download fastest. Leechers with low upload rates took significantly longer to finish downloads. However, the sharing was unfair considering the amount of data shared. Fast uploaders did not gain more data even though they shared it more to others, they only gained it faster.

In the simulations with an underprovisioned initial seed, all leechers finished the download around the same time and thus it appears that the algorithm fails to provide effective incentives for sharing in this case.

6 Cheating in BitTorrent

Although the BitTorrent protocol is designed to encourage sharing, it is possible to use modified client software, such as the one described in [18], to download entire files without contributing at all. The authors developed a client, named BitThief, which does not announce the pieces it has downloaded and thus continuosuly pretends to be a newly joined peer. This way it never has to upload anything, but it can still download from seeders and from leechers which periodically accept to transfer data to it without reciprocation (optimistic unchoking). Unlike regular BitTorrent clients, BitThief does not try to get the rarest pieces first, but will accept any random piece so that all unchoke periods are taken advantage of. The BitThief client also tries to gain a long list of peer addresses by announcing itself to the tracker more frequently in the beginning of the downloading than normally. This improves download rates simply because there are more potential peers which will unchoke the client.

On the Internet there are file sharing community sites which offer lists of high quality torrents. To keep the quality high, these sites usually keep the number of their users limited and require invited registration. They also require the users to have high upload to download ratios and ban users which do not contribute enough. These ratios are calculated from the upload and download data in the announcements the peers give to the community's private torrent tracker. However, because this data is not verified, it has led to the development of software which manipulates the reported upload data to seem higher. An example of such software is GreedyTorrent [4], which works as a proxy between a BitTorrent client and a torrent tracker. The software offers the user the opportunity to modify the reported upload rate to be a multiple of either the real upload rate, or the real download rate. Thus, if a community site requires the upload to download ratio to be 1:1, the user may simply set the reported upload rate to be the real download rate multiplied by one and avoid getting banned, no matter how negligible the user's contribution is in real.

7 Conclusion

In this paper we studied the state of reseach into how game theory can be applied to P2P systems. The research we focused on mostly considered P2P systems with low contribution costs, such as file sharing ones. We found out that game theory is used to predict what kind of strategies the participants will use to gain maximum utility from the networks. In particular, it is used to design effective incentives to prevent free-riding, which is the biggest challenge in such systems. At the same time, the performance of the network can be evaluated based on the predicted behaviour of the users. This is important because even if the incentives would eliminate free-riding, it would be of little use if the network's functioning would become otherwise ineffective or impossible altogether.

We found that most research concentrates on incentive models which are based on reciprocation schemes, in which some form of a punishment is imposed on those peers who are detected contributing too little. However, research on the more complex monetary schemes has been done also.

Although currently the incentive system in BitTorrent is not strong enough to make free-riding impossible, these networks still thrive because of altruism, and also because the cost of uploading data to others is low, many users probably will not even bother looking for ways to cheat. However, to ensure optimal efficiency of torrent networks, and especially of P2P networks with higher contribution costs, it is necessary to continue searching for robust game theoretic solutions which offer no opportunities for free-riding or any kind of cheating.

References

[1] Gnutella. http://gtk-gnutella.sourceforge.net.

[2] Bittorrent. http://www.bittorrent.com.

[3] Skype. http://www.skype.com.

[4] Greedytorrent. http://www.greedytorrent.com.

[5] Emule. http://www.emule-project.net.

[6] J. Andreoni. Giving with impure altruism: Applications to charity and ricardian equivalence. *The Journal of Political Economy*, 97(6):1447–1458, 1989.

[7] C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in p2p systems. In *P2P '03: Proceedings of the 3rd International Conference on Peer-to-Peer Computing*, page 48, Washington, DC, USA, 2003. IEEE Computer Society.

[8] C. F. Camerer. *Behavioral Game Theory*. Princeton University Press, 2003.

[9] B. Cohen. Incentives build robustness in bittorrent. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[10] M. Feldman and J. Chuang. Overcoming free-riding behavior in peer-to-peer systems. *SIGecom Exch.*, 5(4):41–50, 2005.

[11] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 102–111, New York, NY, USA, 2004. ACM.

[12] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. *Selected Areas in Communications, IEEE Journal on*, 24(5):1010–1019, May 2006.

[13] P. Golle, K. Leyton-Brown, and I. Mironov. Incentives for sharing in peer-to-peer networks. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 264–267, New York, NY, USA, 2001. ACM.

[14] R. Gupta and A. K. Somani. Game theory as a tool to strategize as well as predict nodes behavior in peer-to-peer networks. *Parallel and Distributed Systems, International Conference on*, 1:244–249, 2005.

[15] D. Hughes, G. Coulson, and J. Walkerdine. Free riding on gnutella revisited: the bell tolls? *Distributed Systems Online, IEEE*, 6(6):–, June 2005.

[16] S. Jun and M. Ahamad. Incentives in bittorrent induce free riding. In *P2PECON '05: Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 116–121, New York, NY, USA, 2005. ACM.

[17] A. Legout, N. Liogkas, E. Kohler, and L. Zhang. Clustering and sharing incentives in bittorrent systems. In *SIGMETRICS '07: Proceedings of the 2007 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 301–312, New York, NY, USA, 2007. ACM.

[18] T. Locher, P. Moor, S. Schmid, and R. Wattenhofer. Free riding in bittorrent is cheap. In *Proc. of HotNets-V*, pages 85–90, November 2006.

[19] T.-W. J. Ngan, D. Wallach, and P. Druschel. Incentives-compatible peer-to-peer multicast. In *Workshop on Economics of Peer-to-Peer Systems*, 2004.

[20] M. J. Osborne. *A course in game theory*. MIT Press, 4th edition, 1994.