

Incentives for Sharing in Peer-to-Peer Networks

Philippe Golle Kevin Leyton-Brown Ilya Mironov

Computer Science Department, Stanford University
{pgolle,kevinlb,mironov}@cs.stanford.edu

ABSTRACT

We consider the *free-rider* problem that arises in peer-to-peer file sharing networks such as Napster: the problem that individual users are provided with no incentive for adding value to the network. We examine the design implications of the assumption that users will selfishly act to maximize their own rewards, by constructing a formal game theoretic model of the system and analyzing equilibria of user strategies under several novel payment mechanisms. We support and extend upon our theoretical predictions with experimental results from a multi-agent reinforcement learning model.

1. INTRODUCTION

Peer-to-peer (P2P) file-sharing systems combine sophisticated searching techniques with decentralized file storage to allow users to download files directly from one another. The first mainstream P2P system, Napster, attracted a great deal of public attention for the P2P paradigm as well as tens of millions of users for itself. Napster specialized in helping its users to trade music; P2P networks also allow users to exchange other kinds of digital content.

The work of serving files in virtually all current P2P systems is performed for free by its users. Since users do not benefit from serving files to others, many users decline to perform this altruistic act. In fact, two recent studies of the Gnutella network have found that an overwhelming proportion of its users contribute nothing to the system [1, 8]. The phenomenon of selfish individuals who opt out of a voluntary contribution to a group's common welfare has been widely studied, and is known as the *free-rider* problem [5, 9]. The communal sharing of information goods in "discretionary databases" and the resulting free-rider problem has also been studied before the advent of P2P systems [10]. This problem is not simply theoretical. Some P2P systems plan to charge users for access in the near future. However, a system run for profit may not receive the level of altruistic 'donations' that power a free community. There is therefore both a need and an opportunity to improve such P2P file-sharing systems by using an improved incentive scheme to increase the proportion of users that share files, making a

greater variety of files available. This would increase the system's value to its users, and hence make it more competitive with other commercial P2P systems.

In the following section, we introduce our formal game theoretic model. Section 3 presents the Napster system, which we use as a motivating example throughout this paper. In sections 4 and 5, we propose two classes of novel payment mechanisms, analyzing user strategies and the resulting equilibria. Finally in section 6, we use a multi-agent reinforcement learning model to validate our analytical results and to explore further properties of our mechanisms.

2. PROBLEM DEFINITION

We now turn to a more formal, game theoretic characterization of the problem. (Readers unfamiliar with game theoretic analysis may consult [3, 7].) First, we describe the game that we use to model the file sharing scenario. In our model, usage of the system is divided into time periods of equal duration. For example, time periods might represent one month. There are n agents who participate in this system; we denote them as a_1, \dots, a_n . Each agent has two independent actions available in each time period:

1. **Sharing:** Agents select what proportion of files to share. In our model, sharing takes three levels: σ_0 (no sharing), σ_1 (moderate sharing) or σ_2 (heavy sharing).
2. **Downloading:** Each agent must also determine how much to download from the network in each period. We model downloads with agents choosing between three levels: δ_0 (no downloads), δ_1 (moderate downloads) or δ_2 (heavy downloads).

An agent a_i 's strategy in time period t is denoted $S(i, t) = (\sigma, \delta)$, or $S(i)$ when the time period is unambiguous.

2.1 Agent Utility

Agents' utility functions describe their preferences for different outcomes. The following factors concern agents:

- **Amount to Download (AD):** Agents get happier the more they download.
- **Network Variety (NV):** Agents prefer to have more options from which to select their downloads.
- **Disk Space Used (DS):** There is a cost to agents associated with allocating disk space to files to be shared.
- **Bandwidth Used (BW):** Similarly, there is a cost to agents associated with uploading files to the network.
- **Altruism (AL):** Some agents derive utility from the satisfaction of contributing to the network.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'01, October 14-17, 2001, Tampa, Florida, USA.
Copyright 2001 ACM 1-58113-387-1/01/0010 ...\$5.00.

- **Financial Transfer (FT):** Agents may end up paying money for their usage of the network, or conversely they may end up getting paid.

We assume that agents have quasilinear utility functions. Furthermore, we assume that agents are risk neutral. We can thus write the equation for agent a_i 's utility function as $U_i = [f_i^{AD}(AD) + f_i^{NV}(NV) + f_i^{AL}(AL)] - [f_i^{DS}(DS) + f_i^{BW}(BW)] - FT$. Each f function is concerned with a particular variable (e.g., bandwidth used) and an agent; it describes that agent's preference for different values of the variable, in money. There is no f function for the variable FT because this variable represents an amount of money that is transferred to or from the agent. Without restricting ourselves to particular f functions, we can make several observations that justify the signs of the terms above. First, f^{AD} , f^{NV} and f^{AL} must be monotonically increasing, with minimum value 0, as these variables only ever contribute positive utility. DS and BW only contribute negative utility, explaining the subtraction of f^{DS} and f^{BW} above. Finally, we assume that neither f^{DS} nor f^{BW} is superlinear.

We say that two agents a_i and a_j have the same *type* if they have the same utility function; i.e., if $f_i = f_j$ for all five f functions. To simplify our game theoretic analysis in the first part of this paper we often make the assumption that all agents have the same type. In section 6 we approach the file sharing problem experimentally; this approach allows us to discuss the convergence of agent strategies under a wide variety of different agent types.

2.2 Equilibria

As is central to any game theoretic model, we assume that agents are economically rational, and that they act to maximize their expected utility, given their beliefs about the actions that other agents will take and their knowledge about the way that their payoffs are calculated. We denote the joint strategies of all agents in time period t as $\Sigma(t) = \{S(1, t) \dots S(n, t)\}$, or simply as Σ when the time period is unambiguous. Following the usual definition, we say that Σ is a *weak Nash equilibrium* when no agent can gain by changing his strategy, given that all other agents' strategies are fixed. Similarly, Σ is a *pure Nash equilibrium* when every agent would be strictly worse off if he were to change his strategy, given that all other agents' strategies are fixed. Finally, an agent has a *dominant strategy* if his best action does not depend on the action of any other agent.

2.3 Assumptions and Observations

In our analysis, we restrict ourselves to file sharing systems that make use of centralized servers. These servers maintain a database of the files currently available on the network and connect download requests with available clients.

We assume that the servers are able to determine the identity of files provided by users, which may be needed both to pay royalties to the appropriate copyright holder and to detect users who make false claims about the files they share. File identification may be achieved by a cryptographic watermarking scheme (see, e.g., www.sdmi.com); alternately, users who spoof files could be penalized.

One likely payment model for peer to peer systems is some kind of flat rate membership fee per time period. We do not explicitly consider this option anywhere in the discussion that follows, as it has no impact on the equilibria that arise

from any mechanism (although it can affect agents' decisions about participation). All the mechanisms discussed here are *compatible* with the addition of flat rate pricing. The fact that flat fees are unrelated to agents' behavior implies that they still give rise to a free rider problem.

3. THE NAPSTER SYSTEM

In this section we analyze the Napster system that operated from May 1999 through July 2001, since it is probably the best-known peer-to-peer application. Napster is one of the simplest mechanisms that can be represented by our model: regardless of the actions of agents, Napster imposes no financial transfers. Using the model described in section 2, we start with an equilibrium analysis that disregards the 'altruism' component of agents' utility functions; we then go on to consider altruism.

Unsurprisingly, $\Sigma = \{(\sigma_0, \delta_2), \dots, (\sigma_0, \delta_2)\}$ is an equilibrium. As all agents have the same type, it is enough to analyze the choice made by a single agent. Assume that agents other than a_i follow the strategy $S = (\sigma_0, \delta_2)$, and consider agent a_i 's best response. Since a_i is not altruistic, his utility is strictly decreased by sharing files; he will thus choose the action σ_0 which leaves his utility unchanged. Downloading will usually increase a_i 's utility; however, since no other agent is sharing we have $NV = 0$ and so his utility is zero regardless of how much he intends to download. The action δ_2 is therefore a best response, and Σ as given above is an equilibrium. Indeed, we can see that the strategy $S = (\sigma_0, \delta_2)$ is dominant. If all other agents choose σ_0 then S yields the same (maximal) payoff as (σ_0, δ_0) and (σ_0, δ_1) ; if any other agent does share then S yields strictly higher revenue than any other strategy. Because Σ is an equilibrium in dominant strategies, it is the only equilibrium.

We have identified a unique equilibrium in which nothing gets shared and there is nothing to download. Yet songs were plentiful and actively traded on Napster. We identify two incentives that could account for users' willingness to contribute. First, Napster offered its service free of charge and went to great lengths to foster a sense of community among its users. This may have been sufficient to encourage users to altruistically contribute resources that cost them very little. Second, Napster offered a (modest) disincentive for non-contribution: by default, the Napster client shared all songs that an agent downloaded. We represent both of these incentives through the variable (AL).

In the analysis of this situation, we consider two types of agents. First, altruistic agents are those whose reward for altruistic behavior (AL) exceeds its cost in terms of disk space (DS) and expected bandwidth usage (BW). We assume that f functions for these agents are such that they would prefer the action σ_2 to either the action σ_1 or σ_0 regardless of the value of BW . These agents still gain utility from downloads: following an argument similar to the one given above, (σ_2, δ_2) is a dominant strategy for altruistic agents. The second type of agents are those for whom the cost of altruistic behavior exceeds its benefit. These agents are essentially the same as those described in the previous section: although they may receive some payment for altruistic behavior, it will be insufficient to alter their behavior. They thus have the dominant strategy given above: (σ_0, δ_2) .

This analysis is arguably a description of the state of affairs on the Napster system. Some proportion of agents were sufficiently altruistic to share files and did so; other

agents were not altruistic and shared nothing. Regardless of their level of altruism, agents were unrestrained in their downloads. We can now see that Napster did experience a free rider problem: regardless of the contributions of others, selfish agents had no incentive to share.

We now turn to an examination of several alternative mechanisms that overcome the free rider problem through the imposition of financial transfers. In order to avoid relying on altruism we assume that agents have no altruistic motivation, and so drop the $f^{AL}(AL)$ term from agents' utility functions for the remainder of the paper.

4. MICRO-PAYMENT MECHANISMS

We wish to encourage users to balance what they take from the system with what they contribute. A natural approach is to charge users for every download and to reward users for every upload. In this section, we propose and analyze a micro-payment mechanism designed according to this principle, as well as a variant of the basic mechanism.

Let us start with a detailed description of our micro-payment mechanism. For each user the server tracks the number δ of files downloaded, and the number ν of files uploaded during the time period. At the end of each period, each user is charged an amount $C = f(\delta - \nu)$. We assume that f is linear with a coefficient α representing the cost/reward per file (e.g., \$0.05), so that the global sum of all micro-payments is 0. Individual users, however, may reduce their monthly charges or even make a profit by uploading more than they download.

Before considering the equilibria that arise under this mechanism, we must make some assumptions so that the mechanism can be represented in our model. Let σ^{-i} be the total number of units shared by agents other than a_i , and δ^{-i} be the total number of units downloaded by agents other than a_i . If agent a_i chooses the action¹ (σ_s, δ_d) then we express the expected value of FT (a_i 's expected payment to the system) as $E[FT] = \alpha \left(d - \delta^{-i} \frac{s}{\frac{n-2}{n-1}\sigma^{-i} + s} \right)$. This reflects the assumption that the central server matches downloaders uniformly at random with shared units, with the constraint that no agent will download from himself. Finally, we make two assumptions about agents' relative preferences for different outcomes. First, we assume that $f^{AD}(1) > \alpha$: the utility agents gain from downloading one file exceeds the micro-payment charged for downloading one file. Second, we assume that $f_i^{DS}(1) + f_i^{BW}(1) < \alpha$: the disutility agents incur from sharing one file and uploading it once is less than the micro-payment that they are credited for uploading it.

We can now consider the equilibria that result from the micro-payment mechanism. A unique, strict equilibrium is $\Sigma = (S_1 = (\sigma_2, \delta_2), \dots, S_n = (\sigma_2, \delta_2))$. Since we have assumed $f^{AD}(1) > \alpha$ agents have an incentive to download as much as possible—their marginal profit per file is reduced, as compared to the case discussed in section 3, but it remains positive. Thus δ_2 dominates δ_1 and δ_0 . If all agents other than a_i follow the strategy $S = (\sigma_2, \delta_2)$, and a_i follows the strategy $S_i = (\sigma_j, \delta_2)$, a_i can calculate his expected utility for the different values of j . He will have $FT = \alpha(2 - 2(n-1) \frac{s}{2n-4+s})$. Given our assumption about

¹To simplify the exposition we assume that σ_s denotes sharing exactly s units, and likewise δ_d denotes downloading d units. This assumption is not needed for our results.

the cost of uploading a file, a_i will strictly prefer the strategy $S_i = (\sigma_2, \delta_2)$; thus we have shown that Σ is a strict equilibrium. Now we show uniqueness of the equilibrium. Note that it is dominant for all agents to choose δ_2 , as described above. Thus δ^{-i} is $2n-2$ in all equilibria for all i . Since $f_i^{DS}(1) + f_i^{BW}(1) < \alpha$, sharing is worthwhile for an agent if every unit shared yields at least one unit of expected uploads. Substituting $s = 2$ into the expression for expected number of uploads from the equation above, we find that it is thus worthwhile for an agent to choose the action σ_2 when $2(n-1) \frac{2}{\frac{n-2}{n-1}\sigma^{-i} + 2} \geq 2$. Rearranging, we find that σ_2 is the most profitable strategy as long as $\sigma^{-i} \leq 2(n-1)$. This condition must always hold since there are only $n-1$ agents other than i and each agent can only share up to 2 units; hence Σ is a unique equilibrium.

Users strongly dislike micro-payments: having to decide before each download if a file is worth a few cents imposes mental decision costs [6]. To address this problem we introduce a quantized micro-payment mechanism where users pay for downloads in blocks of b files. At the end of a time period, the number of files downloaded by a user is rounded up to the next multiple of b , and the user is charged for the number of blocks used. The pricing mechanism for serving files is unchanged. Note that when $b = 1$ we return to the original micro-payment mechanism, while we approach a purely flat-rate pricing plan as b grows. We omit discussion of this class of mechanisms for space reasons; in short, the same equilibrium holds as discussed above, except that users will download a number of files evenly divisible by b .

5. REWARDS FOR SHARING

In the full version of our paper (to appear at WELCOM'01) we consider mechanisms that make use of an internal currency called "points." (Similar ideas have been used by a variety of web services, e.g. *www.mojonation.net*.) Agents are allowed to buy points either with money or with contributions to the network, but they are not allowed to convert points back into money. Above, we focused on influencing users' consumption by penalizing downloads and rewarding uploads. Here we consider rewarding agents in proportion to their level of sharing rather than the number of actual uploads they provide, while still penalizing downloads. Specifically, agents who share are paid an amount proportional to $\int M(t)dt$, where $M(t)$ is the amount of data in megabytes available for download at time t , and the integral is taken over one time period. We provide an equilibrium analysis of several point-based mechanisms and discuss various implementation considerations.

6. EXPERIMENTS

The previous sections analyzed the existence of equilibria for all our mechanisms under simplifying assumptions. Here we test our mechanisms in simulations that more accurately reflect the real world. We enrich our theoretical model by introducing different types of files and agents, and by considering risk-averse agents.

We consider files of several kinds and agents of several types. Recall that the type of an agent is determined by the agent's utility function; in our experiments agents differ according to their preferences for different kinds of files. More specifically, agent utility functions differ as follows:

- **Altruism:** $f(AL) = \alpha AL$ where α is drawn uniformly from $[\alpha_{\min}, \alpha_{\max}]$.
- **Disk space:** the function $f(DS)$ is set to emulate an agent with maximal storage space d , where d is chosen uniformly from $[d_{\min}, d_{\max}]$.
- **File type preferences:** the term $f(AD)$ is decomposed into $\sum_i \beta f_i(AD_i)$, where each i represents a different kind of file. Agents' preferences for each kind of file are reflected by different f_i functions. The factor β is chosen uniformly at random for each agent.

We model agents' utility for money as $U(x) = A \ln(1 + \frac{x}{A})$. As A tends to infinity, U becomes linear; this allows us to observe changes as agents go from risk-averse to risk-neutral. This model is supported by experimental evidence [4].

6.1 Learning Algorithm

We take an approach similar to that of fictitious play [2] to model the behavior of agents. Agents behave as if other agents' strategies were fixed (i.e., as though other agents do not act strategically), and make a best response based on their observations of other agents' actions. Although agent behavior is not strategic in this model, strategy convergence corresponds to a Nash equilibrium. An agent can acquire knowledge either of the joint distribution of other agents' strategies, as in a fictitious play model, or of the expected payoffs associated with its own strategies. In a sufficiently symmetric and regular world populated by sufficiently many agents, the joint distribution can safely be neglected. As P2P systems typically involve very large numbers of agents, agents in our model attempt to learn the payoffs associated with their own strategy, without modeling other agents.

Agents use the temporal difference (TD) Q-learning algorithm to learn these best responses. This algorithm learns the expected utilities of (state,action)-pairs (called Q-values). We use the standard update equation for TD Q-learning, $Q(a, s) \leftarrow (1 - \alpha)Q(a, s) + \alpha(P(a, s) + c \cdot \max_{a'} Q(a', s'))$, where a is the action that the agent took, s is the current state, s' is the new state and $P(a, s)$ is the payoff of the current round (both are chosen probabilistically by the model as a function of other agents' behavior). The decay $0 < \alpha < 1$ and the future income discount $0 < c < 1$ are fixed.

6.2 Experimental Results

First, our simulations confirm the existence of equilibria for the micro-payment and point-based mechanisms, as our analysis predicted. Figure 1 shows that strategies converge to an equilibrium. Second, we demonstrate that our model




Figure 1: Strategy convergence (logarithmic scale).

is complex enough to exhibit non-trivial effects. Fig. 2 shows the behavior of non-altruistic agents in the presence of altruistic agents under the point-based mechanism. As the proportion of altruistic agents increases from 0 to 1, non-altruistic agents discover that they can download more and

therefore have to share more to compensate for the point cost of their downloads. Third, we tested the robustness of our simulations. Overall, we found that our simulation was very robust, producing qualitatively similar results under very different sets of parameters for the number and types of files and the size of the action space for agents. As an example, two runs of the experiment described above, with agents given respectively 9 and 35 actions in their strategy spaces, produced essentially the same result (Fig. 2). Finally, we studied the influence of risk-aversion on agent's behavior in the micro-payment scheme (Fig. 3). We plot the number of files shared in the system as a function of A , agents' value for money. As A decreases, agents become more risk averse. Risk averse agents tend to cut their spending and scale down their contributions to the system because of their uncertainty about how many other agents will download their shared files.

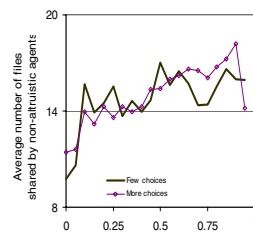


Figure 2: Files shared as a function of the proportion of altruistic agents.




Figure 3: Risk-aversion in micro-payment mechanism.

7. CONCLUSION

The free-rider problem is a real issue for P2P systems, and is likely to become even more important in commercial systems. We have given a simple game theoretic model to analyze agent behavior in centralized P2P systems and shown that our model predicts free riding in the original Napster mechanism. We have proposed and analyzed several different payment mechanisms designed to encourage file sharing in P2P systems. Finally, we presented experimental results supporting our theoretical analysis.

Acknowledgements

We would like to thank Mark Lillibridge for his extensive creative and technical contributions to this paper.

8. REFERENCES

- [1] Eytan Adar and Bernardo A. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.
- [2] D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [3] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [4] C. Grayson. *Decisions Under Uncertainty: Drilling Decisions by Oil & Gas Operators*. Ayer Company, 1979.
- [5] G. Marwell and R. Ames. Experiments in the provision of public goods: I. resources, interest, group size, and the free-rider problem. *American J. of Sociology*, 84, 1979.
- [6] A. M. Odlyzko. The history of communications and its implications for the internet, 2000. www.research.att.com/~amo.
- [7] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [8] P. Gummadi S. Saroiu and S. Gribble. Measurement study of peer-to-peer file sharing systems. TR UW-CSE-01-06-02, 2001.
- [9] J. Sweeny. An experimental investigation of the free-rider problem. *Social Science Research*, 2, 1973.
- [10] B. Thorn and T. Connolly. Discretionary data bases. *Comm. Research*, 14(5), October 1987.