Incentives for Sharing in Peer-to-Peer Networks

Philippe Golle[1,⋆], Kevin Leyton-Brown[1,⋆],
Ilya Mironov[1,⋆⋆], and Mark Lillibridge[2]

[1] Computer Science Department, Stanford University
{pgolle,kevinlb,mironov}@cs.stanford.edu
[2] Systems Research Center, Compaq Computer Corporation
mark.lillibridge@compaq.com

Abstract. We consider the *free-rider* problem in peer-to-peer file sharing networks such as Napster: 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 this work with results from experiments with 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 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, as do most of its competitors; P2P networks also allow users to exchange other digital content.

The work of serving files in virtually all current P2P systems is performed for free by the systems' 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 a very large proportion of its users contribute nothing to the system [2, 11]. 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 [8, 12]. 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 [13].

This problem is not simply theoretical. Some P2P systems plan to charge users for access in the near future, both in order to make money for their investors and to pay any needed royalties. 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

⋆ Supported by Stanford Graduate Fellowship
⋆⋆ Supported by NSF contract #CCR-9732754

using an 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 so make it more competitive with other commercial P2P systems.

In the following section, we introduce our formal game theoretic model. Section 3 describes 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 turn to a more formal, game theoretic characterization of the problem. (Readers unfamiliar with game theoretic analysis may consult [5, 10].) First, we describe the game that we use to model the file sharing scenario during one time period (e.g., one month). n agents participate in the system; we denote them $a_1, \ldots, a_n$. Each agent a_i's strategy, denoted $S_i = (\sigma, \delta)$, consists of two independent actions:

1. **Sharing:** Agents select what proportion of files to share. In our model, sharing takes three levels: σ_0 (none), σ_1 (moderate) or σ_2 (heavy).
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 (none), δ_1 (moderate) or δ_2 (heavy).

2.1 Agent Utility

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

- **Amount Downloaded (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.
- **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; that is, each agent's utility functions is a sum of arbitrary functions, each of which maps one of the above variables to a dollar value. Furthermore, we assume that agents are risk-neutral, and so agents' utility for money is linear. We can thus write the equation for agent a_i's utility function as:

$$U_i = \left[f_i^{AD}(AD) + f_i^{NV}(NV) + f_i^{AL}(AL)\right] - \left[f_i^{DS}(DS) + f_i^{BW}(BW)\right] - FT.$$

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. Likewise, DS and BW only contribute negative utility, explaining the subtraction of f^{DS} and f^{BW} above. Finally, we make two assumptions about agents' relative preferences for different outcomes:

$$f^{AD}(k) > k\beta \tag{1}$$

$$f^{DS}(k) + f^{BW}(k) < k\beta \tag{2}$$

First, in inequality (1) we assume that the monetary equivalent of the utility agents gain from downloading files at level k is more than $k\beta$, for some constant β. Second, in inequality (2) we assume that the monetary cost to agents of sharing files at level k and uploading them at level k is less than $k\beta$.

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: 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 as $\Sigma = \{S_1 \dots S_n\}$. 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 *strict 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 identities of files provided by users, which may be needed both to pay royalties to the appropriate copyright holders and to detect users who make false claims about the files they share. File identification may be achieved by a cryptographic watermarking scheme [1, 7]; alternately, users who spoof files could be penalized.

One likely payment model for peer-to-peer systems is a flat rate membership fee. 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; note especially that the fact that flat fees are unrelated to agents' behavior implies that such pricing does not help avoid a free-rider problem.

3 The Napster System

We analyze the Napster system that operated from May 1999 through July 2001, since it is probably the best-known peer-to-peer application. This is one of the simplest system 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, $\Sigma = \{(\sigma_0, \delta_2), \ldots, (\sigma_0, \delta_2)\}$ 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 that 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; when no other agent shares his utility is zero regardless of how much he intends to download. We can therefore 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 unique.

We have identified a unique equilibrium in which nothing is shared and there is nothing to download. Yet songs were plentiful and actively traded on Napster. We identify two reasons that users might have contributed. First, Napster offered its service free of charge and went to great lengths to foster a sense of community among its users, notably through such features as chat-rooms, a newsletter, and messaging between 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. This could be circumvented, but only by manually moving songs to another directory after download or explicitly shutting down the Napster service. Again, because the donation of resources cost users very little, many users may not have bothered to "opt out". 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.[1] 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 way users behaved 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. This conclusion coheres with the empirical research cited in the introduction claiming that only a small proportion of Gnutella users share any files. Likewise, it supports our claim that Napster experienced a free-rider problem: regardless of the contributions of others, selfish agents had no incentive to share. As with all such problems, the situation in this case appears somewhat paradoxical, because all agents would be better off if they all shared (due to the resulting increase in NV, the variety of files available on the network) than they are under the unique equilibrium.

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 term $f^{AL}(AL)$ from agents' utility functions from here until section 6, in which we present our experimental results. (Of course, all of our results also hold for agents motivated by altruism.)

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 then for every upload. In this section, we propose and analyze

[1] More realistically, we could have assumed three types of agents: those whose level of altruism led them to take each of the three levels of sharing. We analyzed the simpler case to simplify the exposition; the analysis of the case with three agent types proceeds in the obvious way.

a micro-payment mechanism designed according to this principle, as well as a variant of the basic mechanism.

We begin with a detailed description of our micro-payment mechanism. For each user the server tracks the number δ of files downloaded, and the number υ of files uploaded during the time period. The server is aware of all such transfers since it processes all download requests. Note also that there exist standard cryptographic protocols (fair exchange, [3]) to ensure that both parties agree on whether their exchange was aborted or ended successfully. At the end of each period, each user is charged an amount $C = g(\delta - \upsilon)$. We assume that g is linear with a coefficient β representing the cost/reward per file (e.g., $0.05), and that the value of this coefficient is such that inequalities (1) and (2) hold. Note that the global sum of all micro-payments is 0; individual users, however, may make a profit by uploading more than they download.

Before considering the equilibria that arise under this mechanism, we must simplify it so that it can be represented in our model.[2] 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 action[3] (σ_s, δ_d) then we express the expected value of FT (a_i's expected payment to the system) as

$$E[FT] = \beta \left(d - \delta^{-i} \frac{s}{\frac{n-2}{n-1}\sigma^{-i} + s} \right) . \tag{3}$$

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. Note that β is the coefficient representing the cost per net unit downloaded.

Proposition 1. $\Sigma = \{(\sigma_2, \delta_2), \ldots, (\sigma_2, \delta_2)\}$ *is a unique, strict equilibrium.*

Sketch of proof. Inequality (1) states that $f^{AD}(k) > k\beta$. Therefore, 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 $S_i = (\sigma_s, \delta_2)$, a_i can calculate his expected utility for the different values of s: he will have $E[FT] = \beta \left(2 - 2(n-1) \frac{s}{2n-4+s} \right)$. Given our assumption about the cost of uploading a file, a_i will strictly prefer the strategy $S_i = (\sigma_2, \delta_2)$; thus 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

[2] Although the following analysis makes explicit use of the fact that there are only three levels of sharing and of downloading possible, this restriction has been made only for ease of exposition; a similar (albeit more complicated) proof exists for any number of levels of sharing and downloading.

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

for all i. Inequality (2) states that $f^{DS}(k) + f^{BW}(k) < k\beta$: sharing is worthwhile for an agent if every unit of sharing yields at least one unit of uploading on expectation. Substituting $s = 2$ into the expression for expected uploading from equation (3), 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 always holds since there are $n-1$ agents other than i and each agent can only share up to 2 units; hence Σ is unique. □

Note that the same analysis does not suffice for the case of risk-averse agents. The problem is that agents directly control their number of downloads, but only indirectly control their number of uploads through the number of files shared. Depending on the nature and degree of agents' risk aversion and their particular utility functions, they may prefer to reduce their downloads to reduce their worst-case payments to the network. Since the behavior of risk-averse agents depends so heavily on the particular assumptions we make about them we do not give a formal analysis here; however, we return to this issue in section 6.

4.1 Quantized Micro-Payment Mechanisms

Empirical evidence suggests that users strongly dislike micro-payments: having to decide before each download if a file is worth a few cents imposes mental decision costs [9]. Users often prefer flat pricing plans, even when such mechanisms may increase their expected costs. To address this problem we introduce a quantized micro-payment mechanism where users pay for downloads in blocks of b files, where b is a fixed parameter. 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. In practice, we consider values of b on the order of the number of files that an average user would download per time period.

We do not present an analysis of this class of mechanisms, for two reasons. First, in the abstract these quantized mechanisms are the same as general micro-payment schemes, except that it is irrational for agents to download a number of files that is not an even multiple of b (unless of course an agent has reached the maximum number of files that he desires). The key advantage of this class of mechanisms is that agents are spared the mental decision costs associated with per-download pricing; since we do not explicitly include this cost in agents' utility functions, the elimination of this cost does not affect the analysis. Second, this class of mechanisms does not fit easily into our simplistic model for user actions: as we allow only three levels of downloading, it is unclear what to quantize. From the analysis in section 4 it is easy to see that if we charge agents the same for δ_1 as for δ_2 the original equilibrium is preserved: agents will simply be provided with additional incentive for taking the actions that they would take anyway.

There are, however, some interesting practical issues arising from this class of micro-payment mechanisms. First we expose a way in which agents could gain through collusion, and present two possible remedies. We also consider problems arising from the trading of rare files and suggest one solution.

First we examine an important way in which this mechanism could be attacked. Quantized micro-payment mechanisms have the property that after one file has been downloaded, the marginal cost of downloading the remaining $b-1$ files belonging to the same block is zero. Towards the end of a payment period, users may take advantage of zero-marginal-cost downloads left in their account to download files from friends, in order to cause these friends to receive the payment for serving these files. A coalition of users could agree to download excess files from each other and share the profit. The cost to the server is proportional to the difference between the number of files in a block and the average number of files actually desired by agents. However, this collusion can only reduce profits back to the case of simple micro-payments discussed above, where every download corresponds to an upload credited to another agent.

We can modify the quantized payment mechanism so that it is harder for users to direct their zero-marginal-cost downloads to other specific users. This makes it harder for a coalition to generate money for itself: if a user has no control over who is making a profit out of his downloads, this attack becomes less profitable. We could modify our mechanism in one of two ways. First, the server can reply to each download request with a list of users serving files that match the request, but hide the identities of all the users. A user can choose to download from any of the locations listed, but cannot specifically single out his friends. Second, the server can reply to each download request with a random subset of all the users serving files that match the request. In this case user identities do not need to be hidden.

Observe that these solutions only make it less efficient to direct zero-marginal-cost downloads to friends, but by no means make it impossible. Furthermore, these solutions are less and less effective as the number of users sharing a given file decreases, because a user who stores files that are sufficiently rare will receive a large fraction of all the download requests for these files. Thus we propose that rare files (file for which the number of copies available falls below a threshold) be treated differently from files that are more frequent. We cannot simply refuse to credit users who serve rare files, because this would create a strong disincentive for introducing new files into the system. Instead, the central server can give users no credit for serving rare files, but keep track of all the exchanges of rare files. If a file exceeds the threshold of frequency and becomes sufficiently popular, the users who shared it while it was still rare can be credited retroactively.

5 Rewards for Sharing

Previously, we focused on influencing users' consumption by penalizing downloads and rewarding uploads. We take a different approach here: we continue to penalize downloads, but we now consider rewarding agents in proportion to the amount of material they share rather than the number of uploads they provide.

The mechanism we consider makes use of an internal currency, "points."[4] 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. Since agents cannot "cash out" their points, they might be allowed to maintain a balance from one time period to the next. We do not consider such a rolling balance since we model a single time period; furthermore, in a repeated equilibrium agents would have no incentive for accumulating more points than they spend.[5] Agents' payment for sharing is $\int M(t)\ dt$, where $M(t)$ is a measure of the amount of data made available for download at time t and the integral is taken over the whole time period. Downloading a file costs cm points, where m is a measure of the file's size and c is a constant. Intuitively, c represents the number of hours one file must be shared in order for the cost of one download to be waived.

As above, we must simplify this mechanism in order to analyze it according to our game theoretic model. One point costs β, where β is set to a value such that inequalities (1) and (2) hold. Furthermore, assume that all files have the same size and that agents always share files for the same amount of time (one time period). Each level of sharing in one time period earns one point (e.g., σ_2 is worth two points). We take $c = 1$, so each level of downloading costs one point. Downloaded files are not shared in the same time period as they were downloaded. As above, we assume that downloaders are matched uniformly at random with shared units, and that no agent may download from himself. Thus, if a_i shares at level s then his expected number of uploads, υ_i, is:

$$E[\upsilon_i] = \delta^{-i} \frac{s}{\frac{n-2}{n-1}\sigma^{-i} + s}. \tag{4}$$

Proposition 2. *$\Sigma = \{(\sigma_2, \delta_2), \ldots, (\sigma_2, \delta_2)\}$ is a strict equilibrium.*

Sketch of proof. Consider $n-1$ agents playing the strategy $S = (\sigma_2, \delta_2)$, and an agent a_i who must determine his best response. From inequality (1), $f^{AD}(k) > k\beta$, δ_2 dominates δ_1 and δ_0. Thus a_i will play $S = (\sigma_s, \delta_2)$ and must choose a value for s. If a_i plays σ_0, σ_1 or σ_2 his expected number of uploads (given the other agents' strategies) will be 0, just under 1 or 2 respectively, and thus his expected financial transfer to the system will be 2β, slightly more than β or 0. Inequality (2), $f^{BW}(k) + f^{DS}(k) < k\beta$, tells us that agents prefer to share at level k and upload at level k than to pay the system for k points. Since sharing at level 2 leads to uploading at level 2 on expectation, given the other agents' strategies, a_i's expected utility is maximized by the action σ_2. Therefore Σ is a strict equilibrium. □

However, Σ is not a unique equilibrium. Indeed, point-based schemes have the drawback that they can give rise to a degenerate equilibrium in which all

[4] Similar ideas have been used by a variety of web services, e.g. *www.mojonation.net.*

[5] Agents might be encouraged to accumulate points if high balances were rewarded with faster downloads, early access to popular files or other privileges. However, as this intriguing possibility depends heavily on agents' particular utility functions as well as on details about the file sharing system, we do not pursue it further here.

agents download at the highest level and share nothing at all. This can be proven in several ways; we present the simplest proof, which requires the assumption that points have no value if not redeemed to pay for downloads.

Proposition 3. $\Sigma = \{(\sigma_0, \delta_2), \ldots, (\sigma_0, \delta_2)\}$ *is a strict equilibrium.*

Sketch of proof. Consider $n-1$ agents playing the strategy $S = (\sigma_0, \delta_2)$, and an agent a_i who must determine his best response. As above, δ_2 is dominant by the first assumption. Thus a_i who will play $S = (\sigma_s, \delta_2)$ and must choose a value for s. Since all other agents play σ_0, there exist no files to download. Thus gaining points will yield no utility for a_i, by the assumption in the preamble. Furthermore, since all other agents play δ_2, a_i will be made to serve files for all other agents' download requests, bringing him negative utility. He is therefore best off following strategy S, and so Σ is a strict equilibrium. □

This analysis leaves it unclear what equilibrium will be reached in play. We attempt to provide some answers in our experimental section.

We now consider more practical issues arising from this mechanism. First we consider agents' opportunities for collusion and for other undesirable exploitation of the system, and discuss remedies.

Unlike quantized micro-payments, this mechanism does not interfere with download patterns. Instead it always gives the right incentives for consumption, since there is no way for colluding users to make money by downloading from each other. However, this mechanism alters agents' incentives for sharing files. The key problem is that agents have negative utility for the consumption of bandwidth, which only occurs when shared files are actually downloaded. In order to conserve bandwidth, agents may make their collections available at low-usage times, or alternately offer unpopular files. This may reduce the overall value of the network. A possible remedy is to offer distributors different rewards based on expected download demand. The formula to reward distributors thus becomes $\int M(t)\lambda(t)\ dt$, where $\lambda(t)$ is a scaling factor proportional to expected demand. This ensures that the files are available at the right times. The problem of users preferring to share unpopular files can also be addressed through the introduction of a similar coefficient.

Another challenge is that agents cannot be expected to make (and honor) a commitment to share a file for hours into the future. It is much more likely that agents will start and stop sharing unpredictably, sharing only when their computer is idle. A mechanism that accommodates such behavior is likely to be more useful; however, this accommodation must be balanced by ensuring that agents are not able to cheat by suddenly claiming to lose their idle status as soon as they receive an upload request.

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.

6.1 Experimental Setup

We extend upon our theoretical model in two ways. First, we consider action spaces for agents more fine-grained than the three levels of downloading and sharing discussed so far. Second, we consider files of several kinds and agents of several types. Recall that agents with different types have different utility functions; in our experiments agents differ according to their (fixed) preferences for different kinds of files. Agent utility functions differ as follows:

- **Altruism:** $f(AL) = \rho AL$ where ρ is drawn uniformly from $[\rho_{\min}, \rho_{\max}]$.
- **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 $\mu \sum_i f_i(AD_i)$, 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 μ is chosen uniformly at random in $[\mu_{\min}, \mu_{\max}]$ for each agent.

In the simulation of micro-payment mechanisms, our agents are stateless (they do not keep track of the amount of money they spend or make). In the simulation of point-based mechanisms, we define states according to the number of points accumulated by an agent. Points have no intrinsic value to agents, but an agent who runs out of points must purchase more with money.

All the other parameters of our mechanisms are fixed and equal for all agents. 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-aversion to risk-neutrality. This model of risk aversion is supported by experimental evidence; see, e.g., [6].

6.2 Learning Algorithm

We take an approach similar to that of fictitious play [4] 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. This is because convergence corresponds to the situation where each agent's best response is to maintain his strategy, given the assumption that all other agents are fixed in their strategies. An agent could attempt to learn either the joint distribution of other agents' strategies, as in a fictitious play model, or 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 modelling 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). The best response is the action that gives the highest expected payoff. The Q-learning algorithm assumes that the environment does not evolve over time, but decay enables agents to also do well in a slowly changing environment. 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.

6.3 Experimental Results

First, our simulations confirm the existence of equilibria for the micro-payment and point-based mechanisms in the richer setting described above. These equilibria generalize those described in our analysis, giving evidence that our experimental assumptions are reasonable. Fig. 1 shows strategy convergence:

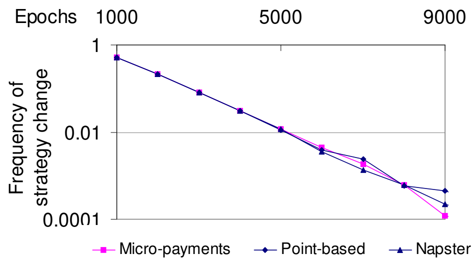

Figure 1: Strategy convergence (logarithmic scale).

Second, we demonstrate that our model 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 the simulations to be quite robust: we observed qualitatively similar results under very different sets of parameters for the number and types of files and for the size of the action space for agents. Agents with a wider choice of actions (more options for downloads and sharing) achieve higher payoffs, but the results remain quantitatively the same. As an example, two runs of the experiment described above, with agents given 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. We observed that 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.

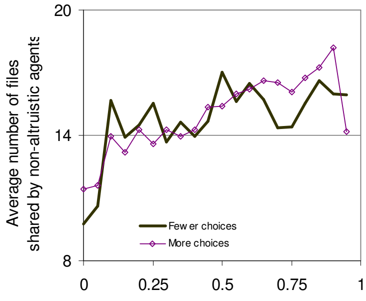

Figure 2: Files shared as a function of the proportion of altruistic agents.

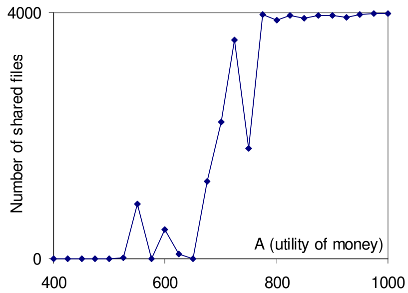

Figure 3: Risk-aversion in micro-payment mechanism.

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 proposed a simple game theoretic model of agent behavior in centralized P2P systems and shown that our model predicts free riding in the original Napster mechanism. We analyzed several different payment mechanisms designed to encourage file sharing in P2P systems. Finally, we have presented experimental results supporting our theoretical analysis.

References

1. Secure digital music initiative. `http://www.sdmi.org`.
2. E. Adar and B. Huberman. Free riding on Gnutella. *First Monday*, 5(10), 2000.
3. N. Asokan. *Fairness in Electronic Commerce*. PhD thesis, University of Waterloo, Ontario, Canada, 1998.
4. D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
5. D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
6. C. Grayson. *Decisions Under Uncertainty: Drilling Decisions by Oil & Gas Operators*. Ayer Company, 1979.
7. B. Macq, editor. *Special Issue on Identification and Protection of Multimedia Information*, volume 87(7), July 1999.
8. 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.
9. A. M. Odlyzko. The history of communications and its implications for the Internet, 2000. Available online at `http://www.research.att.com/~amo`.
10. M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
11. S. Saroiu, P. Gummadi, and S. Gribble. Measurement study of peer-to-peer file sharing systems. Tech Report UW-CSE-01-06-02, University of Washington, 2001.
12. J. Sweeny. An experimental investigation of the free-rider problem. *Social Science Research*, 2, 1973.
13. B. Thorn and T. Connolly. Discretionary data bases. *Comm. Research*, 14(5), October 1987.