

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

The recent and unprecedented surge of public interest in peer-to-peer file sharing has led to a variety of interesting research questions. In this paper, we will address the incentive issues that arise in such file sharing systems. In particular, there is a *free-rider* problem in traditional peer-to-peer networks such as Napster: individual users are provided with no incentive for sharing their own files and thereby adding value to the network. Instead, such services currently rely upon altruistic behavior from their users. We take a different approach, examining the design implications of the assumption that users will selfishly act to maximize their own rewards. We construct a formal game theoretic model of the system and analyze equilibria of user strategies under several novel payment mechanisms. Finally, we support and extend upon the predictions of our game theoretic model by presenting 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, has attracted immense public attention for the P2P paradigm as well as tens of millions of users for itself. Napster specializes in helping its users to trade music, as do most of its competitors; however, P2P applications exist to facilitate the exchange of virtually every kind of digital document.

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, a recent study of the

*Supported by Stanford Graduate Fellowship

†Supported by NSF contract #CCR-9732754

Gnutella network found that more than 70% of its users contribute nothing to the system. [1] 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. The communal sharing of information goods in “discretionary databases” and the resulting free-rider problem received study long before the advent of P2P systems [10]; of course, the free-rider problem has also been extensively studied outside this context [6, 9].

This problem is not simply theoretical. There is reason to believe that some P2P systems will begin charging users for access in the near future. A commercial environment may be even less able to provide altruistic incentives sufficient to motivate users to contribute to the system. A system run for profit may not receive the level of altruistic ‘donations’ that power a free community—users who are charged for the use of a system are likely to feel that they have already contributed in another way. 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. This would make a greater variety of files available, which would increase the system’s value to its users. Also, the files would be distributed across a larger number of machines, increasing parallelism and reliability.¹ The ability to overcome the free-rider problem and to thereby improve the system’s value could be crucial to the success of a newly-commercialized P2P system.

The rest of this paper is organized as follows. Section 2 introduces our formal game theoretic model. Section 3 presents the Napster system, which we will 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

Our results in this paper apply to peer-to-peer file exchange in general; however, for concreteness we will use Napster as a motivating example. Throughout the paper we will consider a file sharing system consisting of centralized servers which maintain a database of the files currently available on the network and which connect download requests with available clients. It is important to note that these servers merely store a list of files available on users’ machines: they are not directly involved in the file-sharing, although they are able to monitor all exchanges.

It is important for the servers to be able to determine the identity of files provided by users. First, such information may be necessary to pay royalties to the appropriate copyright holder, or to determine if a given file may legally be

¹This works especially well because most home PC high-bandwidth Internet connections are asymmetrical, with inbound capacity an order of magnitude greater than outbound. It thus requires multiple home PCs to saturate a PC’s inbound connection.

exchanged on the system. Second, this information makes it possible to detect users who make false claims about the files they share; later in the paper we will discuss mechanisms that reward agents for sharing files, but these mechanisms can succeed only if agents can be prevented from spamming or spoofing the system.

These requirements for file identification are best met by some kind of cryptographic fingerprinting scheme. In the particular case of music, creating such a scheme may be somewhat difficult as digitalizing audio CD tracks is somewhat nondeterministic, producing slightly differing bit strings each time. If such a scheme or other means of automatically identifying files is not available, the best fallback scheme is to penalize users accused of spoofing files. This will be more expensive, though, because people will have to be employed to investigate and resolve complaints.

2.1 Game Setup

We now turn to a more formal, game theoretic characterization of the problem. (Readers unfamiliar with game theoretic analysis may consult [4, 8].) First, we describe the game that we will 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 will denote them as a_1, \dots, a_n . Each agent has two orthogonal actions available in each time period:

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

We can now describe an agent a_i 's strategy in time period t . We say that a_i has a *pure strategy* if his strategy takes the form $S(i, t) = (\sigma, \delta)$. (When the time period is unambiguous, we will write simply $S(i)$ to denote a_i 's strategy.) More generally, a_i can have a *mixed strategy*, in which $S(i, t)$ is a probability distribution over all tuples of sharing levels and download levels.

2.2 Agent Utility

Agents' utility functions describe their preferences for different outcomes. There are a variety of different factors that concern agents:

- **Amount to Download (AD):** Agents get happier the more they download.

- **Network Variety (NV):** Agents prefer to have more choice about what they download.
- **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):** Agents might derive some positive utility from the satisfaction of contributing to the operation of 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 will make the assumption 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. We make the standard assumption 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:

$$U_i = f_i^{AD}(AD) + f_i^{NV}(NV) + f_i^{DS}(DS) + f_i^{BW}(BW) + f_i^{AL}(AL) + FT \quad (1)$$

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 a (positive or negative) amount of money that is paid to the agent; f^{FT} would therefore be the identity function.

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. For ease of exposition in our game theoretic analysis in the first part of this paper we will often make the assumption that all agents have the same type. In section 6 we will approach the file sharing problem experimentally; this approach will allow us to discuss the convergence of agent strategies under a wide variety of different agent types. In this section we will also expand our model to consider categories of files that are preferred by different agents. Under this elaboration, f_i^{AD} will be replaced by functions denoting the value to a particular agent of downloading files from each category; the variable AD will be replaced by variables representing the amount of files downloaded from each category; and agent actions will be expanded to specify an amount to download from each category. Until section 6, however, we will restrict ourselves to the case of a single category.

Although it is not desirable to restrict ourselves by specifying particular f functions, we can make several observations about the shapes of these functions. We can observe that f^{AD} , f^{NV} and f^{AL} must be strictly monotonically increasing, with minimum value 0, as these variables only ever contribute positive utility for any agent. However, in our model we assume that there exist values of k_1 and k_2 such that f_i^{AD} and f_i^{AL} are constant for values greater than k_1 and k_2 respectively: there exists both an amount of files downloaded and an

amount of altruism beyond which a given agent is satiated. We make no such assumption about a limit on the utility that agents receive from the variety of files available on the network. In a similar way, we can observe that f^{DS} and f^{BW} must be strictly monotonically decreasing with maximum value 0, as these variables only ever contribute negative utility for any agent. However, we assume that there exist k_3 and k_4 such that f^{DS} and f^{BW} take the value $-\infty$ for all values greater than k_3 and k_4 respectively: this represents the fact that agents have hard limits on the amount of disk space and bandwidth that they can make available to the network. We make the further assumption that neither f^{DS} nor f^{BW} is superlinear.

We have concluded that some of the variables in agents' utility functions are always positive while others are always negative; only the financial transfer (FT) variable can take both positive and negative values. To make the utility function more intuitive we change the sign of the negative variables, rewriting it as follows:

$$U_i = [f_i^{AD}(AD) + f_i^{NV}(NV) + f_i^{AL}(AL)] - [f_i^{DS}(DS) + f_i^{BW}(BW)] + FT \quad (2)$$

Finally, we consider the relationship between the actions available to an agent and that agent's utility function. Sharing files has a negative impact on an agent's utility, since utility decreases as the amount of disk space (DS) devoted to shared files increases. Furthermore, uploads further reduce utility by consuming bandwidth (BW), and the likelihood of uploads is increased as more files are shared. On the other hand, agents can also gain utility from sharing due to the satisfaction they get from being altruistic (AL). Although the only utility agent a_i receives from sharing files is due to altruism, all other agents $a_j, j \neq i$ receive increased utility due to the increase in the variety of shared files on the network (NV). Agents can increase their own utility by increasing their downloads (AD), as long as $NV \neq 0$ (i.e., files are being shared). Lastly, the designer of the peer-to-peer system can charge or credit agents an amount of money (FT) that depends on their actions. In this paper we will consider several different mechanisms, differing primarily in the way that this amount of money is determined.

2.3 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 will 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 will 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, if Σ is a *strong Nash equilibrium* then every agent would be strictly worse off if he were to change his strategy, given that all other agents'

strategies are fixed. Finally, we say that an agent has a *dominant strategy* if his best action does not depend on the action of any other agent. If this agent's dominant strategy is unique we say that it is *strongly dominant*; otherwise we say that it is *weakly dominant*.

2.4 Flat Rates

One likely payment model for peer to peer systems is some kind of flat rate membership fee per time period, for example as a way of recovering royalty costs or the overhead involved in running servers. We do not explicitly consider this option anywhere in the discussion that follows, because it has no impact on the equilibria that arise from any mechanism. Of course the magnitude of the flat rate charge *does* have an impact on agents' decisions about whether or not to participate; once they have made this decision, however, the presence or absence of such a fee is irrelevant to their efforts to maximize utility. All the mechanisms discussed here are *compatible* with the addition of flat rate pricing; we do not mean to deemphasize its possible importance to commercial applications of peer-to-peer networking. Most importantly, however, 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

Napster is probably the best-known peer-to-peer application, designed to help its users exchange music files. Napster consists of a number of servers² and tens of millions of clients.

Using the model described in section 2, we start with an equilibrium analysis of the Napster system. Initially, we will disregard the 'altruism' component of agents' utility functions, though we will return to it in section 3.1. Napster can be seen as the simplest mechanism that can be represented by our model: regardless of the actions of agents, it imposes no financial transfers. Last, as described above, we will assume that all agents have the same utility function to simplify the exposition. We are therefore concerned with the utility function $U = [f^{AD}(AD) + f^{NV}(NV)] - [f^{DS}(DS) + f^{BW}(BW)]$.

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 pure strategy $S = (\sigma_0, \delta_0)$, 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 the same (zero) regardless of how much he intends to download. The action δ_2 is therefore a best response, and Σ as given above is a weak equilibrium. (The equilibrium is weak

²Although to its users Napster appears to be a single network, it is in fact several distinct networks—each of its servers can serve a maximum of 10000 simultaneous clients.

because agents are indifferent between the actions δ_0 , δ_1 and δ_2 —all choices have the same effect because there are no files available to be downloaded.)

From this analysis we can see that the strategy $S = (\sigma_0, \delta_2)$ is strongly 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) , as explained above; if any other agent does share then S yields strictly higher revenue than any other strategy. Because Σ is an equilibrium in strongly dominant strategies, it is the only equilibrium.

3.1 Altruism

The dominant strategy leads to an equilibrium in which nothing gets shared and there is nothing to download. Yet songs are plentiful and actively traded on Napster. We identify two incentives that could account for users' willingness to contribute. First, Napster offers its service free of charge and has gone 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 well be sufficient to foster a sense of altruism in users, encouraging them to contribute resources that cost them very little. Second, Napster offers a (modest) disincentive for non-contribution. By default, the Napster client shares all songs that an agent has downloaded. This can be circumvented, but only by manually moving songs to another directory after download or explicitly shutting down the Napster service. This stealthy approach to taking resources can succeed for the same reason as Napster's first incentive: resources are in most cases so cheap that many users cannot be bothered to "opt out".

In our model, we represent both of these incentives through the *altruism* variable (AL). We use this variable, therefore, to represent both the utility agents gain from the satisfaction of contributing and the increase in utility that agents experience from *not* opting out of Napster's default of file sharing.

We now consider the equilibrium that will arise when some agents are altruistic. In this section we will 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) will be 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 will thus have the dominant strategy given above: (σ_0, δ_2) .

This analysis is arguably a description of the current state of affairs on the Napster system. Some proportion of agents are sufficiently altruistic to share files³ and do so; other agents are not altruistic and share nothing. Regardless

³More realistically, we could have assumed three types of agents: those whose level of

of their level of altruism, agents are unrestrained in their downloads. This conclusion coheres with the empirical research cited in the introduction which claimed that only a small proportion of Gnutella users share any files[1]. This is the *free rider problem*: regardless of the contributions of others, selfish agents prefer not to share. This problem 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).

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 will assume that agents have no altruistic motivation, and so will drop the $f^{AL}(AL)$ term from agents' utility functions.

4 Micro-Payment Mechanisms

We wish to encourage users to balance what they take from the system with what they contribute to the system. 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 registered user the server tracks the number δ of files downloaded, and the number ν of files uploaded during the time period. Each time a file is successfully exchanged between two parties, the server increments the download count δ of the user who downloaded the file and the upload count ν of the user who uploaded it. Observe that the server is aware of all such transfers since it processes all download requests. Note also that there exist standard cryptographic protocols (fair exchange, [2]) 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 $C = f(\delta - \nu)$. The function f maps the difference between downloads and uploads to a financial transfer from agents to the system. We will assume that f is linear with a coefficient representing the cost/reward per file (e.g., \$0.05), so that the mechanism has the property that the global sum of all micro-payments is zero. Of course, other functions could also be used. Although the global sum of payments is zero, note that our mechanism allows individual users to reduce their monthly charges or even to 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.⁴ Let σ^{-i} be the total number of units shared by agents other than

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.

⁴Although 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.

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 (σ_s, δ_d) then we express the expected value of FT (a_i 's expected payment to the system) as:

$$FT = \alpha \left(d - \delta^{-i} \frac{s}{\frac{n-2}{n-1} \sigma^{-i} + s} \right) \quad (3)$$

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 α is the coefficient representing the cost per net unit downloaded. 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, strong 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 $FT = \alpha(2 - 2(n-1)\frac{s}{2n-2+s})$. Given our assumption about 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 strong equilibrium. Now we will show uniqueness of the equilibrium. Note that it is dominant for all agents to choose δ_2 , as described above. Thus d^{-i} will be $2n-2$ 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 equation (3), we find that it will thus be 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.

We observe 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 they only indirectly control their number of uploads through the number of files they share. Depending on the nature 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 their particular preferences we do not give a formal analysis here; however, we return to this issue in section 6.

4.1 Quantized Micro-Payment Mechanisms

It is well known [7] that users strongly dislike micro-payments (having to decide before each download if a file is worth a few cents imposes mental decision costs). Users often prefer flat pricing plans, even when such mechanisms may increase their expected costs. As one way of addressing this problem, we introduce a variant of our micro-payment mechanism in which 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 this many blocks. 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 will not present an analysis of this class of mechanisms, for two reasons. First, they are essentially the same as general micro-payment schemes, with the difference that it is irrational for agents to download a number of files that is not an even multiple of b (except for the limiting case where the agent reaches the maximum number of files that he desires). 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 above 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.

We now 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 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-margin-cost downloads left in their account to get files from friends, for the sole purpose of letting these friends earn money for serving 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 to the average number of files actually desired by agents.⁵

We propose here a modification to the quantized payment mechanism that makes it harder for users to direct their zero-margin-cost downloads to specific friends. 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:

- The server replies to each download request with a list of users serving files that match the request, but obfuscates the identities of the users. The user can choose to download from any of the locations listed, but has no way to specifically single out his friends.

⁵It is worth pointing out, however, that this “attack” can only bring us back to the case of simple micro-payments discussed above, where every download paid for by an agent corresponds to an upload credited to another agent.

- An alternative solution would be to reply to each download request with a random subset of all the users serving files that match the request.

Observe that these solutions only make it less efficient to direct zero-margin-cost downloads to friends, but by no means make it impossible. 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 to treat rare files differently from files that are more frequent. Let us define a rare file as a file for which the number of copies available is below a certain threshold. A simple approach would be to give users no credit for serving rare files. The problem is that this creates a strong disincentive to introduce new files into the system. To allow for new files, we modify our approach slightly. Users get no credit for serving rare files, but the central server keeps track of the rare files that are exchanged. If a file breaks above the threshold of frequency and becomes sufficiently popular, the users who propagated it while it was still rare are credited retroactively.

5 Point-Based Mechanisms

In this section we consider mechanisms that make use of an internal currency which we will call “points.” Agents are allowed to buy points either with money or with contributions to the network, but agents are not allowed to convert points back into money. Since agents cannot “cash out” their points, the mechanism must allow them to maintain a balance from one time period to the next. We will disregard this rolling balance in our equilibrium analysis, however, as in a repeated equilibrium agents have no incentive for accumulating more points than they spend.⁶

First, we discuss a point-based mechanism which is otherwise similar to the micro-payment mechanisms discussed above.

5.1 Micro-Payments in Points

We describe here a second variant of our original micro-payment mechanism designed to circumvent users’ aversion to micro-payments. In this mechanism all micro-payments are charged and credited in points; for example, downloading a file might cost 1 point while uploading a file might earn 1 point.

There are two significant advantages to dealing with points rather than dollars:

- Like quantized micro-payments, this mechanism lets users trade a fixed amount of dollars for a block of b points. Even though files are paid for with points on a per-file basis, where dollars are concerned, the mechanism is essentially quantized.

⁶One way of encouraging agents to accumulate points would be to reward agents having with high point counts with faster downloads, early access to popular files, or other privileges. However, because 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.

- On the other hand, this mechanism does not suffer from the attack to which the quantized micro-payment mechanism was vulnerable. There are no zero-margin-cost downloads and thus there is no incentive for users to download files other than for the purpose of making direct use of them.

The formal analysis of this scheme proceeds exactly as in the discussion of “pure” micro-payments above, with a unique strong equilibrium of all agents following the strategy $S = (\sigma_2, \delta_2)$. However, a difference arises when agents’ maximum desire for downloads changes. If agents (taking the cost of downloads into account) had no desire to download at a level above δ_1 , this mechanism would give rise to a different equilibrium than would the pure micro-payment mechanism. Intuitively, the latter separates agents’ desire for downloads from their desire to make a profit by uploading so that they will continue to indulge in one if their ability to engage in the other is curtailed. Since points are not redeemable for money, agents have incentive only to accumulate enough points to defray the cost of their downloads. The unique equilibrium in this permuted situation would be for all agents to follow the strategy $S = (\sigma_1, \delta_1)$.

5.2 Rewards for Sharing

All of our previous mechanisms have focused on influencing users’ consumption by penalizing downloads and rewarding uploads. We now take a different approach: although we continue to penalize downloads, we now consider rewarding agents not in proportion to the number of uploads they provide, but in proportion to the amount of material they make available. As in the previous section agents are paid in points, though the number of points charged for a download or awarded for sharing is different. Specifically, agents 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. In practice, the server would also take care to cap the number of downloads directed to a distributor at any time, to avoid overwhelming distributors with limited bandwidth resources. Downloading a file costs cm points, where m is the size of the file in megabytes and c is a constant parameter of the system. Intuitively, c represents the number of hours a newly downloaded file must be shared for the cost of the download to be waived. Finally, agents may buy points with money; we assume that one point can be bought for π .

As above, we must simplify this mechanism in order to analyze it according to our game theoretic model. We will assume that all files have the same size (1 MB), and that agents always share files for the same amount of time (1 time period). Each level of sharing in one time period earns one point (e.g., σ_2 is worth two points). We will 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. We make the same assumptions about uploads as above: if a_i shares at level s then his expected number of uploads, u_i , is expressed as:

$$u_i = \delta^{-i} \frac{s}{\frac{n-2}{n-1}\sigma^{-i} + s} \quad (4)$$

As above, we assume that downloaders are matched uniformly at random with shared units, and that no agent may download from himself. We make two assumptions about agents' preferences that mirror the assumptions we made in section 4. First, $f^{AD}(k) > k\pi$ — agents are willing to download any number of files even if they must pay for them directly. Second, we assume that $f^{BW}(k) + f^{DS}(k) > k\pi$ — agents prefer to share and upload k files than to buy k points.

We now consider agents' behavior in equilibrium. As before, $\Sigma = \{(\sigma_2, \delta_2), \dots, (\sigma_2, \delta_2)\}$ is a strong equilibrium. As above, δ_2 dominates δ_1 and δ_0 . Therefore, we must consider $n - 1$ agents playing the strategy $S = (\sigma_2, \delta_2)$, and a_i who 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 uploads will be 0, just under 1 or 2 respectively, and thus his expected financial transfer to the system will be 2π , slightly more than π or 0. Our second assumption tells us that agents prefer to share k files and perform k uploads than to pay the system for k points. Thus a_i maximizes his utility by choosing the action σ_2 , and so Σ is an equilibrium.

This approach does not interfere with download patterns, as the quantized micro-payment mechanism could do. Instead it always gives the right incentives for consumption, since there is no way to make money out of downloads. 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 download demand. The formula to reward distributors thus becomes $\int M(t)\alpha(t)dt$, where $\alpha(t)$ is a scaling factor proportional to demand. This ensures that the files are available at the right times. The problem of unpopular files can be addressed in exactly the same way. Even a rough formula for the expected number of downloads of each file would solve the problem. There is a trade-off between the efficiency and the complexity of the system. A reasonable middle-ground might be to divide files and times of day into a small number of discrete categories.

Another challenging task is to make this mechanism work well with agents' idle time. 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 sporadically, sharing only when their computer is idle. A mechanism that accommodates such behavior is likely to be more successful; 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 Experimental Results

The previous sections analyzed the existence of equilibria for all our mechanisms in simplifying assumptions. The purpose of this section is to test our mechanisms under simulations that more accurately reflect the complexity of the real world. We enrich the theoretical model by introducing different types of files and agents, and by considering risk-averse agents. In this extended model finding an equilibrium becomes even more difficult.

Our agents use a Q-learning algorithm based on the standard paradigm of trading off exploration and exploitation to learn the behavior of other agents.

6.1 Model

We consider several *types of files*, and several *types of agents* defined by their preferences for different types of files. The type of an agent is defined as the percentage of files of each type stored by that agent. For example, one type of agent might want to keep files of type A and files of type B in equal proportion, whereas another type of agent might only be interested in files of type A. We assume that the preferences of agents are fixed throughout the experiment. The utility of an agent is linear both in the number of files downloaded and in the agent’s preference for the type of files downloaded.

We define our agents with parameters chosen according to the following probability distributions:

- **Disk space:** chosen uniformly at random from the interval $[DS_{\min}, DS_{\max}]$.
- **File type preferences:** chosen from a predefined set of weighted combinations of file types.
- **Altruism:** the utility per file uploaded is chosen uniformly at random from the interval $[AL_{\min}, AL_{\max}]$.
- **Utility factor:** this is a scaling factor applied to the utility of downloads. It is uniformly chosen from $[MA_{\min}, MA_{\max}]$. Even agents who have a low utility factor may download files with the intent of making money redistributing them.

In the simulation of micro-payment mechanisms, our agents are stateless. In the simulation of point-based mechanisms, we define states according to the number of points accumulated by an agent. Points have no intrinsic utility value for an agent. However, an agent who runs out of points has to purchase them with money.

All the other parameters of our mechanisms are fixed once and for all and equal for all agents. In particular, we do not model pricing discrimination. For all agents, we take the utility of money to be a logarithmic function. This assumption appears consistent with experimental evidence [5]:

$$U(x) = A \ln\left(1 + \frac{x}{A}\right),$$

where A is a parameter of the function. As A tends to infinity, the function U becomes linear. This allows us to observe changes in the equilibrium as agents go from risk-averse to risk-neutral.

6.2 Learning algorithm

We take an approach similar to that of fictitious play [3] to model the behavior of agents. Agents behave as if the distribution of other agents' strategies were fixed but unknown to them. This distribution may evolve over time, but at any given moment agents select the best action available to them based on the knowledge they have accumulated in previous rounds.

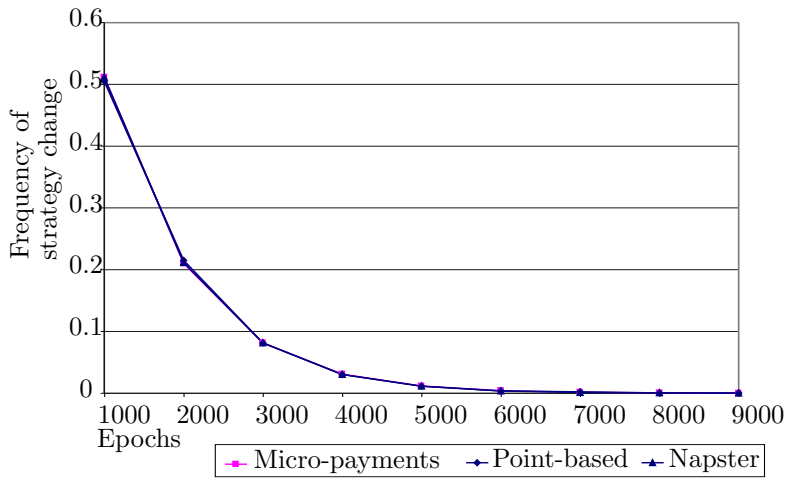


Figure 1: Strategy convergence.

The knowledge that an agent acquires about its environment may come either in the form of a joint distribution of other agents' strategies, as in a fictitious play model, or in the form of 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. We thus focus our attention on learning the payoffs associated with an agent's strategy.

Agents use the Q-learning algorithm to learn the best responses in their environment. This algorithm learns the expected utilities of (state,action)-pairs (called Q-values). The best response is the action that guarantees the highest 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 following update equation for temporal-difference 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.

All agents play a best response to the environment, which they falsely consider to be static. If no agent changes his strategy, then by definition the agents have reached a Nash equilibrium. This is because they have reached a point where each agent's best response is to maintain his strategy, given the assumption that all other agents are fixed in their strategies.

6.3 Results of our Experiments

First and foremost, our simulations confirm the existence of equilibria for the micro-payment and point-based mechanisms, as our analysis predicted (Figure 1 shows the convergence of agents towards an equilibrium).

Our simulation is robust in the sense that we observe essentially the same results under different sets of parameters for the number and types of files and the size of the action space for agents. Agents with a wider choice of actions (more options for downloads and for the proportion of files they share) achieve higher payoffs, but the results remain quantitatively the same. For example, two runs of the same experiment with agents given respectively 9 and 35 actions in their strategy space produced essentially the same result (Fig. 3).




Figure 2: Risk-aversion in micro-payment mechanism.

In our third series of experiments, we studied the influence of risk-aversion on agent's behavior in the micro-payment scheme (Fig. 2) The X -axis is the parameter A of the agents' value of money, the Y -axis is the number of files shared in the system. The lower the parameter A , the more risk-averse agents

are. In accordance with the theoretical analysis, risk-averse agents tend to cut their spending and scale down their contribution to the system.

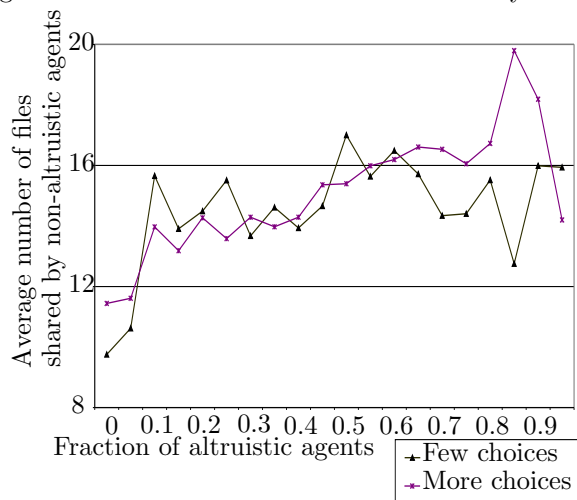


Figure 3: Files shared vs proportion of altruistic agents.

Two experiments show that our model is complex enough to exhibit some non-trivial effects. On Fig. 3 we analyze the behavior of non-altruistic agents in the presence of altruistic agents in the credit-based mechanism. As the proportion of altruistic agents increases from 0 to 1, non-altruistic agents discover that can download more and therefore have to share more to compensate for the point cost of their downloads.

Another experiment⁷ analyzes the behavior of agents in the credit-based mechanism as a function of the price of points that agents have to buy to cover their negative balance. If the price is zero, agents can ignore points and this is effectively the Napster system. As the price goes up, agents start to share more, but as it continues to increase they cut their downloads for fear of overspending.

Acknowledgements

We wish to recognize Mark Lillibridge for his extensive contributions to this paper.

References

- [1] Eytan Adar and Bernardo A. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.

⁷We omit this graph for space reasons.

- [2] N. Asokan. Fairness in electronic commerce. *Ph.D. Thesis, University of Waterloo, Ontario, Canada*, May 1998.
- [3] D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1998.
- [4] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [5] C. Grayson. *Decisions Under Uncertainty : Drilling Decisions by Oil & Gas Operators*. Ayer Company, 1979.
- [6] 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.
- [7] A. M. Odlyzko. The history of communications and its implications for the internet. Submitted to *Computer Networks*, 2000. Available online at <http://www.research.att.com/~amo/doc/networks.html>.
- [8] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [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.