Incentives for Sharing in Peer-to-Peer Networks

write-up by
Georgiana Ifrim*

after a paper by
Philippe Golle, Kevin Leyton-Brown,
Ilyia Mironov, Mark Lillibridge

1 Introduction

Peer-to-peer information sharing environments are undoubtably a well disputed topic of the nowadays Internet. They gave rise to many research questions and in the same time to many hopes in what the success of these kind of systems might mean. They provide an infrastructure in which the desired information can be located and downloaded while preserving the anonymity of both requesters and providers. As recent experience with P2P environments such as Gnutella and Napster shows, together with all the great advantages these kind of systems may bring, we also get many tries of manipulating the networks for different goals.

There are many issues interesting about the dynamics of P2P systems, and in this paper we will refer to the one of file sharing management in such networks. For this, we first consider the fact that a P2P system's users will rationally act as to maximize their benefit from using the system's common resources. Users choose their *strategies* by interacting with the other peers in the network and usually try to manipulate the ratio between what they contribute and what they take from the system. If the mechanisms behind downloading and sharing in the specific P2P system are not designed as to consider these aspects of *selfish* behavior, we are confronted with what is called the *free-riding* phenomenon: peers that take advantage of the network without contributing anything of their own resources.

We have seen this phenomenon in a very large proportion in popular P2P systems, such as Gnutella or Napster where more than half the users just used the system for their own interest, without considering too much the perspective of contributing in any way, and this happened mainly for the reason that none of these systems had a proper (or any) sharing management, that might have imposed some rules. The *legal system* of this kind of P2P networks (if we keep on the line of comparing these networks to the real world) didn't exist at all.

Combining some interesting ideas taken directly from psychological human behavior and economical principles that govern our world, we can try to model the dynamics of P2P file sharing systems by a mapping into the game theory world [3]:

> *Games are characterized by a number of players or decision makers who interact, possibly threaten each other and form coalitions, take actions under uncertain conditions, and finally receive some benefit or reward or possibly some punishment or monetary loss.*

*ifrim@mpi-sb.mpg.de, IMPRS, Saarbruecken, 2004

Now, after we built this bridge between the game theory world and the P2P networks world, based on the resemblance of a game's players to the P2P system's users, we should mention what specifically about game theory we can use, and that is the notion of *Nash equilibrium*.

Why this notion of equilibrium? Because we want our users to *converge* to that moment in time where everybody is pleased with the actions he is taking (the balance between gains and costs is considered 'fair' by the user), enforcing in the same time a flavor of correctness: they have to put some effort into the common welfare. In this ideal situation, everybody should be happy and we could also overcome the well disputed free-rider problem. This goal, already raises some questions: if we inspire our model from an imperfect human society, where there are winners and losers, can we achieve such an ideal goal? If the individual components are selfish can we somehow get *good aggregate behavior*?

The mechanisms proposed by this paper try to give some solutions to these questions and some of the arguments that we will see are very robust and theoretically clear. Still the unpredictable factors of human behavior have to be taken into account.

Thus, the following section will be concerned with defining the formal game theoretic setup and providing some basic definitions from game theory necessary for a basic understanding of the proofs given. Then, in section 3, we will look of how can we use this game theoretic model in order to analyze sharing in a P2P system such as Napster. In section 4 and 5 we introduce two classes of incentive mechanisms for managing sharing in P2P networks, in the game theoretic framework that we designed. Section 6 is devoted to the validation of the theoretical results with some in-sites on further properties of the mechanisms proposed.

Increasing the proportion of users that share files, making a great variety of files available, will also result in increasing the value of the network to its users, making it more robust in competing with other commercial P2P systems.

2 Defining the game theoretic model

We model the file sharing scenario during one time period (e.g. one month) by a non-cooperative game among *rational* and *strategic* players (for further details about game theoretical concepts, the reader can consult [3]):

- We consider n *agents* (peers) to participate in the system; we denote them by: $a_1 \dots a_n$.
- Each agent has a number of possible *strategies* $S_i = (\sigma, \delta)$, where σ stands for the action of *sharing* and δ for the action of *downloading*.
- Each strategy consists of two orthogonal *actions*:
  1. **Sharing**: agents select what proportion of files to share. In our model, sharing takes three levels: σ_0 (none), σ_1 (moderate), σ_2 (heavy). (There are some analysis also for more fine grained action spaces, but for a better understanding of the theoretical proofs, we will only deal with these three levels. A more detailed discussion of these aspects is done in the experiments section).
  2. **Downloading**: agents also should select how much to download from the network in each period. The action space for downloading spans also three levels: δ_0 (none), δ_1 (moderate), δ_2 (heavy).
- The strategies chosen determine the *outcome*.[1]
- Associated with each outcome is a collection of *payoffs*, one to each agent.

[1] This *outcome* can be seen as an output to a set of given inputs.

2.1 Agent Utility

The *utility* of an agent describes the balance between his gains and his costs determined by the participation in the P2P system. If we quantify agent's preferences, we can say that an agent's utility function describes his preferences for outcomes. In our model these preferences are somehow quantified by the following factors:

- **Amount Downloaded (AD):** Agents get happier the more they download.
- **Network Variety (NV):** Agents prefer to have a larger diversity of files from where to choose, in the network.
- **Altruism (AL):** Agents might derive some positive utility from the satisfaction of contributing to the network. This is somehow an abstract quantifier, but it tries to capture exactly the unpredictability of human behavior. This factor will be used in the analysis of Napster file sharing. It will be disregarded from the analysis of the proposed mechanisms, being considered not enough a motivation and too difficult to quantify.
- **Disk Space Used (DS):** There is a cost to agents associated with allocating disk space to files to be shared.
- **Bandwidth Used (BW):** Uploading files to the network also implies a cost for the agents, in terms of the bandwidth used.
- **Financial Transfer (FT):** Agents have to pay for their usage of network's resources or are paid by the system for their contribution, depending on the actions they have chosen.

In the model proposed we assume that agents have a *quasi-linear utility function*, meaning that each agent's utility function can be written as a sum of arbitrary functions. Each of this functions maps the amount expressed by one of the variables mentioned above, into a money quantity. Another assumption about the participating agents, is that they are *risk neutral* and by this we mean that they will not cut down their spendings during their participation to the P2P system, in order to avoid paying too much for the resources they have used. This is a way of saying that in our model, we consider our users to be motivated enough (have enough interest in the system resources they might use) as to have an incentive to download maximally in almost all circumstances (if they are also able to cover their spendings, by contributing). We consider them willing to invest in their participation. The opposite of *risk neutral* agents are the *risk averse* agents. We will deal with this kind of agents in the experiments section. A simple example for a better understanding of the concept *risk averse* is taken from [9]:

One example of risk aversion can be seen on Game Shows. For example, if a person has a 1 in 3 chance of winning $50,000, or can take a sure $10,000, many people will take the sure $10,000.

Thus, we can express an agent a_i's utility function as the following sum:

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

Each f - function corresponds to a particular agent and to a specific variable, and it describes that agent's preference for different values of the variable, in money.

The writing of the utility function is meant to be somehow intuitive: is the balance between the gain (described by the *positive factors*: AD, NV, AL) and the cost (described by the factors that bring negative utility to the agent - *negative factors*: DS, BW). FT is considered to be an amount of money *'positive' or negative'* that is transfered from or to the agent (an agent pays the system or is getting paid by the system). The authors also used some other writings of the utility function in previous versions of this paper, but they decided to stay with the more intuitive one.

The functions associated with AD, NV, AL can be thought of as monotonically increasing functions, with minimal value 0, because these factors only contribute positive utility. Also, DS and BW, only contribute negative utility, explaining the subtraction of f^{DS}, f^{BW}.

Also we can say something more about the relationship between the actions available to an agent and its utility function. Sharing files has a negative impact on an agent's utility, since his utility decreases with the amount of disc space (DS) used for sharing. Furthermore, uploads keep increasing the costs 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 gets for sharing files is due to altruism, all other agents receive increased utility due to the increase in the network variety (NV). Agents can increase their own utility by increasing their downloads (AD), as long as $NV \neq 0$ (i.e. files are being shared). Last, the designer of the peer-to-peer system can charge or credit agents an amount of money (FT) that depends on their actions.

Before going further into the design and analysis of our proposed incentives schemes, we make two more assumptions about agents relative preferences for outcomes:

$$f^{AD}(k) > k \cdot \beta \tag{1}$$
$$f^{DS}(k) + f^{BW}(k) < k \cdot \beta \tag{2}$$

Equation (1) states that the positive utility derived by an agent from downloading k files, should exceed what he has to pay the system for downloading, where β in this case is the cost per file downloaded. This can be again considered to be an incentive for downloading, and is somehow natural because the user is free to choose the files that he will download, files that are important to him and that he considers worthy for the corresponding amount of money. In the case of *heterogeneous peers*, this cost per file β, should vary, there should be different prices for different files (if we want this assumption to hold). The authors don't comment on the choice of this cost β, they use it as a general cost per unit downloaded in the case of a *homogeneous* environment.

The second assumption comes to complete the first one: agents would like their cost of sharing k files and uploading k files to be less than the payment they get from the system for contributing resources; so the credit he gets for contributing those resources should worth the effort. In this case β is considered to be the reward per file uploaded and must be carefully chosen when managing heterogeneous peers.

We often assume our agents to have the same type (same utility function).

2.2 Equilibria

As a central issue for our game theoretic model we considered our agents to be *economically rational*: they act as 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 the n agents by: $\Sigma = \{S_1, \ldots, S_n\}$. We provide some basic game theory definitions that will be needed for developing our analysis. We say that:

1. Σ is a *weak Nash equilibrium* if no agent can gain by changing his strategy given that all other agents' strategies are fixed.

2. 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.

3. An agent has a *dominant strategy* if his best action does not depend on the action of any other agent.

2.3 General Observations

In the design and analysis of the incentive mechanisms, we assume having a central server that maintains a database of the files currently available in the system, matches download requests with available clients, keeps track of the number of download/uploads per user. The way the mechanisms will present are designed make necessary the centralized approach.

There are also some issues of copyright involved, file identification etc. (readers interested in this aspect can directly consult [1]). The mechanisms proposed are compatible with the addition of a flat-rate membership fee, even if the authors don't explicitly consider this option.

3 Motivating example: The Napster System

We want to analyze sharing in the Napster System, using the game theoretic model that we defined. Napster is one of the most popular example of P2P systems, designed for exchanging music files among its users. Because of some legal issues, it operated only from May 1999 through July 2001.

The main characteristic of this centralized P2P system from our game theoretic point of view is the fact that no matter what actions our users took, the system didn't impose any financial transfers ($FT = 0$, in our model). We will first analyze the dynamics of sharing in Napster, by disregarding the altruism (AL) component of agents' utility functions. Then we will also take this factor into account.

Therefore, we are concerned with the utility function:

$$U = [f^{AD}(AD) + f^{NV}(NV)] - [f^{DS}(DS) + f^{BW}(BW)]$$

We first show that, using the assumption above, $\Sigma = \{(\sigma_0, \delta_2), \dots (\sigma_0, \delta_2)\}$ is an equilibrium. We assume that all our agents have the same type, so its enough to analyze the choices made by a single agent. We fix an agent a_i and analyze his behavior in relation to the other peers' strategies. We consider the other $n-1$ agents to play the strategy $S = (\sigma_0, \delta_2)$ and analyze agent a_i's best response. We disregard altruism from our analysis, so the agent a_i has no incentive to share, his utility will be strictly decreased by sharing files (see factor DS). If he chooses the action σ_0, his utility will be unchanged. Thus he will choose as a sharing action σ_0: no share. Downloading will usually increase a_i's utility, but when no other agent shares ($NV = 0$), his utility will be zero, no matter 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 because the 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 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 utility than any other strategy. Because σ is an equilibrium in dominant strategies, it is unique.(This a game theory result, for more information on this theoretical results, please consult [3]).

Thus, we identified an equilibrium in which nothing is shared and there is nothing to download. Still, songs were successfully traded on Napster. Some possible answers to the question of how this system worked after all, could be that Napster offered its service free of charge and tried to foster a sense of community among its users, by such features as chat-rooms, a newsletter and messaging between users. These might be some incentives that made users *altruistically* contribute resources that maybe cost them very little. Second, Napster offered a modest *disincentive for non-contributing*: by default, the Napster client shared all songs (files), than an agent downloaded. This could only be avoided by manually moving files to another directory or specifically 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 don't bother to 'opt out'.

In our model, we represent both these incentives for contributing, by the *altruism* variable AL. We use this factor for representing 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. Here we suppose of course that the utility gained from the increase in network variety and possibly amount downloaded (if everybody shares, you have more to download from), will be somehow above the cost the default sharing option implies (by disk space used and bandwidth used - DS, BW factors). Again, we try to explain the possible ways of human behavior, and of course, we cannot get any guarantees that our assumptions are right.

We now consider the equilibrium that will arise when some agents are altruistic. We consider to model two types of agents: *altruistic agents* - those whose 'reward' for altruistic behavior (AL) exceeds his cost in terms of disk space (DS) and expected bandwidth usage (BW). We assume that f functions for these agents are such hat they would prefer the action σ_2 to either the action σ_1 or σ_0 regardless of the values of BW. These agents still gain utility from downloads due to the AD factor, so (σ_2, δ_2) will be a dominant strategy for altruistic agents. The second type of agents are those for whom the cost for altruistic behavior exceeds its benefit. These agents are essentially the same as those described when we disregarded altruism: they may receive some utility for altruistic behavior, but enough as to change their behavior. So, this type of actions will have as a dominant strategy: (σ_0, δ_2).

This might be a description of how the process of exchanging files worked in Napster: some proportion of agents were sufficiently altruistic as to share files, and these is how resources still existed in the network, and some of them had no incentive to share, so they didn't contribute anything. Regardless of their level of contribution (altruism), they were not restricted in any way in their downloads. This analysis supports the empirical research that the majority of Napster users were free-riders, they didn't share anything at all. The authors make here the observation that this situation is somehow paradoxical: 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), then they are under the unique equilibrium (no sharing, maximally downloading), but this comes again to the ideal perfect society where all people will behave nice and this in reality doesn't happen.

We will discuss in the following sections (3 and 4), about several alternative mechanisms that [1] 'overcome the free-rider problem through the imposition of financial transfers'. In order to avoid relying on altruism, the design of these mechanisms assumes that agents have no altruistic motivation, and so we will drop the term altruism (AL) from the utility equation. Will consider again this factor in the experiments section.

4 Micro-Payment Mechanisms

We would like users to both use and contribute to the system's resources, so this is the goal of this schema. One somehow natural approach to this problem would be to charge users for downloads and reward them for uploads. In this section we will describe the micro-payment mechanism that is designed using this notion of financial transfer between the user and the system. In the next section we'll look to another mechanism: quantized micro-payments, which is developed from this one.

We will first discuss in detail how exactly this mechanism based on micro-payments works: the central server tracks the number of downloads and uploads per user, for a given period of time; let d denote the number of files downloaded and u the number of files uploaded by some user. The server is aware of all such transfers since it processes all download requests. Here we can involve some issues concerning protocols to ensure that both parties agree on whether their transaction was aborted or ended successfully (the reader can consult [1] for more details concerning fair exchange).

At the end of each period, each user is charged an amount $C = f(d-u)$. The function f maps the difference between downloads and uploads to a financial transfer from agents to the system. We assume that f is linear with a coefficient β representing the cost/reward per file (e.g. \$0.5), and that the value of this coefficient is such that inequalities (1) and (2) hold.

Note that this mechanism has the property that the global sum of all micro-payments is zero. Because underlying this class of mechanisms there are financial issues, we have to consider what are the chances that some users take advantage of the network, manipulating the mechanism such that they may make a profit out of this participation in the network. This might happen when some users download less then they upload (in a heterogeneous environment). From a network administrator or investor point of view, this is not very convenient.

What follows next is the analysis of the equilibria that arise in this kind of mechanisms. For a more clear discussion of the proofs, we will make some simplifying assumptions.

Let σ^{-i} be the total number of units shared by agents other than a_i, and δ^{-i} be the total units downloaded by agents other than a_i. If agent a_i chooses the strategy (actions) $S = (\sigma_s, \delta_d)$[2], then his expected payment to the system (expected value of FT) can be expressed as:

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

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. In this context β was considered to be the cost per net unit downloaded. We supposed agent a_i shares s units and downloads d units, and we established that an user is charged an amount linear in the difference between the number of downloads and the number of uploads. Agent a_i directly controls his number of downloads d, but indirectly controls his number of uploads by the number of units shared. So in calculating the expectation of FT, d is fixed, but the number of uploads made by a_i is a random variable and we need to calculate its expectation.

We have to prove that the expected number of uploads made by a_i is:

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

By a simple intuition, if we consider that each of the other $n-1$ agents shares an equal proportion of the files $\frac{\sigma^{-i}}{n-1}$, than the probability that one of the $n-1$ agents downloads from a_i (a_i uploads) is the ratio between the number of units a_i shares (s units) at that moment and the total number of units shared in the system, which is exactly $\frac{n-2}{n-1}\sigma^{-i} + s$

[2]Here σ_s has the meaning of sharing s units, δ_d means downloading d units; these notations were used for the ease of exposition

because we assumed no agent will download from himself (this is only a counting problem). We considered the probability that a_i uploads to an arbitrary agent. Now for calculating the *expected number* of a_i's uploads is enough to consider this probability over all the units downloaded in the system, by the $n-1$ agents other than a_i (we also consider the assumption that downloaders are matched at random with shared units). This results exactly in:

$$E[u] = \delta^{-i} \cdot \frac{s}{\frac{n-2}{n-1}\sigma^{-i} + s}. \tag{5}$$

It follows the discussion concerning the equilibria that arise in this mechanism and the proof of our arguments:

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

Proof[3]:
From the assumption (1): $f^{AD}(k) > k\beta$, we see that 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 in this context:

$$E[FT] = \beta \cdot \left(2 - 2(n-1) \cdot \frac{s}{2(n-2) + s}\right). \tag{6}$$

By analyzing the relation between the number s of units shared by a_i and the expected number of uploads, considering also the assumption made about the cost of sharing and uploading a file: (2) $f^{DS}(k) + f^{BW}(k) < k\beta$, we get that the best action agent a_i can take is σ_2. Thus a_i will strictly prefer the strategy: $S_i = (\sigma_2, \delta_2)$.

Thus Σ s a strict equilibrium. We also have to show that it is unique. We have seen that its dominant for all agents to choose δ_2. Thus, δ^{-i} is $2n-2$ in all equilibria, 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 action σ_2 when $2(n-1) \cdot \frac{s}{\frac{n-2}{n-1}\sigma^{-i}+s} \geq 2$. This results in σ_2 being the most profitable strategy as long as $\sigma^{-i} \leq 2(n-1)$. This condition holds always because there are $n-1$ agents other than a_i and each agent can only share up to 2 units, so Σ is unique. $\square$

A possible explanation of why this argument (proof) works[4] is the fact that [9] in a game, agents strive *to maximize their (expected) utility* by choosing particular courses of action. This comes from the central assumption in game theory, that agents are *economically rational*. By choosing as strategy, no share, no download, the utility of all agents will be zero (all the factors composing the utility - AD, NV, DS, BW, FT are zero).

Still, this analysis works only with the assumption that all agents are risk neutral. 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 of

[3]This proof explicitly uses the fact that there are only 3 levels of sharing and of downloading possible; this restriction was done only for the ease of exposition. The authors sustain there exists also a more complex proof for any number of levels of sharing and downloading

[4]Why shouldn't we consider the possibility of an equilibrium $\Sigma = \{(\sigma_0, \delta_0) \ldots (\sigma_0, \delta_0\}$ in which nobody shares and nobody downloads nothing

agents' risk aversion and their particular utility functions, they may prefer to reduce their number of downloads to reduce their worst-case payment to the system. The authors consider this issue too difficult to analyze formally, and come back to this problem in the experiments section.

4.1 Quantized Micro-Payment Mechanisms

These type of mechanism were proposed as a solution to the empirically result that users strongly dislike micro-payments, the fact that they have to decide before each download if a file worths the money (mental decision costs). There is an empirical established belief that users often may prefer flat pricing plans, even when such mechanisms may increase their expected costs. So for overcoming this possible drawback, we will introduce *quantized micro-payment mechanisms*. The schema is based on the following concept: users pay for downloads in block of b files, where b is a fixed parameter[5] and are paid for uploads in the same manner as in the previous presented mechanism. After 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 that number of blocks. As b increases, we go from the micro-payments mechanism to a flat rate pricing schema.

The authors do not present an analysis of this class of mechanisms and justify this fact by the complexity of the formal analysis. In the abstract these quantized mechanisms are a sub-family of the general micro-payments schemes, so the original equilibrium is preserved: sharing and downloading maximally. The key advantage of this type of mechanisms is considered to be the fact that agents are spared the mental decision costs of per-download pricing.

Still, there exist an analysis of some practical issues. We come again to observing ways in which users might gain through conspiracy and we also see some proposed solutions for these situations.

One case when users might gain by manipulating the system sharing management is based on the following observation: after one file has been downloaded, the remaining $b-1$ files belonging to the same block, have zero-marginal cost. Towards the end of a payment period, users may take advantage of zero-marginal cost downloads left in their account and use these to credit their friends with uploads (downloading from friends will result in a payment these friends get for uploading files to the network). The proposed solutions to this situation would be that the server should reply to download requests by hiding the identities of users or by showing a subset of users randomly chosen.

Another problem approached here is that of users sharing rare files[6]. These users might be advantaged because they will receive the most number of download requests (uploads). The solution proposed is to treat rare files differently from files that are more frequent, mainly by retroactively crediting these files, after they became frequent enough. We have to remember the reader, that even some issues approached in this paper may seem a little strange from the user point of view, we are dealing with mechanisms based on financial principles, so the notion of 'altruism' from the part of the 'investor' is not very likely[7].

5 Points-Based Mechanisms

The first class of mechanisms presented in the previous section, charged users for downloads and rewarded them for uploads. Here we will see a different approach: the users will continue

[5] In practice, b is chosen in the order of the number of files downloaded by an average user, per time period

[6] By rare files we mean files for which the number of copies available falls below a threshold.

[7] These mechanisms are also intended for possible use in commercial systems, this might be a reason for some of the non-users profit oriented analysis done in this paper

to be penalized for downloads, but they will be credited in proportion to the size of material shared, rather than the number of uploads they make. Also this mechanism makes use of an internal currency: 'points', that can be bought with money or contribution to the network, but cannot be sold back. Agents payment for sharing is $\int M(t)dt$, 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. Downloading a file costs $c \cdot m$ points, where m is the size of the file in megabytes and c is a system's constant parameter. c may be taken to be the number of hours a newly downloaded file has to be shared in order to waive its download cost [2]. Again, we must simplify the mechanism in order to analyze it according to our game theoretic model. We assume that one points costs β, where β is chosen such that inequalities (1) and (2) hold. Furthermore, we will asume that all files have the same size (1 MB) and that all agents share 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 2 points). We take $c = 1$, so each level of downloading cost one point. Downloaded files are not shared in the same time as they were downloaded and we also use the assumptions that downloaders are matched uniformly at random with share units and no agent will download from himself. Thus, if we analyze the strategy of a fixed agent a_i, supposing that a_i chooses to share s units, his expected number of uploads[8] is:

$$E[u_i] = \delta^{-i} \cdot \frac{s}{\frac{n-2}{n-1}\sigma^{-i} + s}. \tag{7}$$

Next follows the analysis of the equilibria in this mechanism.

Proposition 2 $\Sigma = \{(\sigma_2, \delta_2) \ldots (\sigma_2, \delta_2)\}$ *is a strict equilibrium.*[9]

Proof:
Why do we need to analyze the expected number of uploads if we consider to penalize downloads and reward sharing? Because of the assumption that *downloaded files are not shared in the same time period as they were downloaded*. So, if an agent a_i shares k files, and other agents download some of those files, then we consider a_i to 'loose' the files uploaded to the network. So we need to analyze what is the expected number of uploads for a given amount of material shared, in order to analyze an agent's expected utility, given the other agents' actions.

We consider again the situation of the $n-1$ agents playing the strategy $S = (\sigma_2, \delta_2)$ and an agent a_i must determine his best response. From inequality (1) $f^{AD}(k) > k\beta$, we get that δ_2 is a dominant strategy. Thus, a_i will play $S = (\sigma_s, \delta_2)$ and must choose a value for s. Again, by σ_s here we mean the action of sharing s units. We analyze the expected financial transfer to the system given the agent's preferences and the other agents' actions. If a_i plays σ_0, σ_1 or σ_2 his expected number of uploads (given the other agents' strategies) will be respectively 0, just under 1 or 2. According to the level of sharing he chooses, his expected financial transfer to the system will be 2β, slightly more than β or 0. From the assumption (2) $f^{DS}(k) + f^{BW}(k) < k\beta$ we see that agents prefer to share k files and upload k files than to pay the system for k points[10]. Since sharing at level 2 leads to uploading at level 2 on expectation, given the other agents' strategies, a_i maximizes his expected utility by the action σ_2. Thus Σ (sharing and downloading maximally) is a strict equilibrium. □

[8] We denote the number of a_i's uploads by u_i

[9] We will see that unfortunately is not unique.

[10] In order to be able to download, first you need to have the points currency, gained by sharing or directly with money.

But, unfortunately, is not a unique equilibrium. Generally, point-based schemes have the drawback that they give rise to a degenerate equilibrium in which nobody shares nothing, but all agents want to download maximally. This is triggered by the fact that points have no value if they are not repurchased for paying downloads.

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

We will not give the proof of this proposition here, but only an intuition.[11] If no agent shares, but all want to download heavily, gaining points will yield no utility if there is nothing to download.

This is indeed a drawback of this class of mechanisms, it is not clear which equilibrium will be reached in play. There are some tries to provide answers for this, in the experimental section.

In the tradition of the previous section, the authors analyze opportunities for collusion or for any other kind of exploitation of the system, and also provide some solutions.

The possibility of users conspiracies in order to simulate 'fake' uploads doesn't make sense in this points based schema. But another problem arises, users are now paid for the amount of shared files, but have a cost implied by uploading files to the network, by the consumption of bandwidth. For this reason they might share unpopular files or share at low-usage times, resulting in a general network value alteration. A possible solution proposed by the authors is to scale the reward for sharing measure by a factor proportional to the expected demand, denoted here by $\lambda(t)$:

$$\int M(t)\lambda(t)dt \tag{8}$$

This might ensure that the files are available at the 'right times'. The problem of users preferring to share unpopular files can also be addressed in a similar way.

There are still some questions unanswered as concerning the fact that P2P networks are dynamic systems, users go and come, we cannot expect them to share files for hours, so some mechanisms that will accommodate such behavior will be more useful.

6 Experimental Results

The previous sections formally analyzed the proposed mechanisms under simplifying assumptions: homogeneous environment (all agents have the same type), only a restricted space of actions for users (three levels of sharing and downloading). In this section we will look at some experiments done in a more complex setup: we consider now different types of agents and different types of files, we take risk aversion into consideration and also the altruism factor. The purpose of this section is to test the proposed mechanisms under simulations that reflect more accurately the complexity of the real world. The experiments use a Q-learning algorithm based on a paradigm of trading *exploration* and *exploitation* to learn the actions that maximize an agent's expected payoff, given the behavior of other agents.

6.1 Experimental Setup

The model presented in section 2 is extended by considering action spaces for agents more fine-grained than the three levels of sharing and downloading; here we consider to analyze different types of agents and different types of files. In the experiments presented, agents differ according to their (fixed) preferences for different kinds of files. The agents utility function differ according to parameters chosen from the following probability distributions:

[11] Readers willing to see a formal proof may consult directly the paper [1].

- **Disk space:** chosen uniformly at random from the interval $[DS_{min}, DS_{max}]$.
- **Altruism:** the utility per file uploaded is chosen uniformly at random from the interval $[Al_{min}, AL_{max}]$.
- **File type preferences:** chosen from a predefined set of weighted combination of file types.
- **Utility factor:** this is a scaling factor applied for utility of downloads. It is uniformly chosen from $[MA_{min}, MA_{max}]$.

The authors make the remark that even agents who have a low utility factor may download files with the intent of making money redistributing them.

All the other mechanisms' parameters are considered fixed and equal for all agents. In the simulation of micro-payment mechanisms, the agents are considered to be stateless, by this meaning that they don't keep track of the amount of money they spend or make. In the points-based mechanisms, agents have states corresponding to the number of points. Points don't have any intrinsic value to agents, but if an agent runs out of points, he must purchase more with money.

For all agents, the utility for money is considered to be a logarithmic function: $U(x) = A \cdot ln\left(1 + \frac{x}{A}\right)$. This *model of risk aversion*, as is called by the authors, is taken from empirical studies. A is a parameter defined as the agent's value for money (how much is he willing to invest in his participation to the system). As A increases towards infinity, the utility for money U, becomes linear: this captures the changes in behavior as agents go from risk aversion (A small) to risk neutrality.

6.2 Learning algorithm

In simulating the learning process, agents behave as if other agents' strategies were fixed, but unknown to them (like the other agents do not act strategically), and choose the action that results in the highest possible payoff based on the knowledge they have accumulated in the previous rounds.

Thus 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.

For realizing the experiments we mentioned that a *multi-agent reinforcement learning* paradigm was used. First, we should give the unfamiliar reader some basic notions of what is reinforcement learning about and one clear explanation is given in [7]:

> *Reinforcement learning addresses the question of how an autonomous agent that senses and acts in its environment can learn to choose optimal actions to achieve its goals. ...Each time the agent performs an action in its environment, a trainer may provide a reward or penalty to indicate the desirability of the resulting state.*

In a multi-agent reinforcement learning paradigm, an agent could attempt to learn either the joint distribution of other agents' strategies or the expected payoffs associated with its own strategies, without modelling the other agents. As the number of participants in a P2P system is usually very large, the authors prefer the second alternative for their experiments.

For simulating the learning process a temporal difference (TD) Q-learning algorithm was used. Q learning is one form of reinforcement learning in which agent learn an evaluation function over states and actions. The Q learning algorithm has the advantage that it can be employed even when the learner has no prior knowledge of how its actions affect its environment [7].

By this algorithm agents learn their best actions in their environment, the actions that maximize their expected payoff. *Temporal difference algorithms* learn by iteratively reducing the discrepancies between the estimates produced by the agent at different times [7].

The Q-learning algorithm assumes that the environment doesn't evolve over time, but the learning equation enables agents to also do well in a slowly changing environment.[12]

6.3 Experiments

The simulations confirm the existence of equilibria for the micro-payments and points-based mechanisms in the more complex setup. The strategy convergence is shown for both mechanisms. The plots are represented on a logarithmic scale.

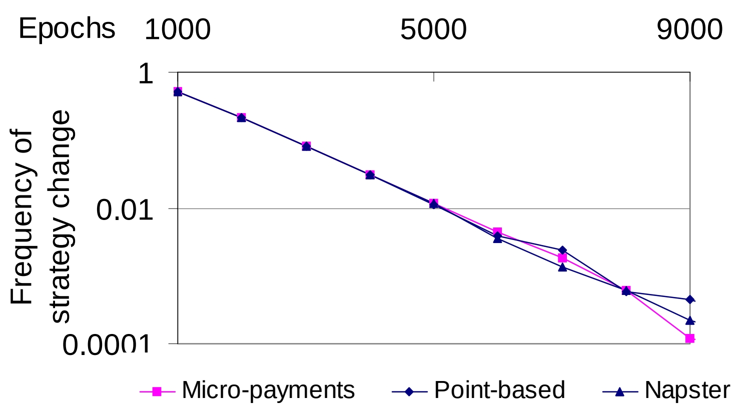

Figure 1: Strategy convergence (logarithmic scale).

There are also some other experiments presented in this section, meant to simulate the behavior of agents under diverse setups. Fig. 2 shows an experiment about the behavior of non-altruistic agents in the presence of altruistic agents, in the points-based mechanism. In this kind of mechanisms you need points currency in order to be able to download files. As the proportion of altruistic agents increases, non-altruistic agents discover they have more files to download and they increase their sharing level in order to gain the currency for downloading.

The robustness of the simulation was analyzed for different sets of parameters for the number and types of files and for the size of the action space for agents.

The results observed from these settings were qualitatively the same: in fig. 2 we can see two runs of the experiment with agents given 9 and 35 actions in their strategy space.
Also the paper presents an experiment concerning the level of agents' risk aversion modeled in the micro-payments mechanisms. For various values of the A parameter, the sum that users want to invest in their participation, we can observe changes in the agent behavior. For smaller values of A we get greater values of risk aversion, and we can observe the number of files shared as a function of the agents' risk aversion. As agent become more risk averse they tend to cut their spendings and scale down their contribution to the network, because of their uncertainty related to the fact that they don't directly control their number of uploads.

[12]For more information on the standard equation of TD Q-learning the interested reader can consult [7].

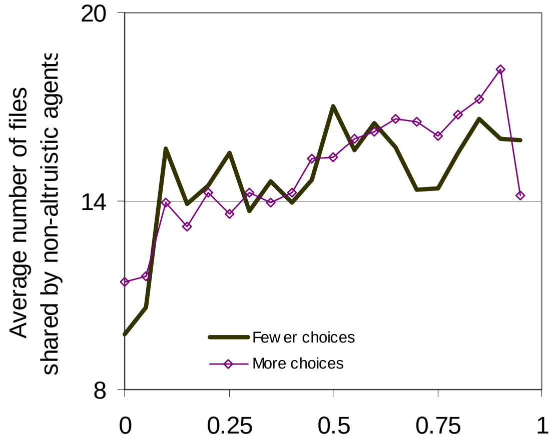

Figure 2: Files shared as a function of the proportion of altruistic agents.

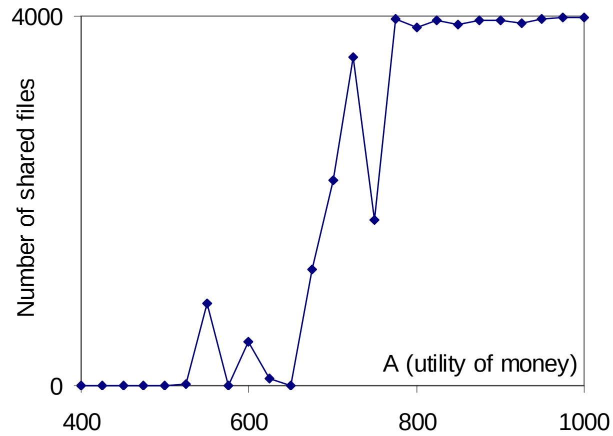

Figure 3: Risk-aversion in micro-payment mechanism.

7 Conclusions

We have seen two classes of mechanisms proposed as a solution to the free-rider problem. These schemas are meant for managing sharing in centralized P2P networks. Users are not altruistic by nature and they will try to take advantage of the systems' resources if there is a way to do that. The formal game theoretic model proposed by the authors has the quality of capturing some of the main features of sharing dynamics in P2P systems, and also allow for a clear theoretical analysis of the equilibria that arise in the mechanisms proposed.
The authors have also analyzed sharing in Napster and gave some explanation about the possible state of actions in this well known centralized P2P system. Even if is very difficult to capture all human behavior's features in the formal model, we also have seen some experiments simulating agents' behavior in possible real life situations.

This paper gives a clear theoretic model of a simplified world of P2P file sharing systems. There still remain some unanswered questions about what the behavior of users will be in a practical, real life situation, that implies unpredictable players' behaviors, triggered by the complexity of human beings. For the formal model we have seen that some simplifying assumptions were done as regarding agents reaction to their environment, but each of this assumptions may not hold in a real P2P system. There is also the problem that P2P systems are by nature dynamic systems and we have only seen models of a static world.

Beyond all these possible drawbacks, this paper came with some interesting ideas applied directly from the world of game theory to the one of P2P networks. Some of them might be used in conceiving file sharing paradigms for commercial P2P systems, some may be represent and they actually do, sources of inspiration for designing better sharing schemas.[13]

[13] A large part of the literature concerned with modelling sharing in P2P systems was inspired from this paper. For details on this the reader may consult [4], [5].

References

[1] P. Golle, K. Leyton-Brown, I. Mironov, M. Lillibridge, **Incentives for Sharing in Peer-to-Peer Networks**, Proc. of the ACM Conference on Electronic Commerce, 2001.

[2] P. Golle, K. Leyton-Brown, I. Mironov, **Incentives for Sharing in Peer-to-Peer Networks**, Proc. of Electronic Commerce'01, 2001.

[3] T. Ferguson, **Game Theory Text**,

`http://www.math.ucla.edu/~tom/Game_Theory/Contents.html`

[4] C. Buragohain, D. Agrawal, S. Suri, **A Game Theoretic Framework for Incentives in P2P Systems**, IEEE P2P, 2003.

[5] S. Kamvar, M. Schlosser, H. Garcia-Molina, **Incentives from Combatting Freeriding on P2P Networks**, Euro-Par, 2003.

[6] K. Ranganathan, M. Ripeanu, A. Sarin, I. Foster, **'To Share or not to Share': An Analysis of Incentives to Contribute in Collaborative File Sharing Environments**, Workshop on Economics of Peer-to-Peer systems, 2003.

[7] T. Mitchell, **Machine Learning**, McGraw-Hill International Editions, 1997.

[8] S. Săroiu, P. K. Gummandi, S. D. Gribble, **A Measurement Study of Peer-to-Peer File Sharing Systems**, Proc. of Multimedia Computing and Networking, 2002.

[9] `http://en.wikipedia.org/wiki/Game_theory`