

A Peer-to-Peer Cooperation Enhancement Scheme and its Performance Analysis

Wei-Cherng Liao, Fragkiskos Papadopoulos, Konstantinos Psounis
 Department of Electrical Engineering, University of Southern California, Los Angeles, USA
 Email: {weicherl,fpapadop,kpsounis}@usc.edu

Abstract—The performance of peer-to-peer systems depends on the level of cooperation of the system’s participants. While most existing peer-to-peer architectures have assumed that users are generally cooperative, there is great evidence from widely deployed systems suggesting the opposite. To date, many schemes have been proposed to alleviate this problem. However, the majority of these schemes are either too complex to use in practice, or do not provide strong enough incentives for cooperation.

In this work we propose a scheme based on the general idea that offering uploads brings revenue to a user, and performing downloads has a cost. We also introduce a theoretical model that predicts the performance of the system and computes the values of the scheme’s parameters that achieve a desired performance. Our scheme is quite simple and very easy to implement. At the same time, it provides very strong incentives for cooperation and improves the performance of P2P networks significantly. In particular, theory and realistic simulations show that it reduces the query response times and file download delays by one order of magnitude, and doubles the system’s throughput.

Keywords: P2P networks, user cooperation, theoretical analysis, realistic simulations.

I. INTRODUCTION

Peer-to-peer (P2P) systems provide a powerful infrastructure for large-scale distributed applications, such as file sharing. While cooperation among the system’s participants is a key element to achieve good performance, there has been growing evidence from widely deployed systems that peers are usually not cooperative. For example, a well known study of the Gnutella file sharing system in 2000 reveals that almost 70% of all peers only consume resources (download files), without providing any files to the system at all [1]. This phenomenon is called “free-riding”.

Despite the fact that this phenomenon was identified several years ago, recent studies of P2P systems show that the percentage of free-riders has significantly increased [2]. This is not because industry and academia have ignored the problem. There is a large body of work on incentive mechanisms for P2P networks, varying from centralized and decentralized credit-based mechanisms, *e.g.* [3], [4], [5], [6], to game-theoretic approaches and

utility-based schemes, *e.g.* [7], [8], to schemes that attempt to identify and/or penalize free-riders, *e.g.* [9], [10], [11], the last two being proposed by the popular KaZaA and eMule systems. The problem of free-riders is hard to tackle because the solution has to satisfy conflicting requirements: minimal overhead, ease of use, and at the same time good amount of fairness and resilience to hacking.

In this paper we propose and study the performance of an efficient algorithm that is very easy to use, it enforces users to be fair, and it can be implemented in a number of ways that tradeoff overhead and resilience to malicious users. According to the algorithm, users use tokens as a means to trade bytes within the system. A user earns K_{up} tokens for each byte he/she uploads to the system and spends K_{down} tokens for each byte he/she downloads from the system. The user also gains K_{on} tokens for each second his/her machine is on the system (*i.e.* it is online). A user can initiate a download only if the number of tokens that he/she has is large enough to download the complete file.

The proposed algorithm relies on the general idea that users should be awarded for offering uploads and staying online, and pay for performing downloads. While others have proposed solutions that use the same general idea in the past, *e.g.* [6], [8], there are a number of questions that either have not been addressed at all or have been studied via simulations only: (i) How should one tune the parameters that dictate the gain from uploads and the cost of downloads? Specifically to our scheme, what is the right value for the parameters K_{on} , K_{up} , K_{down} ? (ii) What is the exact effect of such an algorithm on overall system performance over a wide range of conditions? (iii) Would a *small* number of malicious users, that manage to subvert the scheme, degrade overall performance noticeably? (iv) Is it possible to trade off one performance metric for another by varying the parameters of the algorithm, *e.g.* trade off download delay for total system capacity? Our theoretical analysis of the performance of the resulting system, coupled with extensive realistic simulations, gives concrete answers to all these questions. Interestingly enough, it shows that the query response times and file download delays can be reduced by one order of magnitude while being able to sustain higher user download demands.

An important aspect of any solution to the free-riding problem is if the information about the users’ behavior

This paper is based on “An Efficient Algorithm for Resource Sharing in Peer-to-Peer Networks”, a preliminary version published in the Proceedings of 5th International IFIP-TC6 Networking Conference, Coimbra, Portugal, May 15-19, 2006, by Wei-Cherng Liao, Fragkiskos Papadopoulos, and Konstantinos Psounis.

is determined and maintained locally and without any interaction with the other peers of the system (localized solutions), or it is determined and maintained by either a centralized authority (non-localized centralized solutions), or by the continuous exchange of information among the system's participants (non-localized distributed solutions). Localized solutions are simple and impose very little overhead but they are easy to subvert. Non-localized solutions are hard to subvert but are complex to use in practice. (KaZaA, eMule, and BitTorrent all use localized approaches.) In any case, our proposed scheme can be easily implemented in any way, and we describe later in the paper how to implement it efficiently with each one of the approaches.

The rest of the paper is organized as follows: In Section II we briefly discuss prior work on providing incentives for P2P systems. In Section III we provide a detailed description of the proposed scheme. In Section IV we provide a mathematical model that predicts the performance of the system and gives guidelines on how to set the scheme's parameters. In Section V we present realistic experiments of P2P systems on top of TCP networks. In Section VI we briefly discuss some implementation thoughts and finally conclude in Section VII.

II. RELATED WORK

There has been a large body of work on incentive mechanisms for P2P networks. Three of the most popular localized schemes are the ones implemented by the eMule [11], the KaZaA [12], and the BitTorrent [13] systems. eMule rewards users that provide files to the system by reducing their waiting time when they upload files using a scoring function (called *QueueRank*). Similarly, in KaZaA, each peer computes locally its *Participation Level* as a function of download and upload volumes, and peers with high participation levels have higher priority [10]. A disadvantage of both of these schemes is that they provide relatively weak incentives for cooperation since peers that have not contributed to the system at all can still benefit from it, if they are patient enough to wait in the upload queues. Other problems include that they favor users with high access bandwidth, which may result in frustration or a feeling of unfairness [14], and that they are vulnerable to the creation and distribution of hacked daemons that do not conform to the desired behavior [15]. BitTorrent uses a different scheme that is specific to its architecture. Each peer periodically stops offering uploads to its neighbors that haven't been offering uploads to him/her recently. This scheme is hard to subvert. However, it suffers from some unfairness issues and it only works with "BitTorrent-style" systems, that is, in systems where files are broken into pieces, and downloading a file involves being connected to almost all of ones neighbors in order to collect and reassemble all the pieces of the file.

Non-localized proposals are primarily concerned with creating systems that cannot be subverted. Some of them make use of credit/cash-based systems. They achieve

protection from hackers by either using central trusted servers to issue payments (centralized approach), e.g. [3], [4], or by distributing the responsibility of transactions to a large number of peers (distributed approach), e.g. [6]. Other distributed approaches use lighter-weight exchanged-based mechanisms, e.g. [16], or reputation management schemes, e.g. [5]. These mechanisms are indeed hard to subvert but they are also quite complex to use in practice [16].

In this paper we decouple the issue of how to design an algorithm to prevent free-riding from the issue of how to implement this algorithm in a P2P system. We first propose an efficient scheme that provides very strong incentives for cooperation. We show this via both theory and simulations. Then, we show that the scheme is generally applicable to any P2P system and comment on how to implement it using either a localized, or a non-localized approach. Another important contribution is the theoretical analysis of the performance of a P2P system with and without the proposed scheme. The analysis yields a set of equations that are used to predict the system's performance under a wide range of conditions, and to tune the parameters of the scheme.

III. A SIMPLE AND EFFECTIVE ALGORITHM

As mentioned earlier, the algorithm uses tokens as a means to trade bytes within the system. Each user is given an initial number of tokens M when he/she first joins the network. This allows new users to start downloading a small number of files as soon as they join the system. When a user rejoins the system he/she uses the amount of tokens he/she previously had.

Users spend K_{down} tokens for each byte they download from the system and earn K_{up} tokens for each byte they upload to the system. This forces users to offer files for upload proportionally to the number of files they want to download. Further, users gain K_{on} tokens/sec while being online. This mechanism of accumulating tokens serves two purposes. First, it allows users who are not contacted frequently for uploads to gain tokens by just being online, which is more fair towards users with low access bandwidth [14]. Second, it provides an incentive for users to keep their machines on the system even when they are not downloading a file, which helps to prevent the so-called problem of "low availability" [17]. Note that the value of K_{on} should be relatively small, in order to prevent users from gaining many tokens by just keeping their machines on without providing any uploads. Finally, a user can initiate a download only if the number of tokens he/she currently possesses is greater or equal to the number of tokens required to download the requested file.

This scheme provides strong incentives for cooperation. Free-riders are "forced" to provide some uploads to the system in order to gain tokens fast enough to sustain their desirable download demands. Some free-riders may decide to share their files as soon as they are out of tokens. Others may adopt a more dynamic behavior and decide to

adjust the number of uploads they provide to the system as a function of the number of tokens they currently have. In any case, the change in the free-rider's behavior increases the amount of available system resources tremendously, which, in turn, significantly improves the system's performance, as we shall see in Section V.

IV. A MATHEMATICAL MODEL FOR THE PROPOSED SCHEME

In this section we derive a mathematical model, which can be used to tune the parameters of the scheme and to predict the system's performance. Tuning the parameters of the scheme is important because an arbitrary setting of their parameters may lead to several undesired situations. For example, giving a large value to K_{on} may provide tokens to the free-riders fast enough, so that there won't be any reason for them to start sharing their files with the system. As another example, giving relatively small values to both K_{on} and K_{up} may reduce the token accumulation rate of cooperative users so much such that they cannot sustain their download demands. Further, predicting the performance of the system from the model is beneficial because the alternative is P2P simulations/experiments, and those either involve a significantly smaller number of peers than the number in reality, or are prohibitively expensive.

A. Token Dynamics

We assume a system that implements the proposed scheme, which we call "system with the tokens". Recall that K_{down} and K_{up} are expressed in tokens/byte and K_{on} in tokens/sec. Now, let C_{down} and C_{up} denote the file download and upload speeds of a user (access line bandwidth), both expressed in bytes/sec. The user spends $K_{down}C_{down}dt$ tokens if he/she is downloading files from other peers during time $(t, t + dt)$. Also, he/she earns $K_{on}dt$ tokens if he/she is online during time $(t, t + dt)$ and $K_{up}C_{up}dt$ tokens if other users are downloading files from the user under study during time $(t, t + dt)$. Let $T(t)$ denote the number of the user's tokens at time t , with $T(0) \geq 0$. We can then write the following differential equation:

$$\frac{dT(t)}{dt} = K_{on}I_{on}(t) + K_{up}C_{up}I_{up}(t) - K_{down}C_{down}I_{down}(t), \quad (1)$$

where

$$I_{on}(t) = \begin{cases} 1 & \text{if the user is online in } (t, t+dt) \\ 0 & \text{otherwise} \end{cases},$$

$$I_{up}(t) = \begin{cases} 1 & \text{if the user provides uploads in } (t, t+dt) \\ 0 & \text{otherwise} \end{cases},$$

$$I_{down}(t) = \begin{cases} 1 & \text{if the user performs downloads in } (t, t+dt) \\ 0 & \text{otherwise} \end{cases}.$$

Taking expectations on both sides of Equation (1), and interchanging the expectation with the derivative on the

left hand side¹, we get:

$$\frac{dE[T(t)]}{dt} = K_{on}P_{on}(t) + K_{up}C_{up}P_{up}(t) - K_{down}C_{down}P_{down}(t), \quad (2)$$

where $P_{on}(t)$ is the probability that the user is online at time t , $P_{up}(t)$ is the probability that the user provides uploads to the system at time t , and $P_{down}(t)$ is the probability that the user performs downloads from the system at time t . Note that Equation (2) can be regarded as a fluid model describing the token dynamics.

$P_{on}(t)$, $P_{up}(t)$, and $P_{down}(t)$ depend on how the user behaves given the number of tokens that he/she has at some point in time, and on his/her download demands. Along these lines, one can define user profiles, e.g. for non-free-rider and free-rider users, and solve the corresponding differential equation. The solution can be used to study how the expected amount of tokens of the particular class of users evolves as a function of time, for different values of the scheme's parameters and for different download and upload speeds. We next demonstrate this by considering free-rider users, under a simplistic profile that captures their main behavior. (One can also perform a similar study for non-free-rider users.)

B. A Free-rider Profile Example: The Linear Model

As mentioned earlier, free-riders are motivated to provide uploads to the system when they do not have enough tokens to sustain their download demands and they may lose their willingness to provide uploads, as the amount of tokens they possess is larger than the amount of tokens they need. To capture the main characteristics of this behavior we introduce a simple model, which we call the *Linear Model*.

According to the *Linear Model*, a free-rider provides uploads to the system with a high (constant) probability, say P_{upmax} , when his/her amount of tokens is less than some threshold, say T_{th1} . When his/her amount of tokens reaches T_{th1} , the probability that he/she provides uploads to the system is linearly decreasing as the amount of tokens he/she possesses continues to increase. (This is the reason why we call it the *Linear Model*.) The decrease continues until the amount of tokens he/she has reaches some threshold, say T_{th2} . After T_{th2} , we assume that the free-rider provides uploads to the system with some low (constant) probability, say P_{upmin} . This behavior is depicted in Figure 1 and characterized by Equation (3).²

$$P_{up}(t) = \begin{cases} P_{upmax} & \text{if } T(t) \leq T_{th1} \\ P_{upmin} & \text{if } T(t) \geq T_{th2} \\ P_{upmax} \left(1 - C_1 \frac{T(t) - T_{th1}}{T_{th2} - T_{th1}} \right) & \text{otherwise} \end{cases}, \quad (3)$$

¹Taking into account that $T(t)$ is bounded in practice, we can use the bounded convergence theorem [18] to justify the interchange.

²Determining more realistic free-rider profiles of their upload behavior as a function of the amount of tokens they possess, is an interesting and challenging research problem in its own right, and it is out of the scope of this paper.

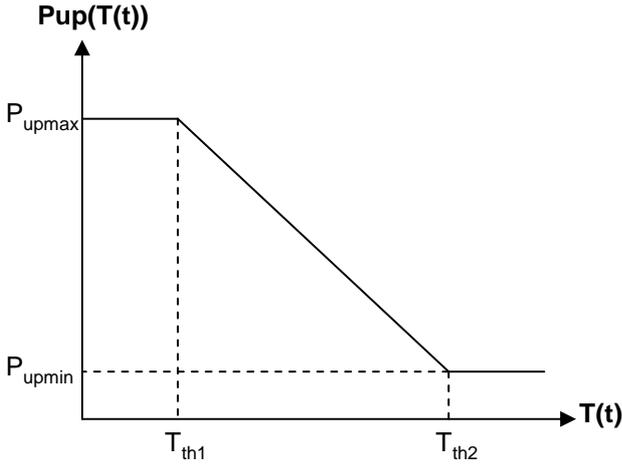


Figure 1. The Linear Model for P_{up} .

where $C_1 = \frac{P_{upmax} - P_{upmin}}{P_{upmax}}$.

For ease of exposition, we assume that $P_{on}(t) = 1$, i.e. users are always online, and that $P_{down}(t) = P_{down}$, i.e. independent of time t . Then, Equation (2) along with Equation (3) yield:

$$\frac{dE[T(t)]}{dt} = \begin{cases} K_{on} + C_2 - C_3 & \text{if } T(t) < T_{th1} \\ K_{on} - C_3 & \\ +C_2 \left(1 - C_1 \frac{T(t) - T_{th1}}{T_{th2} - T_{th1}}\right) & \text{otherwise} \end{cases} \quad (4)$$

where

$$\begin{aligned} C_2 &= K_{up}C_{up}P_{upmax}, \\ C_3 &= K_{down}C_{down}P_{down}. \end{aligned}$$

Solving the above differential equation with an initial condition $T(0) = 0$, gives:

$$E[T(t)] = \begin{cases} \frac{(K_{on} + C_2 - C_3)t}{C_4T_{th1} + K_{on} + C_2 - C_3} & \text{if } t < t_0 \\ \frac{C_4}{C_3 - K_{on} - C_2} + \frac{C_4e^{C_4(t-t_0)}}{C_4e^{C_4(t-t_0)}} & \text{otherwise} \end{cases} \quad (5)$$

where

$$\begin{aligned} C_4 &= \frac{K_{up}C_{up}(P_{upmax} - P_{upmin})}{T_{th2} - T_{th1}}, \\ t_0 &= \frac{T_{th1}}{K_{on} + K_{up}C_{up}P_{upmax} - K_{down}P_{down}C_{down}}. \end{aligned}$$

We can now assign values to K_{on} , K_{up} , K_{down} , C_{up} , C_{down} , P_{upmax} , P_{upmin} , T_{th1} , T_{th2} , and P_{down} , and use Equation (5) to study the dynamics of the expected amount of tokens of a free-rider. For example, Figure 2 shows how the expected number of tokens evolves as a function of time, when $K_{on} = 1000$ tokens/sec, $K_{up} = 0.5$ tokens/byte, $K_{down} = 1$ tokens/byte, $C_{up} = 1.5$ Mbps, $C_{down} = 1.5$ Mbps, $P_{upmax} = 1$, $P_{upmin} = 0.1$, $T_{th1} =$

10^6 tokens, $T_{th2} = 10^7$ tokens, and $P_{down} = 0.3$. From the plot we can observe that, initially, the expected amount of tokens increases linearly as a function of time, which represents the first part of Equation (5) (where $t < t_0 = 84$). After the amount of the free-rider's tokens reaches threshold $T_{th1} = 10^6$ (which occurs at time $t_0 = 84$), the token accumulation rate starts decreasing, since the free-rider now provides uploads to the system with a lower probability. Finally, the user adapts to an equilibrium point, where the token spending rate equals the token earning rate, and the expected amount of tokens settles to a steady state value, which we refer to as T_{ss} .³

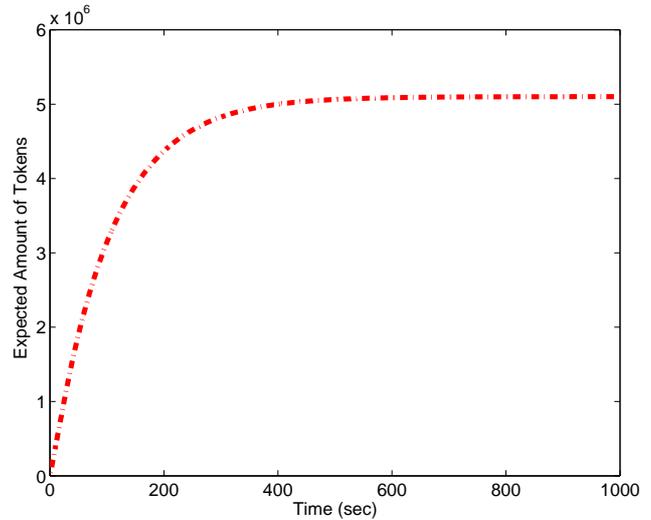


Figure 2. Expected amount of tokens for a free-rider vs. time.

Notice that T_{ss} can be obtained directly by either letting $t \rightarrow \infty$ in Equation (5), or by setting Equation (4) equal to zero, i.e. without the need of solving the differential equation first. It is given by:

$$T_{ss} = \frac{(K_{on} + C_2 - C_3)(T_{th2} - T_{th1})}{C_1C_2} + T_{th1}. \quad (6)$$

Now, substituting Equation (6) into Equation (3), one can also get an expression of the steady state upload probability. This probability, denoted by P_{upss} , is given by:

$$P_{upss} = \frac{K_{down}P_{down}C_{down} - K_{on}}{K_{up}C_{up}}. \quad (7)$$

Equations (6) and (7), which are much simpler than Equations (5) and (3), can be used for “back-of-the-envelope” calculations, and for gaining a better intuition on the long-run system behavior.

We now proceed to study the impact of different values of K_{up} and K_{on} on token dynamics. First, let's keep K_{on} , and the rest of the parameters fixed, and vary K_{up} . The results are shown in Figure 3. From the plots, we observe that as K_{up} increases, T_{ss} increases and P_{upss} decreases. This can be easily justified by looking at Equations (6) and (7). Intuitively, as K_{up} increases, free-riders are

³Similar results hold for other values of the model parameters.

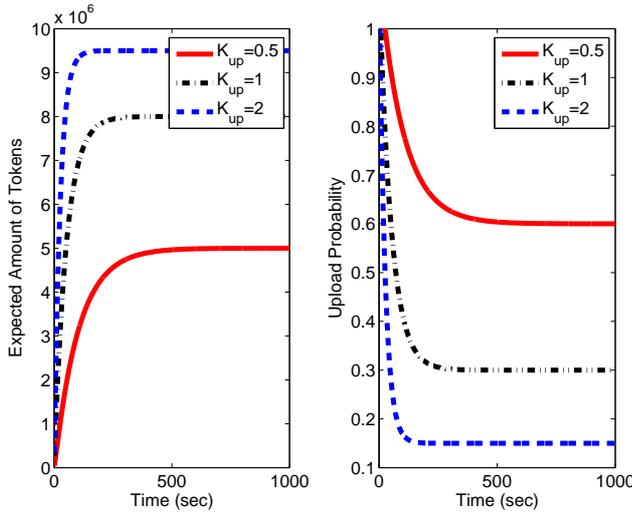


Figure 3. The effect of different values of K_{up} .

accumulating tokens faster and are able to settle to a larger T_{ss} . At the same time, they need to provide fewer uploads in order to gain the tokens they need to sustain their download demands, which yields a smaller P_{upss} .

Now, let's study the impact of the proposed scheme when we vary K_{on} , keeping the rest of the parameters fixed. The results are shown in Figure 4. First, from the

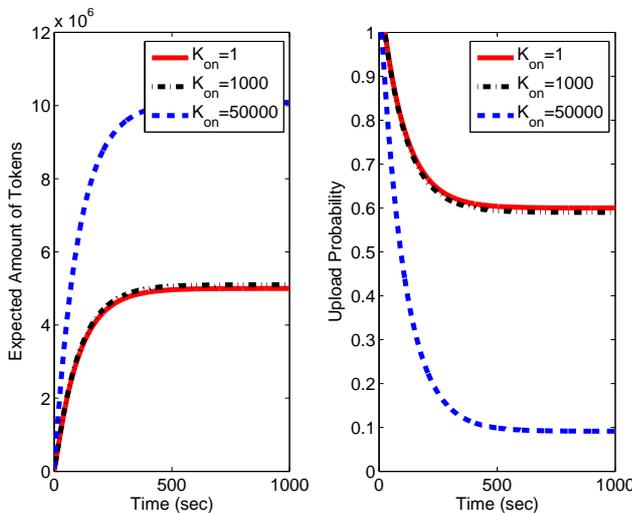


Figure 4. The effect of different values of K_{on} .

plots we observe that the effect of choosing different values of K_{on} is subtle when K_{on} is much smaller than the corresponding (long-run) token earning rate from uploads, which is $K_{up}C_{up}P_{upss}$. This is expected, as small K_{on} values do not have a great influence on the mechanism by which free-riders can accumulate tokens, and thus T_{ss} and P_{upss} , remain almost unaltered. On the other hand, a comparatively large value of K_{on} can have a great impact on the token accumulation mechanism, and thus on T_{ss} and P_{upss} . In particular, as we can see from the plots, it can significantly increase T_{ss} and decrease P_{upss} . As before, this can be easily justified by looking

at Equations (6) and (7). Intuitively, with larger values of K_{on} , free-riders can accumulate tokens faster and are able to settle to a larger T_{ss} . And this can be accomplished by just staying online, without the need of providing many uploads to the system, which yields a smaller P_{upss} .

C. Steady State Performance Analysis

So far, we have seen how one can use Equation (2) in order to study the token dynamics for a specific class of users (both transient and steady state behavior), and we have demonstrated the effect of using different values of the scheme's parameters on the token accumulation and the upload probability. We are now interested in using Equation (2) to derive a mathematical model that can be used for predicting steady state performance metrics (such as user download/upload rates, which we will define shortly), and for tuning the scheme's parameters in order to achieve a target performance.

To study the steady state we set $\frac{dE[T(t)]}{dt} = 0$ and drop the time dependence from the probabilities in Equation (2). Note that the existence of a steady state can be easily justified for a free-rider independently of his/her exact profile, by taking into consideration that in the long run he/she will spend as many tokens as he/she gains.⁴

Now, let R_{up} be the long-run average rate of file upload requests per second that the user handles, which we refer to as the *upload rate*. Also, let R_{down} be the long-run average rate of file download requests per second that the user initiates, which we refer to as the *download rate*. Last, let S denote the average file size in the system in bytes. Then, it is easy to see that $P_{up} = \frac{R_{up}S}{C_{up}}$ and $P_{down} = \frac{R_{down}S}{C_{down}}$. Equation (2) in steady state yields:

$$K_{on}P_{on} + K_{up}R_{up}S - K_{down}R_{down}S = 0.$$

Taking the average over all free-riders yields:

$$K_{up} = K_{down} \left(\frac{E[R_{down}|FR]}{E[R_{up}|FR]} \right) - \frac{K_{on}P_{on}^{FR}}{E[R_{up}|FR]S}, \quad (8)$$

where P_{on}^{FR} is the long-run probability that a free-rider is online.⁵

Equation (8) relates the parameters of the scheme, K_{on} , K_{up} , and K_{down} , with the average download and upload activity of free-riders. We will later use it to select the parameter values that yield a target performance. But first, we need to compute the average download and upload rates.

1) *User Download Rate (R_{down}):* Let N be the number of peers in the system and let a proportion α of them be free-riders. Assume that free-riders are uniformly distributed over the system. Also, assume that both non-free-riders and free-riders have the same download demands. In particular, they have the same query request

⁴Considering the existence of a steady state for a non-free-rider is a bit of more involved. As we will shortly see he/she may or may not have a steady state. Nevertheless, this will not be important for the system dynamics.

⁵Similarly, we denote the long-run probability that a non-free-rider is online by P_{on}^{NF} .

rate, denoted by R_q queries/sec, and the same preference over files, that is, each query is for file i with some probability $Q_f(i)$ irrespectively of the query issuer.

Now, in order to proceed, we have to define an upload profile for free-riders, which determines how they respond to query requests given the number of tokens they have (i.e. when do they initiate an upload?). To make the analysis tractable, we assume the following simple upload profile: Free-riders respond to a query request only if the amount of tokens they currently have is less than the amount required to download the file they currently desire. (Recall that non-freeriders always respond to query requests.)

Let $P_{ans}(i)$ be the probability that a query request for file i is successfully answered. Now, recall that in the system with the tokens a user can initiate a download only if the amount of tokens he/she has is larger than the amount required to download the file. Let P_{tkn}^{FR} and P_{tkn}^{NF} denote respectively the probability that a free-rider and a non-freerider have enough tokens to initiate a download. Then, we can express the average download rate of free-riders and non-freeriders as follows:

$$E[R_{down}|FR] = \sum_i R_q \cdot Q_f(i) \cdot P_{ans}(i) \cdot P_{tkn}^{FR} \cdot P_{on}^{FR}, \quad (9)$$

$$E[R_{down}|NF] = \sum_i R_q \cdot Q_f(i) \cdot P_{ans}(i) \cdot P_{tkn}^{NF} \cdot P_{on}^{NF}, \quad (10)$$

where the summation is taken over all files i . Clearly, the average download rate over all users in the system is:

$$E[R_{down}] = E[R_{down}|FR] \cdot \alpha + E[R_{down}|NF] \cdot (1 - \alpha). \quad (11)$$

To complete the calculation of the download rates, note that R_q , $Q_f(i)$, P_{on}^{FR} , P_{on}^{NF} , and α are given quantities. (There exist a large body of work in measurement studies of P2P systems, e.g. [19], [20], from which one can deduce typical values for these quantities.) Hence, what remains is to compute P_{ans} , P_{tkn}^{FR} , and P_{tkn}^{NF} . We start by deriving a relation between P_{tkn}^{FR} and P_{tkn}^{NF} . First, recall that in steady state the token earning rate equals the token spending rate for each free-rider. A free-rider responds to a query request only when he/she doesn't have enough tokens and certainly when he/she is online, i.e. with probability $(1 - P_{tkn}^{FR})P_{on}^{FR}$. Since a non-freerider always responds to a query request when he/she is online, it is easy to see that the token earning rate of free-riders over that of non-freeriders equals $\frac{(1 - P_{tkn}^{FR})P_{on}^{FR}}{P_{tkn}^{NF}}$. Now, the token spending rate is proportional to the download rate, and Equations (9) and (10) imply that the token spending rate of free-riders over that of non-freeriders equals $\frac{P_{tkn}^{FR} P_{on}^{FR}}{P_{tkn}^{NF} P_{on}^{NF}}$. Assuming that non-freeriders are also in steady state (in which case Equation (8) also holds if the average is taken with respect to non-freeriders only), we can equate these two ratios and write $P_{tkn}^{NF} = \frac{P_{tkn}^{FR}}{(1 - P_{tkn}^{FR})}$. Clearly, this equality is valid for $P_{tkn}^{FR} \leq 0.5$. In particular when $P_{tkn}^{FR} = 0.5$, $P_{tkn}^{NF} = 1$, which implies that non-freeriders always have enough tokens to initiate downloads. For

$P_{tkn}^{FR} > 0.5$, the last equality no longer holds. In this case the token earning rate of non-freeriders will be larger than their token spending rate, which implies that their amount of tokens will continuously increase. However, this still suggests that $P_{tkn}^{NF} = 1$. We can now write:

$$P_{tkn}^{NF} = \min \left(1, \frac{P_{tkn}^{FR}}{1 - P_{tkn}^{FR}} \right). \quad (12)$$

Now lets find a relation for $P_{ans}(i)$. First, assume that due to congestion at the overlay layer [21], each message (either a query request or a query response) has a probability p of being dropped at some peer.⁶ Then, if L is the average number of overlay hops until a query is answered, $P_{drop} = 1 - (1 - p)^L$ is the probability that the query response is lost. Next, observe that if $K \leq N$ is the average number of peers that a query request can reach if all users were online, the request can be answered by an average of $K \cdot (P_{on}^{FR} \cdot (1 - P_{tkn}^{FR}) \cdot \alpha + P_{on}^{NF} \cdot (1 - \alpha))$ peers, which we call K_{eff} . Finally, let $P_f(i)$ be the probability that a peer has file i . We can then write:

$$P_{ans}(i) = 1 - (1 - P_f(i) \cdot (1 - P_{drop}))^{K_{eff}}. \quad (13)$$

2) *User Upload Rate (R_{up}):* The total number of downloads equals the total number of uploads in any system and thus the expected download and upload rates over *all* nodes are also equal. This does not mean that all peers provide uploads. For example, in a system that does not implement the proposed scheme, $E[R_{down}] = E[R_{up}]$, but we know that only non-freeriders provide uploads, i.e. $E[R_{up}|FR] = 0$, and hence $E[R_{up}|NF] = \frac{E[R_{down}]}{(1 - \alpha)}$. On the other hand, in the system with the tokens each free-rider answers to a query request with probability $P_{on}^{FR}(1 - P_{tkn}^{FR})$. As a result, this system behaves as if there are $N \cdot (P_{on}^{NF} \cdot (1 - \alpha) + P_{on}^{FR} \alpha \cdot (1 - P_{tkn}^{FR}))$ non-freeriders. It is easy to see that the expected upload rate of each non-freerider is now given by:

$$E[R_{up}|NF] = \frac{P_{on}^{NF} \cdot E[R_{down}]}{P_{on}^{NF} \cdot (1 - \alpha) + P_{on}^{FR} \cdot \alpha \cdot (1 - P_{tkn}^{FR})}. \quad (14)$$

And, since $E[R_{up}] = E[R_{down}]$, the expected upload rate of each free-rider equals:

$$E[R_{up}|FR] = \frac{P_{on}^{FR} \cdot (1 - P_{tkn}^{FR}) \cdot E[R_{down}]}{P_{on}^{NF}(1 - \alpha) + P_{on}^{FR} \alpha \cdot (1 - P_{tkn}^{FR})}. \quad (15)$$

D. Choosing the right values for K_{on} , K_{up} , and K_{down}

We use P_{tkn}^{FR} as the design parameter of our system since it dictates how often free-riders offer uploads, which, in turn, specifies the average amount of available resources in the system. We are given the query- and file-popularity probability functions $Q_f(i)$, $P_f(i)$, the query request rate R_q , the user statistics P_{on}^{FR} , P_{on}^{NF} , and information about the overlay network. (For example,

⁶This assumption is introduced to make the model more general. A well designed system usually has $p \approx 0$, which is accomplished by setting the buffer size of the TCP socket sufficiently large.

information about the overlay network includes the percentile of free-riders α , the socket buffer sizes that dictate the drop probability p , and the structure of the overlay graph as well as the search algorithm that dictate the number of peers that a query reaches K and the average path length between a query issuer and a query responder L .) We want to find a set of values for K_{on} , K_{up} and K_{down} that will satisfy a target P_{tkn}^{FR} , and, in turn, a target system performance.

First, observe from Equation (8) that it is the *relative* values of K_{on} , K_{up} , and K_{down} that are important for the proper operation of the system. Recall also that K_{on} should be sufficiently smaller than the token spending rate of free-riders. This is to prevent them from accumulating enough tokens by just staying online without offering any upload. Thus, we should have $K_{on} \ll K_{down}E[R_{down}|FR]S$.

With the above observations in mind we proceed as follows in order to satisfy the target P_{tkn}^{FR} :

- (i) Fix K_{down} to some arbitrary value,
- (ii) use Equation (12) to compute P_{tkn}^{NF} , (To guarantee that cooperative users will not be penalized, P_{tkn}^{NF} should be close to 1.)
- (iii) use Equations (9) and (13) to compute the value of $E[R_{down}|FR]$, and Equations (15), (11) and (10) to compute $E[R_{up}|FR]$,
- (iv) assign a value to K_{on} which is one order of magnitude smaller than $K_{down}E[R_{down}|FR]S$, (The specific value turns out not to affect the performance sizeably.) and
- (v) use Equation (8) to find the right value for K_{up} .

Conversely, if we are given the values of K_{on} , K_{up} , and K_{down} we can use our equations to predict quantities like $E[R_{down}|FR]$, $E[R_{down}|NF]$, $E[R_{up}|FR]$ and so on.⁷

In the next Section we verify the accuracy of our analysis via experiments on top of TCP networks, and show the impact of the proposed scheme on system's performance.

V. EXPERIMENTS

A. Simulation Setup

For our experiments we use GnutellaSim [22], a packet-level peer-to-peer simulator build on top of ns-2 [23], which runs as a Gnutella system. We implement the file downloading operation using the HTTP utilities of ns-2.

We use a 100-node transit-stub network topology as the backbone, generated with GT-ITM [24]. We attach a leaf node to each stub node. Each leaf node represents a peer. The propagation delays assigned to the links of the topology are proportional to their length and are in the order of *ms*. We assign capacities to the network such that the congestion levels are moderate. The capacity assigned to a peer's access link is 1.5Mbps.

In order to test the algorithm on a general gnutella-like unstructured P2P network we use Gnutella v0.4, which

uses pure flooding as the search algorithm and does not distinguish between peers. The *TTL* for a query request message is set to 7 (the default value used in Gnutella).

For simulation purposes we implement the following user behavior: each user initiates query requests at the constant rate of 1 query every 20sec. Once a timeout for a query request occurs, the corresponding query is retransmitted. The maximum number of retransmissions is set to 5,⁸ and the timeout to 60sec.

There are 1000 distinct files in the system, $i = 1 \dots 1000$. A query request is for file i with probability proportional to $\frac{1}{i}$ (Zipf distribution). The number of replicas of a certain file is also described by a Zipf distribution with a scaling parameter equal to 1, and the replicas of a certain file are uniformly distributed among all peers. These settings are in accordance with measurement studies from real P2P networks [19], [20]. We distinguish two systems: (i) the original system which does not implement the proposed algorithm, and (ii) the system with the tokens. In both systems, 85% of peers are free-riders in accordance to the percentage reported in [2]. Finally, the file size is set to 1MB.

We first perform simulations for the following two scenarios: (i) when $P_{on} = 1$, i.e. when all peers initially join the system and never go offline, and (ii) when $P_{on} < 1$, i.e. when peers dynamically join and leave the system according to P_{on} . Then, we study what the impact on system's performance is, when some malicious users subvert the proposed scheme.

B. Simulation Results for $P_{on} = 1$

1) *Download and Upload Rates:* For various values of the design parameter P_{tkn}^{FR} we compute the corresponding values of K_{on} , K_{up} and K_{down} according to the procedure described in the previous Section. We then assign these values to all users of the system and compare the theoretical download and upload rates with the experimental results. Figures 5 and 6 show respectively the expected download and upload rate over all non-free-riders, over all free-riders, and over all users of the system, as a function of P_{tkn}^{FR} .

The horizontal line in Figure 5 represents the expected download rate of a user in the original system. (Clearly, in the original system $E[R_{down}] = E[R_{down}|FR] = E[R_{down}|NF]$.) The horizontal line in Figure 6 represents the expected upload rate of a non-free-riders in the original system. (Recall that in this system $E[R_{up}|FR] = 0$.)

It is clear from the plots that analytical and simulation results match. Further, we can make several interesting observations. First, notice that as P_{tkn}^{FR} increases, the download rate for both classes of users first increases and then starts decreasing until it reaches the value of the original system. Second, observe that while the upload rate of free-riders behaves in a similar manner, the upload rate of non-free-riders continuously increases until

⁷Note that we can also use Equations (9)...(15) to compute upload/download rates in a system that does not implement the scheme, by setting $P_{tkn}^{FR} = 1$.

⁸This corresponds to the situation where a user quits searching for a certain file, after reattempting an unsuccessful search, e.g. by using a different keyword, for 5 consecutive times.

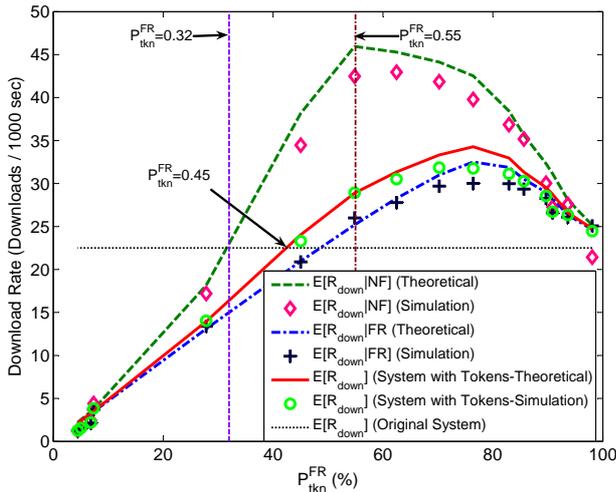


Figure 5. User's expected download rate.

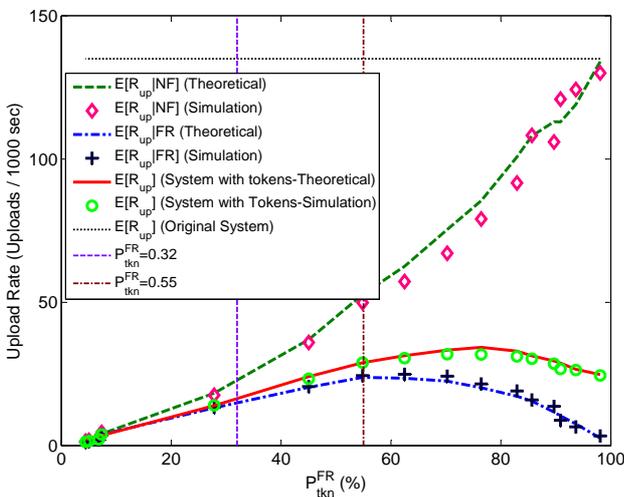


Figure 6. User's expected upload rate.

it reaches its original value. Based on these observations we divide the plots into three regions. The first region corresponds to $P_{tkn}^{FR} < 0.32$. In this region, both classes of peers are constrained to a lower download rate compared to the original system, since the probability of having tokens to initiate a new download after a successful query is pretty low. Notice that for $P_{tkn}^{FR} = 0.32$, and hence for $P_{tkn}^{NF} = 0.47 < 1$, cooperative users can at least sustain the same download rate they had in the original system. The second region corresponds to $0.32 \leq P_{tkn}^{FR} \leq 0.55$. In this region, users accumulate tokens at a higher rate than before. Since there are more responses than in the original network, users can use the extra tokens to initiate more downloads. Notice that cooperative users earn tokens faster than free-riders since they always respond to query requests. At $P_{tkn}^{FR} = 0.55$, non-free-riders achieve their maximum download rate, which is approximately twice the one they had in the original system. Finally, the third region corresponds to $0.55 < P_{tkn}^{FR} \leq 1$. In this region free-riders accumulate tokens faster than before

and reduce their query response rate since they do not need to provide as many uploads as before. This causes cooperative users to handle more uploads. Further, since the query response rate regulates the download rate, the latter also decreases. At $P_{tkn}^{FR} = 1$, the two systems have approximately the same performance, as expected.

2) *Impact On Delays:* Figures 7 and 8 show respectively the average query response time (that includes retransmissions) and the average download delay as a function of P_{tkn}^{FR} . The plots can be divided in the same three regions as before. For $P_{tkn}^{FR} < 0.32$, the low user download rate imposes a low load into the network. This yields the low delays. For $0.32 \leq P_{tkn}^{FR} \leq 0.55$, as the user download rate increases, the load in the network and hence the delays also increase. Note that the query and download delays are still significantly smaller than in the original system, despite that the download rate, and hence the load, is higher. This is because a significant portion of the load is now handled by the free-riders. For $0.55 < P_{tkn}^{FR} \leq 1$ the delays continue to increase even though the download rate decreases. This is because free-riders provide fewer and fewer uploads. As P_{tkn}^{FR} approaches 1, the performance of the two systems is approximately the same.

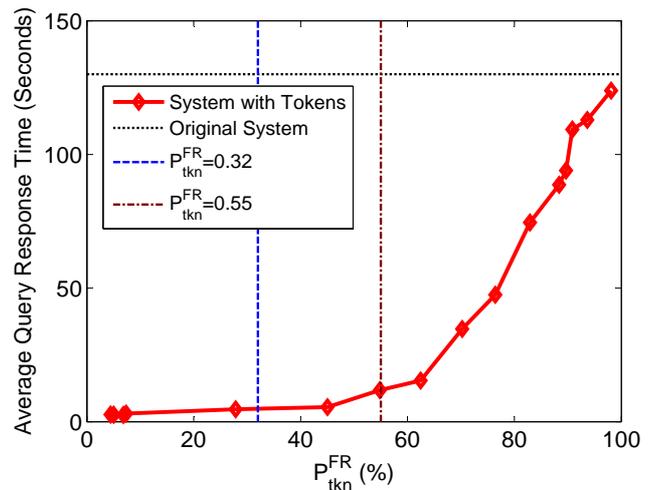


Figure 7. Average query response time.

To fairly compare the delays between the two systems, we should consider the case where the load is the same, i.e. where $E[R_{down}] = 22$ downloads/1000sec. This value corresponds to $P_{tkn}^{FR} = 0.45$, and as we can see from the plots this corresponds to approximately one order of magnitude lower query and file download delays. This is a gigantic amount of improvement on the system's performance.

As a final note, the best operating region is the second, where $0.32 \leq P_{tkn}^{FR} \leq 0.55$. In this region, we can either choose to operate the system at $P_{tkn}^{FR} = 0.32$, where cooperative users can sustain the same download demands as in the original system, or sacrifice a bit from the performance improvement with respect to reduced delays to support higher user demands.

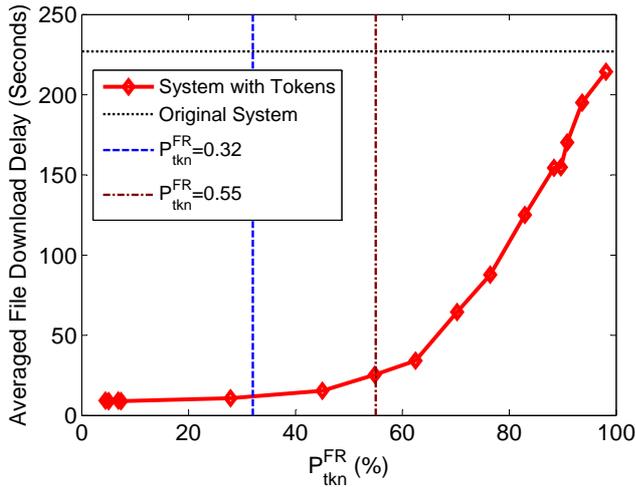


Figure 8. Average file download delay.

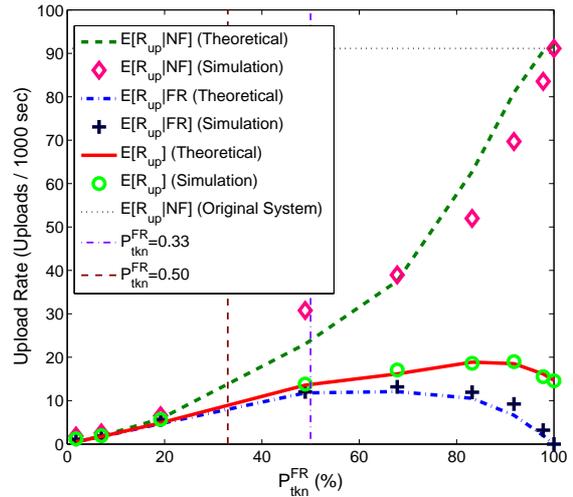


Figure 10. User's expected upload rate.

C. Simulation Results for $P_{on} < 1$

We now study the impact of the proposed scheme in the more realistic scenario where $P_{on} < 1$. In particular, we now set $P_{on}^{NF} = P_{on}^{FR} = 0.5$.⁹ As before, we first compare the theoretical and simulation results for the expected download and upload rates, and then show the impact of the scheme on system delays.

1) *Download and Upload Rates:* Figures 9 and 10 show respectively the expected download and upload rate over all non-freeriders, over all free-riders, and over all users of the system, as a function of P_{tkn}^{FR} .

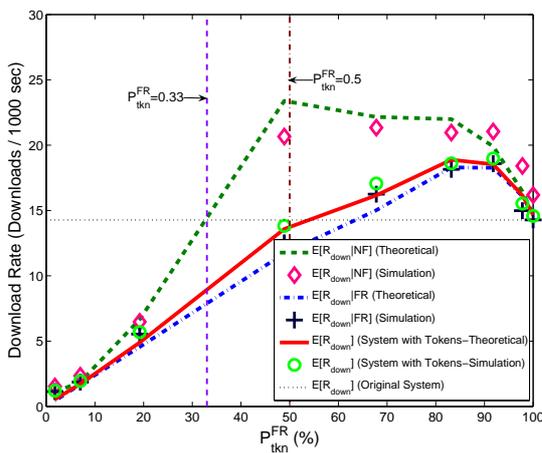


Figure 9. User's expected download rate.

As before, the horizontal line in Figure 9 represents the expected download rate of a user in the original system. And, the horizontal line in Figure 10 represents the expected upload rate of a non-freerider in the original system. We again see from the plots that analytical and simulation results match. Further, we can divide the plots

⁹ P_{on}^{FR} and P_{on}^{NF} do not have to be equal. Similar results hold for other values of these parameters.

into regions and justify the behavior in each region, just like we did before.

2) *Impact On Delays:* Figures 11 and 12 show respectively the average query response time and the average download delay as a function of P_{tkn}^{FR} . We again observe a similar behavior as with the case where $P_{on} = 1$.

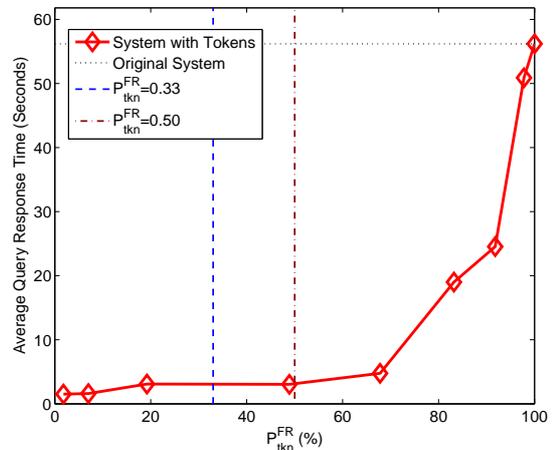


Figure 11. Average query response time.

As we can see from Figures 9, 10, 11, and 12, the performance improvement from utilizing the proposed scheme can still be quite significant, even if users dynamically join and leave the system. In particular, under the appropriate parameter tuning, one can again reduce the delays by one order of magnitude, while being able to sustain higher user download demands.

D. The Impact of Malicious Users on System's Performance

So far we have seen that system's performance can be tremendously improved when adopting the proposed

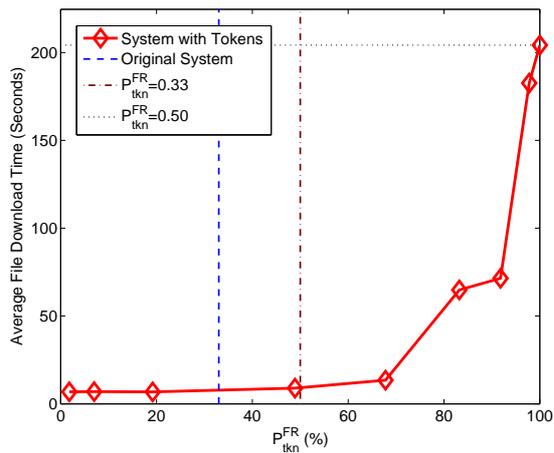


Figure 12. Average file download delay.

scheme. However, this was for the case where all users were well-behaved. We are now interested in studying how much performance degeneration is observed if a small percentage of users are malicious. In particular, let's assume that 10% of free-riders have managed to circumvent the proposed scheme, so that they can download files as long as they receive query responses from the system, independently from the amount of tokens they currently have.

Figures 13 and 14 show respectively the user expected download rate and the average file download delay in a system with and without malicious users. From the plots we observe that performance degeneration due to the presence of malicious users is almost negligible. Notice however that, in general, the user download rate in the system with malicious users is smaller than the rate in a system without malicious users, and that the average file download delay in the system with malicious users is larger. These observations can be easily justified. First, malicious users do not respond to query requests, which means a lower user download rate. And, since fewer users provide uploads, the file download delay increases.

VI. IMPLEMENTING THE SCHEME

As mentioned before, this scheme can be implemented either locally or non-locally. Implementing this scheme locally is quite simple. The local P2P client takes care of bookkeeping by increasing the user's tokens for each acknowledged byte he/she uploads and for being online, and by decreasing the tokens for each byte the user downloads. However as we have already mentioned, this approach is quite vulnerable to hacked clients.

There are several directions for making the hacking of localized solutions hard. One can utilize encryption techniques *e.g.* [25] that make unauthorized modifications to data (such as the scheme's parameters) hard. In addition, one can also use technologies like DRM (Digital Rights Management) in order to protect the entire client's code from being altered *e.g.* [26], and re-distribute new

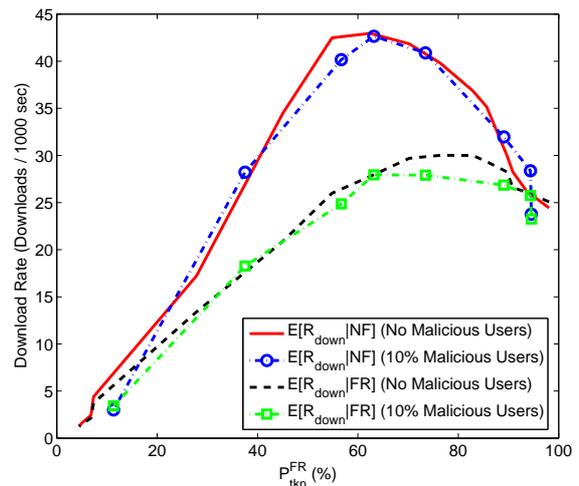


Figure 13. User's expected download rate.

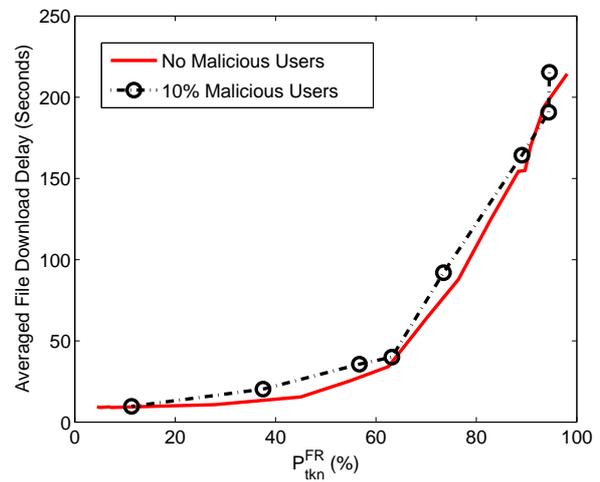


Figure 14. Average file download delay.

clients frequently in order to minimize the number of hacked clients that can be connected to the network. Further, one could also employ techniques such as tamper-proofing and self-checking in order to verify the client's code integrity during the join process and/or on every download request *e.g.* [27], [28], [29], [30]. Of course, the only way to guarantee that all P2P clients are original is to have a trusted platform where both the hardware and the operating system can be trusted [31]. This is clearly not an option in practice. However, interestingly enough, both theory and simulations dictate that our scheme is quite resilient to a *small* number of hacked clients. In particular, as we have seen earlier, the system performance is virtually unchanged when the hackers comprise less than 10%. Hence, all one needs to do is to make it hard for users to use hacked clients. The scheme can be also implemented in a secure non-localized centralized manner, where peers exchange messages with a centralized trusted authority that updates and maintains

their amount of tokens. Peers would communicate with the centralized authority once they finish downloading to report the source node and the file size, periodically while being online, and to get permission to initiate a new download. This is similar to the main idea that many centralized “cash-based” systems, e.g. [3], [4], follow. Finally, the scheme can be also implemented in a secure non-localized decentralized manner, e.g. by utilizing the framework suggested in [6].

VII. CONCLUSION

In this paper we studied a simple algorithm that provides strong incentives for cooperation in file sharing P2P networks. We derived a mathematical model that describes the system’s dynamics and which can be used for parameter tuning and performance prediction. We demonstrated the effectiveness of the algorithm via experiments with TCP networks.

We have several future work directions. First, we plan to perform larger scale experiments, and to implement the scheme in an operational P2P network. Further, we plan to extend our analytical methodology to compute other important performance metrics, e.g. the improvement on the expected download delays and response times. And finally, we would like to extend our study to include other kinds of P2P networks, e.g. such as “BitTorrent-style” systems.

REFERENCES

- [1] E. Adar and B. Huberman, “Free riding on gnutella,” http://www.firstmonday.dk/issues/issue5_10/adar, October 2000 (accessed Aug. 2005).
- [2] D. Hughes, G. Coulson, and J. Walkerdine, “Free riding on gnutella revisited: the Bell Tolls?” *IEEE Distributed Systems Online Journal*, vol. 6, no. 6, June 2005.
- [3] “Mojonation,” <http://www.mojonation.net/Mojonation.html> (accessed Aug. 2005).
- [4] J. Ioannidis, S. Ioannidis, A. Keromytis, and V. Prevelakis, “Fileteller: paying and getting paid for file storage,” in *Proc. of 6th International Conference on Financial Cryptography*, March 2002.
- [5] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, “EigenRep: Reputation management in P2P networks,” in *Proc. of 12th International World Wide Web Conference (WWW 2003)*, May 2003.
- [6] V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer, “KARMA: A secure economic framework for P2P resource sharing,” in *1st Workshop on Economics of Peer-to-Peer Systems*, June 2003.
- [7] C. Buragohain, D. Agrawal, and S. Suri, “A game-theoretic framework for incentives in P2P systems,” in *Proc. of International Conference on Peer-to-Peer Computing*, Sep 2003.
- [8] L. Ramaswanmy and L. Liu, “Free-riding: A new challenge to peer-to-peer file sharing systems,” in *Proc. of the 36th Hawaii international conference on system sciences*, 2003.
- [9] M. Feldman, C. Papadimitriou, I. Stoica, and J. Chuang., “Free-riding and whitewashing in Peer-toPeer systems,” in *SIGCOMM Workshop*, 2004.
- [10] “KaZaA participation level,” http://www.kazaa.com/us/help/glossary/participation_ratio.htm (accessed Aug. 2005).
- [11] “The emule project,” <http://www.emule-project.net/> (accessed Aug. 2005).
- [12] “KaZaA media desktop,” <http://www.kazaa.com/> (accessed Aug. 2005).
- [13] “Bittorrent,” <http://www.bittorrent.com/protocol.html> (accessed Aug. 2005).
- [14] H. Bretzke and J. Vassileva, “Motivating cooperation in peer to peer networks,” in *Proc. of User Modeling UM03*, June 2003.
- [15] “Hack KaZaA participation level,” <http://www.davesplanet.net/kazaa/> (accessed Aug. 2005).
- [16] K. Anagnostakis and M. Greenwald, “Exchanged-based incentive mechanisms for peer-to-peer file sharing,” in *Proc. of 24th International Conference on Distributed Computing Systems*, 2004.
- [17] R. Bhagwan, S. Savage, and G. M. Voelker, “Understanding availability,” in *Proc. of 2nd IPTPS*, 2003.
- [18] R. Durrett, *Probability: Theory and Examples*. Duxbury Press, 2nd edition, 1996.
- [19] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. Gribble, and H. M. Levy, “An analysis of internet content delivery systems,” in *Proc. of the Fifth Symposium on Operating System Design and Implementation (OSDI)*, December 2002.
- [20] J. Chu, K. Labonte, and B. N. Levine, “Availability and locality measurements of peer-to-peer file sharing systems,” in *Proc. of SPIEITCom: Scalability and Traffic Control in IP Networks*, July 2002.
- [21] M. A. Qi He, “Congestion control and message loss in gnutella networks,” in *Proc. of Multimedia Computing and Networking*, 2004.
- [22] “Packet-level Peer-to-Peer Simulation Framework and GnutellaSim,” <http://www.cc.gatech.edu/computing/compass/gnutella/> (accessed Oct. 2005).
- [23] “Network simulator,” <http://www.isi.edu/nsnam/ns> (accessed Sep. 2005).
- [24] K. Calvert, M. Doar, and E. W. Zegura, “Modeling internet topology,” *IEEE Communications Magazine*, 1997.
- [25] “Data encryption standard,” <http://www.itl.nist.gov/fipspubs/fip46-2.htm> (accessed Oct. 2005).
- [26] T. Sander, *Security and Privacy in Digital Rights Management*. Springer, 1st Edition, 2002.
- [27] D. Aucsmith, “Tamper resistant software: An implementation,” in *Proc. 1st International Information Hiding Workshop*, May 1996.
- [28] H. Chang and M. Atallah, “Protecting software code by guards,” in *Proc. of 1st ACM Workshop on Digital Rights Management*, May 2002.
- [29] Y. Chen, R. Venkatesan, M. Cary, R. Pang, S. Sinha, and M. Jakubowski, “Oblivious hashing: A stealthy software integrity verification primitive,” in *Proc. of 5th International Information Hiding Workshop*, October 2002.
- [30] C. Collberg and C. Thomborson, “Watermarking, tamper-proofing, and obfuscation - tools for software protection,” *IEEE Transactions on Software Engineering*, vol. 28, no. 6, June 2002.
- [31] S. W. Smith, *Trusted Computing Platforms: Design and Applications*. Springer, 1st Edition, 2004.
- [32] W.-C. Liao, F. Papadopoulos, and K. Psounis, “An efficient algorithm for resource sharing in peer-to-peer networks,” University of Southern California, Tech. Rep. CENG-2005-15, 2005.



Wei-Cherng Liao is a Ph.D. student in Electrical Engineering at the University of Southern California. He received his first degree and the M. Sc. degree from the department of Electrical and Computer Engineering of National Taiwan University, Taiwan, in June 1998 and June 2000 respectively. His research interests include modeling, and performance prediction/analysis of computer networks and the incentive scheme of P2P system.



Fragkiskos Papadopoulos is a Ph.D. student in Electrical Engineering at the University of Southern California. He received his first degree from the department of Electrical and Computer Engineering of National Technical University of Athens, Greece, in June 2002 and the M. Sc. degree in Electrical Engineering, specializing in computer networks, from University of Southern California in May 2004. His research interests include modeling, simulation, and performance prediction/analysis of computer networks and P2P systems. He is a recipient of the Fulbright Scholarship.



Konstantinos Psounis is an assistant professor of Electrical Engineering and Computer Science at the University of Southern California. He received his first degree from the department of Electrical and Computer Engineering of National Technical University of Athens, Greece, in June 1997, the M.S. degree in Electrical Engineering from Stanford University, California, in January 1999, and the Ph.D. degree in Electrical Engineering from Stanford University in December 2002.

Konstantinos models and analyzes the performance of computer networks, sensor and mobile systems, and the web. He also designs methods and algorithms to solve problems related to such systems. He is the author of more than 30 research papers on these topics. Konstantinos has received faculty awards from NSF and the Zumberge foundation, has been a Stanford graduate fellow throughout his graduate studies, and has received the best-student National Technical University of Athens award for graduating first in his class