Optimal media sharing policies in peer-to-peer networks

Ritesh Sood and Mihaela van der Schaar

Department of Electrical and Computer Engineering
University of California, Davis
Davis, CA 95616. USA

ABSTRACT

Multimedia content distribution through a distributed system, a peer-to-peer (P2P) network for instance, is attractive since it harnesses the resources available with the numerous peers in the network. Another advantage of such a system is that the potentially available resources scale in proportion to the demand as more and more peers join the system. Recent studies have concentrated mainly on such aspects of these distributed networks as querying, indexing, etc. These studies however take for granted the voluntary contribution of resources by peers in the system. Empirical evidence however points to the contrary, *i.e.* in existing P2P systems, a substantial fraction of peers do not contribute resources to the system, while benefiting from the services it provides at the expense of the contributing peers. In this paper we analyze a P2P system in a game-theoretic setting in which games involving content exchange are played repeatedly. The model takes into account the manner in which a peer adapts his contribution to the system depending on the benefit he has derived from the system so far and expects to derive in the long run. The model enables us to formulate an optimization problem that yields optimal content sharing strategies that a peer should adopt in order to maximize his net benefit by participating in the system.

1. INTRODUCTION

The problem of multimedia content distribution over computer networks (wired and wireless) is accentuated by the relatively large bandwidth requirements for this class of data. Multimedia content providers resort to multiple servers at the edges of the network that provide content to clients in their vicinity. This approach however does not address the issue of scale effectively. A possible solution to this problem is to harness the resources that are available with other peers in a network, *vis-a-vis* a peer-to-peer (P2P) file sharing system.[1, 2] Such an approach may be effectively employed for scalable video since different layers of the multimedia content may be provided by different peers depending on their upload bandwidth, processing power, etc.

Existing P2P systems are however known to suffer from the *free-rider* problem, *i.e.*, a large majority the of peers do not contribute to the network while enjoying the benefits that it provides. In the Gnutella system for example, as much as 25% of the peers do not share files at all.[3, 4] The reason commonly cited for this phenomenon is that users do not derive any immediate benefit by providing their services to the network, and hence lack the incentive to participate.

In Ref. 5, file sharing between peers in a P2P network has been analysed using a game theoretic approach. A peer is regarded as rational decision maker whose decision process is guided by the his desire of maximizing a suitably defined utility function. Under the adopted set of assumptions, the analysis leads to the conclusion that the unique equilibrium point for a P2P system is one in which *no* peer is willing to share. In order to resolve this apparent paradox (since files are plentiful on P2P systems such as Napster), the authors conclude that some peers in the system are driven by altruistic motives and it is their selfless contribution that keeps the system going. Subsequently, certain payment mechanisms are proposed and it is proved that desirable equilibrium points are attained under such mechanisms, in which widespread file sharing results. A payment mechanism however requires an infrastructure in order to facilitate these transactions and safeguard the parties against fraud. Also,

Further author information: (Send correspondence to Ritesh Sood.)
Ritesh Sood: E-mail: rsood@ucdavis.edu, Telephone: 1 530 754 9861
Mihaela van der Schaar: E-mail: mvanderschaar@ece.ucdavis.edu, Telephone: 1 530 754 6281
This research was supported in part by "IT@Intel Research".

Applications of Digital Image Processing XXVII,
edited by Andrew G. Tescher, Proceedings of SPIE Vol. 5558
(SPIE, Bellingham, WA, 2004) · 0277-786X/04/$15 · doi: 10.1117/12.564454

since the payment mechanism is required to oversee each and every transaction that is carried out, the cost of deploying and maintaining the required infrastructure will be prohibitive considering the size of a typical peer-to-peer network. Moreover, existence of a centralized authority runs contrary to the principle characteristic of a P2P network, *viz.*: a distributed, self-organizing system.

An alternative approach for alleviating the free-rider problem has been adopted in Ref. 6. Here the authors seek to provide incentives for file-sharing through the *differential service* model in which the benefit that a peer can derive from the system, measured in terms of content downloaded from other peers, is proportional to his contribution: shared disk space, upload bandwidth, etc. File-sharing is modeled as a non cooperative game where the players seek to maximize their benefit while keeping their contribution to the minimum.

In this paper, we model file-sharing in a P2P network in a game theoretic setting as in the references cited above. Our model however augments the ones used previously in two important ways: extention of the model in the temporal domain; and strategies that are affected by a peer's experience in the network which in turn dynamically affect the strategies of other peers.

We contend that as a rational decision maker, the decision that a peer makes regarding contributing his resources to the system is guided not only by his perception of immediate gain, but his expectation of long term gain resulting from the present choice of actions. A peer is aware of the fact that his contribution may encourage other peers in the network to also contribute, which in the long run will result in his own overall benefit. Thus it is this foesight on the part of the peers that is responsible for the success of a P2P system like Napster, despite the absence of a payment mechanism, and not altruistic behavior on the part of some peers alone. The concept of *long term gain* immediately requires a notion of time, which we incorporate by modeling a P2P network as a dynamic system with multi-act file sharing games.

We now come to the second important thesis of this paper that builds upon the dynamic model mentioned in the previous paragraph. A peer participates in many sessions over a period of time as either a server or a client. He responds to the network by incorporating in his future strategies his experience with the network so far. This experience comprises of the resources he has contributed compared to the benefit he has derived. Since his strategy affects the experience of other peers which in turn affects his own experience, a realistic model of P2P system is a closed-loop system with feedback.[7] Thus a faithful model must take into account this feedback effect and assume such strategies as are affected by experience and learning, rather than assume that peers employ open-loop strategies. These strategies are known as behavioral strategies[8] in game theory literarure.

The following example illustrates the fallacious conclusions that are we are lead to if we only assume open-loop strategies for the peers:

Example 1 Consider the following simple formulation of file-sharing in a 2 player P2P system:

- Benefit of downloading file from peer $(b) = 2$
- Cost of uploading file to peer $(c) = -1$
- Benefit(cost) of not downloading (uploading) $= 0$
- Net utility $u = b + c$

Table 1 represents this bimatrix game. Considering the payoff matrix for peer 1 (row player) Table 1(a), it can be seen that row 2 is his security strategy. Similarly, the security strategy for player 2 (column player), Table 1(b), is the second column. These security strategies correspond to the unique Nash equilibrium pair (not-share, not-share). This equilibrium point is therefore also an admissible solution.[8,9] The desirable strategy pair (share, share) is not an equilibrium point, since a player can increase his pay-off by unilaterally deviating from it.

Even if we assume that this game is played out time and again, *i.e.* a multi-act game, but only assume open-loop *mixed-strategies*, we reach the same equiilibrium point since mixed strategies include pure strategies as a special case. However, if we also include behavioral strategies, the game may be played out as follows: keeping in mind his long term gain, p_1 might decide to take risk and accept p_2's request. Receiving a positive response, p_2 might also decide to honor p_1's request, and so on.

Table 1. Bimatrix game of Example 1

(a) Payoff for p_1

		p_2 share	p_2 not share
p_1	share	1	-1
p_1	not share	2	0

(b) Payoff for p_2

		p_2 share	p_2 not share
p_1	share	1	2
p_1	not share	-1	0

Keeping in mind the foregoing discussion, in this paper, a P2P system has been modelled as a non-cooperatrive multi-act dynamic game with behavioral strategies. The organization of the rest of the paper is as follows: Section 2 models the system with the above characteristics and describes the evolution of the system with time. The concpets of strategy, contribution and benefit are made precise. Section 3 formulates a peer's long term expected benefit as an optimization problem, the solution to which yields his optimal strategy.

2. SYSTEM MODEL

We consider a P2P system with K peers and M contents. In the following we use the words *peers* and *players* interchangeably whichever, is more appropriate to the context. Let $\mathcal{P} = \{1, \ldots, K\}$ denote the set of peers and $\mathcal{C} = \{1, \ldots, M\}$ denote the set of contents. $\mathcal{P}_{-i}$ denotes the set of peers other than i. The system is initialized with the peers having some of these M contents while not having others. The system evolves as successive games involving content exchange are played by the peers. These games are labeled $n = 1, 2, \ldots$. At the beginning of each game, peers request each other for content that they don't possess. Thus a particular peer makes requests and is requested by other peers. Unless stated otherwise, we use $j \in \mathcal{P}$ to denote the requesting peer and $i \in \mathcal{P}$ to denote the serving peer. A peer's requests are either accepted or declined and *vice-versa*. These games are played repeatedly until certain termination criteria are met. The termination criteria are specified later in the paper.

The distribution of these contents among the peers at the beginning of the n^{th} game is represented by the vector $\boldsymbol{y}_n = (y_{1n}^1, \ldots, y_{Kn}^1, \ldots, y_{1n}^M, \ldots, y_{Kn}^M)$ where

$$y_{kn}^m = \begin{cases} 1 & \text{if peer } k \text{ has content } m \text{ before the } n^{th} \text{ game} \\ 0 & \text{otherwise.} \end{cases}$$

The state of peer i at n is represented by $\boldsymbol{y}_{in} = (y_{in}^1, \ldots, y_{in}^M)$. There are 2^{KM} possible $\boldsymbol{y}$'s. These represent the system states.

It is assumed that at each time interval peers in the network have information regarding the contents possessed by other peer. Let $r_{jn}^m \in \{i : y_{in}^m = 1,\ i \in \mathcal{P}_{-j}\}$ denote the peer requested by peer j for content m at time n among those peers who have this particular content. $r_{jn}^m = 0$ is assumed to mean that j does not request content during interval n despite not having it. Then the vector of requests of for peer j is

$$\boldsymbol{r}_{jn}(\boldsymbol{y}_n) = (r_{jn}^m,\ m \in \mathcal{C}_{jn})$$

and that for the system is

$$\boldsymbol{r}_n(\boldsymbol{y}_n) = (\boldsymbol{r}_{1n}, \ldots, \boldsymbol{r}_{Kn}),$$

where $\mathcal{C}_{jn} = \{m : y_{jn}^m = 0\} \subseteq \mathcal{C}$.

From a state $\boldsymbol{y}_n$, peers generate requests $\boldsymbol{r}_n(\boldsymbol{y}_n)$. Some of these requests are accepted, while others are rejected. The decisions regarding accepting/rejecting requests are determined by the peers' strategy to requests, which we denote *serving strategies.* These are discussed in the following section. Thus depending on the requests and serving strategies, the system moves to state $\boldsymbol{y}_{n+1}$

2.1. Serving strategies

In this subsection and the next, we use boldface type to denote random quantities and ordinary typeface to denote the values taken by the random quantities. In a state y, a peer is faced with a number of requests from other peers for uploading files from among the ones that he possesses, depending on the state the system is in. The peer might respond to each requests in either of two ways: accept it or reject it. These choices denote the *action* that might be taken. More precisely, let $\boldsymbol{u}_i^{jm}(s)$ be a random variable that denotes the action taken by i in response to j's request for content m at when the state is y. Let $\boldsymbol{u}_{iy}^{jm} = 1$ when the request is accepted, while $\boldsymbol{u}_{iy}^{jm} = 0$ when the request is refused. Let

$$P\{\boldsymbol{u}_{iy}^{jm} = 1\} = p_{iy}^{jm}, \text{ and} \tag{1}$$

$$P\{\boldsymbol{u}_{iy}^{jm} = 0\} = q_{iy}^{jm} = (1 - p_{iy}^{jm}). \tag{2}$$

That is, the probability of i accepting j's request for content m when the system is in state y is p_{iy}^{jm}. The corresponding probability of rejecting this particular request is q_{iy}^{jm}.

A strategy π_{iy}^{jm} is a probability distribution on the alternative actions of i in state y to the pertinent request. Thus π_{iy}^{jm} prescribes values for p_{iy}^{jm}, and therefore also q_{iy}^{jm}; subject to the system being in state y. Let $\pi_{iy} = (\pi_{iy}^{jm})$, the tuple of strategies of peer i to all requests in state y.

Definition 2.1. *The history up to time at which the nth action is taken is*

$$H_n \triangleq (y_1, r_1, u_1, y_2, r_2, u_2, \ldots, y_{n-1}, r_{n-1}, u_{n-1}, y_n)$$

Consider the sequence of strategies $(\pi_{i1}, \pi_{i2}, \ldots, \pi_{in})$

Definition 2.2. *The strategy*[10] *for player i is the sequence $\Pi_i = (\pi_{i1}, \pi_{i2}, \ldots, \pi_{in})$ such that $\pi_{in}(\boldsymbol{H}_n)$ is a tuple of response strategies to requests in state y if the last element of H_n specifies $y_n = y$.*

In view of the above definition, the probability of $\boldsymbol{u}_{in}^{jm}$ taking values 1 and 0 at n, is given respectively by $P_{\pi_{in}}\{\boldsymbol{u}_{in}^{jm} = 1\}$ and $P_{\pi_{in}}\{\boldsymbol{u}_{in}^{jm} = 0\}$.

If the randomization between alternative actions for time period $n+1$ does not depend on the randomization for periods 1 through n, then the definition of strategies given above corresponds to the *behavioral strategies* of game theory.[8, 9] Also note that Def. 2.2 admits strategies that are not Markov. In words, a strategy for player i permits the nth decision to depend on more of the past history than only present state y_n.

2.2. State Transition Probabilities

Given the state of the system at the n^{th} game $y_n : y_{jn}^m = 0$, and $r_{jn}^m = i$ *i.e.* j does not possess m and requests i for it; the serving strategy of i toward this particular request (along with the strategies to other requests) determines the transition probability to the next state $\boldsymbol{y}_{n+1}$. Thus

$$P\{\boldsymbol{y}_{j\ n+1}^m = 1 \mid \boldsymbol{y}_{jn}^m = 0, \boldsymbol{r}_{jn}^m = i\} = P_{\pi_{in}}\{\boldsymbol{u}_{in}^{jm} = 1\} \text{ and similarly} \tag{3}$$

$$P\{\boldsymbol{y}_{j\ n+1}^m = 0 \mid \boldsymbol{y}_{jn}^m = 0, \boldsymbol{r}_{jn}^m = i\} = P_{\pi_{in}}\{\boldsymbol{u}_{in}^{jm} = 0\}. \tag{4}$$

Using Bayes rule, from (3) and (4) we get

$$P\{\boldsymbol{y}_{j\ n+1}^m = 1 \mid \boldsymbol{y}_{jn}^m = 0\} = P_{\pi_{in}}\{\boldsymbol{u}_{in}^{jm} = 1\} P\{\boldsymbol{r}_{jn}^m = i\}, \text{ and} \tag{5}$$

$$P\{\boldsymbol{y}_{j\ n+1}^m = 0 \mid \boldsymbol{y}_{jn}^m = 0\} = P_{\pi_{in}}\{\boldsymbol{u}_{in}^{jm} = 0\} P\{\boldsymbol{r}_{jn}^m = i\} + P\{\boldsymbol{r}_{jn}^m = 0\}. \tag{6}$$

Assuming that serving strategies for a particular request, specified by the requesting peer and requested content, are take independently of other requests, and also that the serving strategies of the peers are independent of one another, (5) and (6) determine the transition probability $P\{\boldsymbol{y}_{n+1} = y_{n+1} \mid \boldsymbol{y}_n = y_n, \boldsymbol{r}_n = r_n\}$.

event	$\boldsymbol{r}_1^2$	$\boldsymbol{r}_2^1$	$\boldsymbol{u}_1^{31}$	$\boldsymbol{u}_2^{12}$	$P\{y_{n+1}\|y_n\}$
1	0	1	p_1		$\frac{1}{4}p_1$
2	2	1	p_1	q_2	$\frac{1}{4}p_1q_2$

Table 2. Transition Probabilities

We now consider the requests that peers make to other peers for possible content download. Considering $\boldsymbol{r}_{jn}^m$ as a random variable, it takes values from the set $\{i : y_{in}^m = 1,\ i \in \mathcal{P}_{-j}\}$ with probability $P\{\boldsymbol{r}_{jn}^m = i\}$. These request probabilities are in general different for different contents depending on how important content m is to peer j. Also, it is plausible that there is some correlation in $\boldsymbol{r}_j^m(n)$ from one game to the next. For instance if i declines j's request at for content m at n, j might request i again in the $(n+1)^{th}$ game, or might choose to request some other peer. However, to keep the analysis simple, we make the following assumptions:

- The random variables $\boldsymbol{r}_{jn}^m$ are assumed to be independent from one time interval to the next.
- $P\{\boldsymbol{r}_{jn}^m = i_l \mid \boldsymbol{y}_{jn}^m = 0,\ \boldsymbol{y}_{i_1 n}^m = 1, \ldots, \boldsymbol{y}_{i_L n}^m = 1\} = \frac{1}{L},\ l = 1, \ldots, L.$ *i.e.*, among the L peers known to have the content that peer j is interested in, a request is sent to any one of these L peers with equal probability.
- $P\{\boldsymbol{r}_{jn}^m = 0\} = P\{\boldsymbol{r}_{jn}^m \neq 0\} = \frac{1}{2}$. That is, despite not having content m, it is requested/not-requested with probability $\frac{1}{2}$.

Having specified the serving strategies and the manner in which peers make requests, the transition probability $P\{\boldsymbol{y}_{n+1} = y_{n+1} \mid \boldsymbol{y}_n = y_n\}$ between the states is completely specified.

Example 2 Consider a 2 content, 2 user P2P system. The current state of the system is $y_n = (10\ 01)$ that is, each peer has one of the contents, but not the other. The transition probabilities of moving from this states to $y_{(n+1)} = (11\ 01)$ listed in Table 2. In these state transitions, peer 2 requests content 1 and his request is granted. Peer 1 either does not request content 2 (transition 1), or he requests, but is refused (transition 2). The rest of the possible states and their associated transitions are illustrated in Figure 1. As mentioned before, and discussed in detail the next section, the two events denoted by events 1 and 2 affect the serving strategies of the peers in the state y_{n+1}, and therefore the transition probabilities out of that state. Thus the system is *not* a Markov chain. Also note that although the the possible transition probabilities have been summed in Figure 1, this has only been done for the purpose of readability. Indeed the two transitions lead to different states which share a common content distribution y_{n+1}.

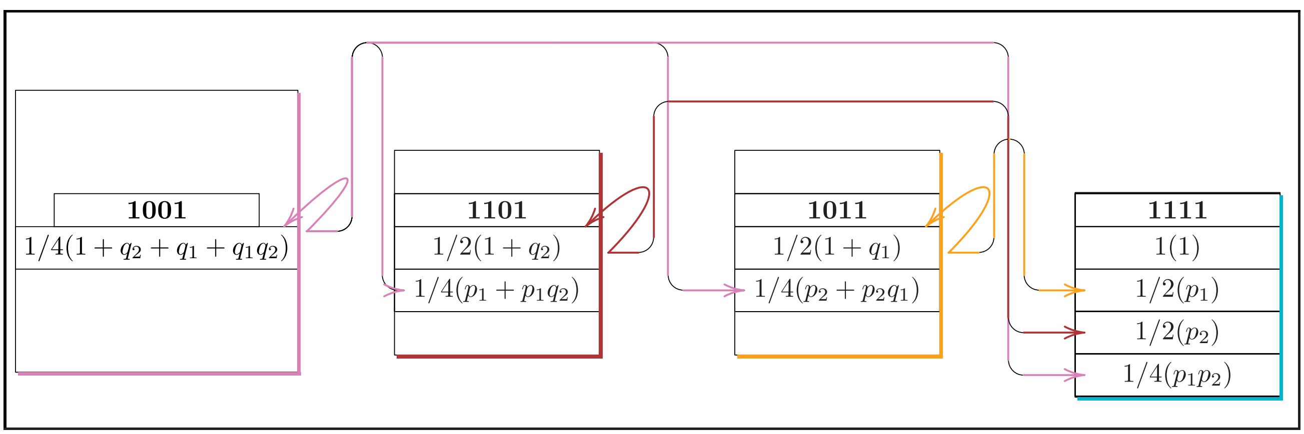

Figure 1. Evolution of the system. The expression at the end of an arrow is the transition probability between states connected by it. The topmost entry in bold is the system state $\boldsymbol{y}$

3. OPTIMAL STRATEGIES

Starting from an initial state, the terminal state of interest $\boldsymbol{y}_j^f$ for an individual peer j would be when he possesses all the contents; $\boldsymbol{y}_j^f : \boldsymbol{y}_{jn} = (1, \ldots, 1)$. From a system-wide perspective, the terminal state is $\boldsymbol{y}^f : \boldsymbol{y}_n = (1, \ldots, 1)$, *i.e.* all peers possess all the contents. The second case includes the first. However among the histories leading to $\boldsymbol{y}^f$, the ones that reach $\boldsymbol{y}_j^f$ earlier than $\boldsymbol{y}_i^f, i \in \mathcal{P}_{-j}$ are of interest to peer j. Thus j wants to minimize $N : \boldsymbol{y}_{jN} = \boldsymbol{y}_j^f$. Following the discussion in Sec. 1 regarding how a peer's strategies might affect the strategies of other peers; qualitatively, he can achieve his goal by adopting a strategy where he accepts a relatively higher number of requests, since this will encourage other peers also to accept a higher number of requests. However accepting a higher number of requests will incur a higher cost in terms of upload bandwidth, processing power, etc. Thus there is a trade off between reaching the desired goal in a short time and expending resources. A strategy is thus optimal in the sense that it strikes a desired balance between these conflicting requirements.

As noted in Definition 2.2, the strategy depends not only on the present state, but on the entire history leading up to it. For mathematical tractability, we assume that peer i adopts the same serving strategy π_{in}^{jm} for requests from all peers $j \in \mathcal{P}_{-i}$ and contents in $\mathcal{C}'_{jn}$, where $\mathcal{C}'_{jn} = \{m : y_{jn}^m = 1\} \subseteq \mathcal{C}$. π_{in}^{jm} however changes with n according to i's experience with the network. Let a_{in} denote the number of contents uploaded to requesting peers in interval n and b_{in} denote the number of contents downloaded (successful requests). Thus $\beta_{in} = \frac{b_n - a_n}{b_n + a_n}$, $\beta_{in} \in [-1, 1]$ is a measure of peer i's experience with the network during n. We assume that i updates his strategy according to the following equation:

$$p_{i\ n+1} = \begin{cases} p_{in} + (1 - p_{in})\alpha_i\beta_{in} & \text{if } \beta_{in} > 0, \\ p_{in} + p_{in}\alpha_i\beta_{in} & \text{if } \beta_{in} < 0, \\ p_{in} & \text{if } a_{in} = b_{in} = 0 \end{cases} \tag{7}$$

where $\alpha_i \in [0\ 1]$. A favorable response from the network prompts i to accept more requests, while a poor response elicits a strategy with reduced contribution. Note that (2) is in fact the first derivative of p_{in} and governs the dynamics of the serving strategy. Also, holding α_i fixed, and $\beta_{in} > 0$, the incremental change p_{in} in lower when p_{in} is close to 1 than when it is close to 0. This is desirable, since if a peer is already contributing a lot, a positive response will have little effect on his strategy compared to when his strategy is conservative. Alternatively, a negative experience will have little effect if p_{in} is small . The parameter α_i controls the magnitude of the incremental change in the peer's strategy in response to his immediate experience. Thus an optimal value of α_i needs to be determined in order to meet the desired objective.

An element of *feedback* is introduced by adopting (2), since peer i's strategy affects the experience of other peers' experience in subsequent intervals and consequently their own strategies, which in turn affects i's experience, *ad-infinitum*.

Starting from a given state $\boldsymbol{y}_1$, there are a number of histories $H_{N_k}, k = 1, 2, \ldots$ (2.1) that lead to the final state $\boldsymbol{y}_{N_k} : \boldsymbol{y}_{jN_k} = \boldsymbol{y}_j^f$. With each such state, there is associated the probability of reaching that state, P_k which is the product of transition probabilities between states $(\boldsymbol{y}_{1_k}, \ldots, \boldsymbol{y}_{N_k})$. Given a strategy π_{in}, these transition probabilities are given by (5) and (6). Let w_j^{lm} be the number of contents uploaded by j in going from state $\boldsymbol{y}_l$ to state $\boldsymbol{y}_m$. Then $W_{jk} = \sum_{(l,m) \in H_{N_k}} w_j^{lm}$ is the total number of contents uploaded in history H_{N_k}, which is indicative of the cost incurred. Starting from an initial content distribution $\boldsymbol{y}_1$ and a particular choice of α_j, the expected cost in reaching the terminal state is

$$\mathbb{E}\ W_j = \sum_k P_k W_{jk}. \tag{8}$$

The expected number of steps to reach the terminal state $\mathbb{E}\ N_j$ is

$$\mathbb{E}\ N_j = \sum_k P_k N_{jk}. \tag{9}$$

From (2) we see that P_k is a polynomial in α_j of degree $N_{jk} - 1$. A peer may then choose a value of α to reach a desired balance between the two quantities $\mathbb{E}\ N$ and $\mathbb{E}\ W$.

4. RESULTS

In this section, we present simulation results for the ideas developed in the preceding sections. A 3 peer ($K = 3$), 2 content ($M = 2$) system has been considered. The initial state is $\boldsymbol{y} = [01\ 00\ 01]$; peer1, the peer of interest starts with no content. The system is symmetrical in the other two peers, in the sense that each has one of the contents but not the other. For peer 1, a state $\boldsymbol{y} = [1 * *1 * *]$ is a terminal state, where $*$ denotes either 0 or 1. Two cases for the initial probabilities of acceptance $p_i = 0.5,\ i = 1, 2, 3$ have been considered.

Figure 2 shows the expected number of steps to termination and expected downloads (incurred cost) for peer 1 as a function of α_1. The system is initialized with, $p_i = 0.5,\ i = 1, 2, 3$ *i.e.* the peers are neutral to risk initially, and $\alpha_2 = \alpha_3 = 0.3$. The plot verifies the intuitive expectation that the number of steps to termination and incurred cost should be inversely related. Figure 3 shows similar results for $\alpha_2 = \alpha_3 = 0.8$. Thus given the α values for the other peers, peer 1 may choose a suitable value for α_1 in order to achieve the desired trade-off between the cost and steps to termination. Figures 4 and 5 illustrate quantities of interest to peer 1 as a function of α_1 and α_2. α_3 is held fixed at 0.5. This technique can be extended to any number of peers and contents.

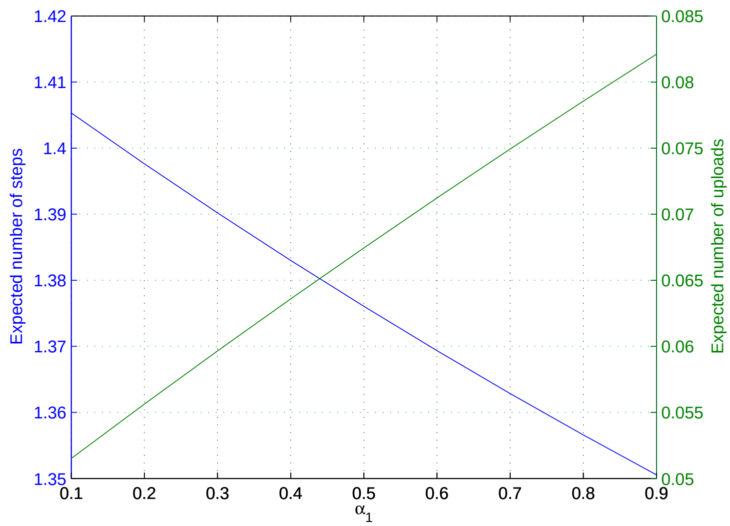

Figure 2. Expected number of steps to termination and expected cost with α_1. $\alpha_2 = \alpha_3 = 0.3$

5. CONCLUSIONS AND FUTURE WORK

In this paper we have modeled content sharing over a P2P network as a non-cooperative multi-act dynamic game with behavioral strategies. The system is represented by states that depict the distribution of contents among the peers. The decisions regarding a peer's contribution to the system are governed by his strategy. The state representation enables us to formulate this dynamic system in a Markov chain like structure with the transition probabilities among states given by the peers' strategies. A peer adapts his strategy according to his experience with in the network, and his expected net benefit from participating in the system. By enumerating all possible state transitions we find the expected number of steps to the desired terminal state and the expected cost as a function of the peers' strategies. This enables us to determine the strategy that a peer may adopt in order to achieve the desired trade-off between these two quantities.

The method of enumerating all possible histories is however computationally inefficient. Techniques from Markov Decision Process (MDP) with due modifications maybe employed in order to ease the computational burden. This paper has concentrated on determining optimal strategies for one peer. The task of determining system-wide optimal strategies $\alpha_1, \ldots, \alpha_K$ for all the peers is yet to be considered.

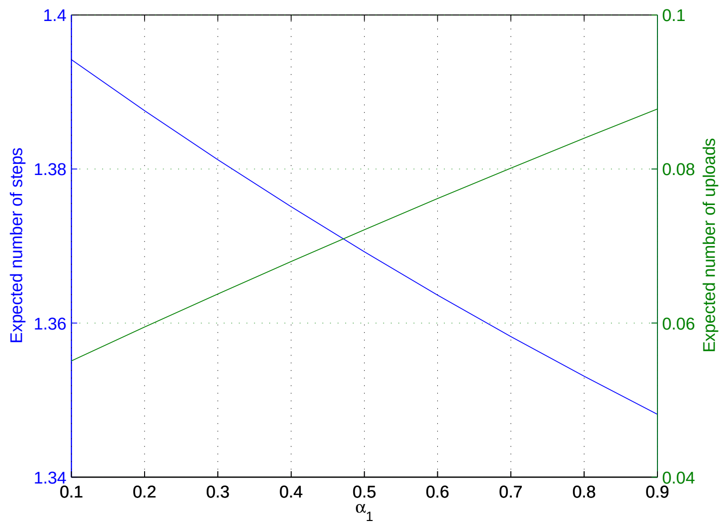

Figure 3. Expected number of steps to termination and expected cost with α_1. $\alpha_2 = \alpha_3 = 0.8$

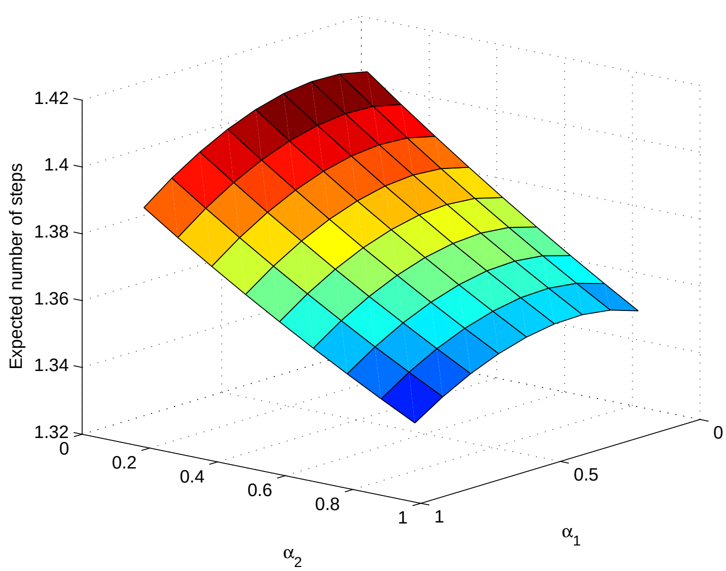

Figure 4. Expected number of steps to termination with α_1 and α_2. $\alpha_3 = 0.5$

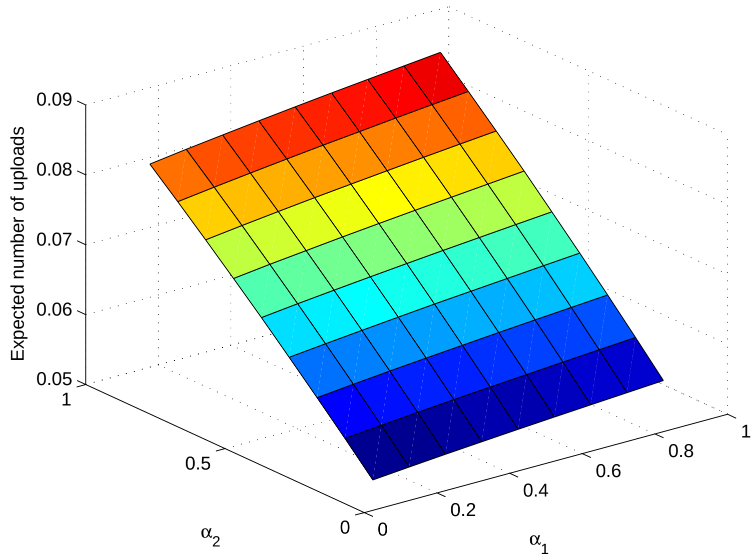

Figure 5. Expected uploads with α_1 and α_2. $\alpha_3 = 0.5$

REFERENCES

1. T. Nguyen and A. Zakhor, "Distributed video streaming over the internet," in *MMCN*, (San Jose, California), January 2002.
2. D. Xu, M. Hefeeda, S. Hambrusch, and B. Bhargava, "On peer-to-peer media streaming," in *ICDCS*, (California), July 2002.
3. S. Saroiu, P. K. Gummadi, and S. D. Gribble, "A measurement study of peer-to-peer file sharing systems," in *Multimedia Computing and Networking*, 2002.
4. E. Adar and B. Huberman, "Free riding on gnutella." First Monday, 2000.
5. P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge, "Incentives for sharing in peer-to-peer networks," in *Proc. 2001 ACM Conference on Electronic Commerce*, 2001.
6. C. Buragohain, D. Agrawal, and S. Suri, "A game theoretic framework for incentives in p2p systems," in *Third International Conference on P2P Computing (P2P 2003)*, (Linkoping, Sweden), 2003.
7. A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*, Oxford University Press, New York, NY, 1995.
8. T. Basar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, SIAM, Philadelphia, PA, 2 ed., 1999.
9. R. B. Myerson, *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA, 2002.
10. D. P. Heyman and M. J. Sobel, *Stochastic Models in Operations Research*, vol. 2, Dover Publications Inc., Mineola, NY, 2004.