Stochastic Optimization for Content Sharing in P2P Systems

Mihaela van der Schaar, *Senior Member, IEEE*, Deepak S. Turaga, *Member, IEEE*, and Ritesh Sood

***Abstract*—Available resources in Peer-to-Peer (P2P) systems depend strongly on resource contributions made by individual peers. Empirical data shows that in the absence of incentives, a majority of the participating peers do not contribute resources. Modeling interactions between individual peers is often difficult as the number of peers in the system can be very large, and the relationships among them can be very complex. In this paper, we propose a new solution for P2P systems, where peers upload and download content to and from the contributing peers based on agreed-upon/determined sharing rates. We propose a P2P solution that deters free-riders by imposing constraints on participating peers; specifically, a peer is allowed access to new content only as long as its own content contribution exceeds an adaptively set threshold. The constraints are enforced either by a central authority (e.g., a tracker) or by a decentralized coalition of peers in a swarm, social network, etc. We derive optimal upload policies for the peers given their estimated future download requirements and their previous contribution (credit) to the other peers. Our results show considerable improvement in the cost-benefit tradeoff for peers that deploy such an optimal policy as compared to heuristic upload policies. We also propose mechanisms based on which the coalition of peers can provide incentives or penalties to participating peers to adjust their policies such that the availability of content and/or number of peers contributing content is maximized.**

***Index Terms*—Content sharing, optimal upload policies, resource contributions in P2P networks.**

I. Introduction

CONTENT distribution through Peer-to-Peer (P2P) networks has recently gained widespread popularity as it provides a distributed framework that is ideal for the dissemination of large files such as multimedia data and software programs, without relying on a dedicated infrastructure (content servers, networks, etc.). Deploying a set of dedicated servers for the purpose of distributing content to many users is inefficient and prone to failures and congestion. Another advantage of P2P networks is their scalability, as available resources scale with demand [15]. While file sharing (downloading) is the predominant activity on existing P2P networks, telephony, live broadcast and streaming multimedia P2P applications are also emerging [1], [2], [5]. Furthermore, P2P networks have also been used to provide support for distributed directory services, storage, and grid computation.

While earlier research on P2P networks has focused mainly on system design, efficient streaming techniques and traffic measurements for real-time applications [21], [22], several other important issues pertaining to the evolutionary dynamics (i.e., upload-download) of these networks have also been recently examined [3], [14], [17], [18]. Since the functioning of a P2P network depends on the resources (content) contributed by participating peers, it is essential to ensure that peers who benefit from the network also contribute to it [2], [3]. Central to the successful evolution and stable behavior of a P2P system is the issue of providing incentives (and sometimes imposing constraints) to the participating peers to contribute resources, i.e., upload content, as otherwise "the tragedy of commons" becomes inevitable [20]. This topic has been extensively investigated [3]–[9], [14] and proposed solutions range from microeconomic payment mechanisms to providing differentiated services for peers based on their prior contributions to the P2P network.

Currently, two types of P2P system architectures exist [15]: a partially centralized system, referred to as a tracker-based system, and a fully decentralized system, referred to as a trackerless system. In tracker-based systems, a "tracker" (a central directory server) facilitates communication between peers, while in trackerless system [25] relevant content is tracked by the "swarm" of peers in a distributed manner. A "swarm" is a group of connected peers that share one or more content items in which they are interested. In the remainder of the paper, we will use the terms "swarm" and "coalition" interchangeably. Note that coalitions can be formed using rules that consider the specific characteristics of the peers, i.e., upload bandwidths, content availability, content interests etc. In this paper we do not investigate such coalition formation and peer matching, but details on these issues may be obtained from [31]. Nevertheless, both system architectures need to deploy explicit upload-download policies and provide incentives to participating peers to prevent them from becoming "free-riders." For instance, several current P2P systems employ a simple differentiated services model that stipulates that the amount of data a peer is allowed to download is proportional to the amount of data that it has already uploaded [3]. More sophisticated differentiated service models can be constructed by taking into account upload/download capabilities, computational resources, disk-space, and value of the content as perceived by the various users. For instance, the value of content [4] may increase based on whether it is semantically important (popular) or if its availability is restricted to only certain limited locations on the P2P network. However, in such

Manuscript received March 18, 2007; revised July 30, 2007. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Anna Hac.

M. van der Schaar is with the Department of Electrical Engineering, University of California, Los Angeles, CA 90024 USA (e-mail: mihaela@ee.ucla.edu).

D. S. Turaga is with the IBM T. J. Watson Research Center, Hawthorne, NY 10532 USA.

R. Sood is with the Electrical and Computer Engineering, University of California, Davis, CA 95616 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TMM.2007.911288

systems (e.g., BitTorrent [3]), no centralized resource allocation strategy exists and each peer is responsible for maximizing its own utility (download rate, percentage of accepted downloads, etc.). In [3], peers do this by downloading from whomever they can and deciding which peers to upload to using a "tit-for-tat" strategy. Hence, peers decide to cooperate by contributing their resources and not to cooperate by denying download requests temporarily from other specific peers. These upload and download algorithms are not part of the protocol, however, as mentioned in [3], effective upload and download policies are necessary to ensure the welfare of each peer and also the overall benefit and growth of the network. Game theory has also been deployed to model peer behavior in noncooperative P2P environments [24]. However, this work is mainly aimed at showing that, at the Nash equilibrium, free-riders are prevented from joining the P2P system, and not at developing operational content sharing strategies for the various peers. In contrast with issues such as equilibrium characterization, we aim at describing and improving the *dynamic* behavior of peers. Our main interest lies in understanding the emergent behavior of interacting peers that use simple adaptation rules to sequentially adjust and optimize their upload/download strategies. The outcome of these interactions need not converge to equilibrium, i.e., perpetual adaptation of strategies may persist. These interactions may involve a large number of peers, which continuously adapt their strategies based on several parameters such as their content demand etc. Hence, from the perspective of any single peer, the P2P environment is *nonstationary*. This unfortunately inhibits strategic optimization and convergence to equilibrium through repeated interactions. In this paper we focus on developing solutions that allow individual peers to determine the amount of resources they should contribute (e.g., uploaded content) to other peers in the P2P coalition or swarm [16], [28] to ensure that they can download content from other peers when desired in a dynamic, nonstationary setting. At the same time, expending resources that significantly exceed the necessary amount for download, is wasteful and needs to be avoided. Additional challenges arise due to the stochastic nature of a peer's download requirements. Note that this problem has many similarities with inventory management [11] especially with uncertain model distributions, and we will discuss how some of those solutions may be adapted in this case.

Content and resource sharing among peers may be modeled as a noncooperative interaction between rational decision makers. This model enables the study of simultaneous upload and download policies for a subset of network peers that maximize their own payoff noncooperatively. For instance, in [8], a system with multiple peers sharing multiple content resources has been modeled as a Markov chain, and the evolution of the system is analyzed numerically. However, this numerical method is beset with difficulties that arise due to the exponential growth of the number of states and the attendant increase in the required computational power and memory, thereby necessitating a sub-optimal solution.

In this paper, we consider an alternative approach where, instead of tracking interactions among sets of peers, we view the content upload-download of individual peers solely in terms of their interaction with an abstract "environment" (that physically includes other peers in a swarm). To prevent free riders, we assume that participation in a swarm of peers is allowed only if each peer agrees to maintain certain content sharing ratios, i.e., upload versus download fraction [16], [28]. These content sharing ratios are set based on the upload and download behavior of individual peers. Enforcement is performed either by a tracker, or collectively by the coalition of peers in the distributed tracker-less scenario.

We focus on developing optimal upload-download policies for individual peers in terms of the amounts of content (resources) they contribute to the coalition, and content they receive in return, given their specific content sharing ratio. Specifically, we study the problem of determining optimal upload strategies for individual peers, given estimates of their future content needs (download requirements), previous contributions, and their sharing ratio. We formulate this as a stochastic optimization problem and determine a closed form solution. The derived optimal upload strategy depends on the peer's intrinsic characteristics, and peer tunable parameters such as its desired service level and its level of *altruism*. We measure service level in terms of the expected shortfall of content that a peer experiences in response to its demand (measured as a fraction of the demand). Peer *altruism* is characterized by the willingness of the peer to upload more content than necessary, while satisfying its requirements. Our proposed solution for upload-download is not dependent on the specific peer selection policy as in [5], [8]. As we will see in the following sections, this abstraction of individual peer interacting with the outside enables us to derive optimal sharing policies analytically, under a general set of assumptions, thereby allowing the deployment of the proposed solutions across a variety of existing P2P systems and applications.

We then study the design of optimal sharing constraints such that peers are provided incentives to upload and download content and maximize the overall P2P system benefit. It should be noted that this paper is aimed at studying possible resource exchange strategies and system interactions in a general manner, and does not aim at providing a complete implementation framework for a P2P content distribution application. However, the proposed solution can be deployed in a variety of P2P systems [1], [3], [15], [23] to derive optimal upload policies for a variety of applications. We demonstrate the performance of our algorithms by deploying them in a real P2P system architecture that is based on BitTorrent [26]. We use this implementation to present results on the content upload-download behavior, clustering of peers based on upload rate as well as altruism, and built-in deterrents to free-riders. The main contributions of this paper are as follows.

- *Optimal resource exchange strategies for individual peers*. We derive optimal upload policies for each peer given its estimated future download requirements, its previous contribution, and its desired level of service.
- *Imposing constraints and providing incentives*. We determine what incentives, in terms of the sharing fraction, can be provided to individual peers based on their upload/download behavior to maximize the overall coalition benefit (total amount of uploaded and downloaded content). Our solution differs from other incentive mechanism

research for P2P systems [3], [9], [16], which are mainly aimed at deterring free-riding. Importantly, in our system, the P2P system performance depends not only on the peers fulfilling their upload requirements, i.e., not being free-riders or partial free-riders, but also the amount of data they actually upload/download[1]. These constraints and incentives are designed based on the specific peer behavior.

This paper is organized as follows. In Section II, we describe our P2P system model. In Section III, we analytically derive the optimal upload policies for peers given their demand distribution, and the desired level of service. In Section IV, we describe mechanisms for sharing ratio and incentive design. We present several experimental results on a real P2P system deployment in Section V, and conclude in Section VI.

II. System Model for the Investigated P2P System

We model the content upload-download of individual peers solely in terms of their interaction with the "environment" that physically includes other peers in a swarm or coalition. We assume that the peer interaction with this environment is analyzed within the context of control intervals. The duration of a control interval corresponds to the amount of time taken to download an individual content chunk. A "chunk" is a designated segment of the shared file with size determined by the content creator or seeder. The beginning of a control interval corresponds to the download request for the next chunk of content. In terms of our formulation and analysis, the notions of control interval and chunk are conceptually interchangeable, and in the rest of the paper we use the term control interval for ease of explanation.

In this system model, a peer interacts with the system by adjusting the amount of content that it uploads at the beginning of each control interval to provision for its future predicted demand, given its desired level of service. As in [16] and [19], we assume that there are always sufficient requests from other peers to take advantage of the content/resources uploaded by any peer based on its policy. Of course, the validity of this assumption depends on the specific system, its scale, the available content at the various peers, and the presence of storage support for uploaded resources [16], [26]. In this paper, we will only focus on the upload/download policy and not on coalition formation, content negotiation or on content streaming protocols for which multiple other solutions already exist [4]–[6], [13], [15], [23].

Denote the cumulative amount of content uploaded and downloaded[2] by the peer at the start of the nth control interval (the end of the $(n-1)$th control interval, i.e., after $n-1$ chunks have been downloaded) by U_n and D_n, respectively. Let the ratio between the content uploaded and downloaded be denoted as $\alpha_n = (U_n/D_n)$. This ratio characterizes the behavior of the peer until the nth control interval, and is labeled the upload-download fraction or content sharing ratio. The constraint imposed on each peer is such that the peer is required to maintain $\alpha_n \geq \alpha$ for $\forall n \in \mathbb{N}^+$. As will be shown later in this paper, α can be adapted depending on the behavior of the individual peer or in order to maximize the benefit for the coalition of peers. Moreover, in the considered P2P system, we assume that several peers can regularly produce or create new content or regularly acquire royalty-free content, which they can then make available online. Furthermore, there exist service providers such as media companies that often distribute their content, and serve as the initial content seeders, without expecting upload from other peers in response. Hence, the total upload and download in the system do not necessarily need to be in equilibrium for the system to operate efficiently. Let u_n denote the amount uploaded by the peer at the beginning of the nth control interval. Then, the maximum amount of resource d_n that the peer can download in this control interval can be determined as

$$d_n \leq \gamma(U_n + u_n) - D_n \tag{1}$$

where $\gamma = \alpha^{-1}$. Let the state at the beginning of time n be

$$s_n = \gamma U_n - D_n. \tag{2}$$

s_n can be interpreted as the savings accumulated by the peer, and hence (1) can be rewritten as

$$d_n \leq s_n + \gamma u_n. \tag{3}$$

u_n can also be interpreted as the resource production of the peer in the nth control interval, since the peer incurs a cost in terms of upload bandwidth, or other contributed resources. d_n represents the consumption/benefit of the peer in the same interval. We also define the maximum consumption in this control interval as

$$a_n = s_n + \gamma u_n. \tag{4}$$

a_n is analogous to the net purchasing power of the peer in control interval n. Let z_n denote the peer's *desired download* in the control interval n. It is clear that $d_n \leq z_n$. Importantly, z_n is a random quantity, since the peer can only estimate what its future desired download is going to be. The strategy of the peer for resource exchanges is characterized by the following parameters.

- r is the benefit of downloading a chunk. Note that this benefit could be different for various peers, different downloaded content chunks, and can vary over time. For the ease of formulation, we consider r as constant. This is the case when benefit is expressed in terms of bandwidth or storage or computational cost rather than content importance, and a fixed number is assigned to it per chunk. The proposed formulation can be easily modified to consider content-adaptive, time-varying benefit. Hence, the net benefit from download for this peer, in the nth control interval is $r(\min(a_n, z_n))$.
- c is the per-unit cost[3] of uploading, which can be determined as a composite function of the upload-bandwidth/storage/consumed power etc. As was the case for the benefit, we consider fixed c, and the proposed formulation

[1]The value of the content uploaded/downloaded is not explicitly considered here but it can be easily included in our proposed solution.

[2]Note that different content data may have different size and hence, can take different time periods to be uploaded.

[3]The cost and benefit parameters are determined specifically per peer, to indicate the different value different peers may place on the same underlying content. Additionally, these parameters may be determined on the fly to account for the popularity and quality of content being exchanged. While we do not explicitly investigate algorithms to tune these parameters dynamically, such tradeoffs can be implemented in the considered P2P architecture.

can be easily modified to consider content-adaptive, time-varying costs. The net cost incurred by the peer in the nth control interval is cu_n.

- In case the desired download exceeds the maximum allowed download, p is the per-unit penalty for not having sufficient resources (upload) to fulfill the desired level of downloads. The net penalty is $p(\max\{z_n - a_n, 0\})$. The parameter p controls the service level of the peer, i.e., a high value for p corresponds to high penalties for not uploading the desired amount, thereby requiring the peer to upload enough to support its download needs. We will show later how this parameter can be determined by the peer given its desired service level.
- Conversely, if at the end of the nth control interval, the amount uploaded exceeded the minimum required to satisfy the desired download ($\alpha_{n+1} > \alpha$), the policy leading to this situation is also penalized, as the peer contributed resources without deriving an immediate benefit, thereby wasting resources. The incurred penalty equals $h(\max\{a_n - z_n, 0\})$. Since the current amount uploaded minus the amount downloaded is carried over to the next time interval, the incurred excess cost may be viewed as an insurance against future contingencies. However, when the parameter p is selected to meet the desired service level, the parameter h may be viewed as an altruism parameter. Altruism represents the willingness of the peer to upload more data than required. Hence, an altruistic peer is one that minimally penalizes excess upload, i.e., an "ideal" altruistic peer is one with $h \to 0$, while an "ideal" selfish peer is one with $h \to \infty$ [16]. As will be shown later, this will also have an impact on the service level that the peer can derive from the system.

Note that while the parameters r and c are intrinsic to the peer (e.g., determined by its connection to the network, etc.), the parameters p and h are peer-tunable based on their desired service level, and altruism.

If we denote $x^+ \triangleq \max\{x, 0\}$, the state that the peer evolves to after the nth control interval is

$$s_{n+1} = \max\{s_n + \gamma u_n - z_n, 0\} \triangleq (a_n - z_n)^+ \quad (5)$$

where s_1 is the initial state of the peer. This state could be assigned to the peer by the system, or it could depend on the type of content or resources the peer can contribute to the system as opposed to the other peers. Based on the above model, we define the one-step utility function that needs to be considered by the peer in determining the optimal resource exchange policy

$$J_n(a_n, z_n, u_n) = r(\min\{a_n, z_n\}) - cu_n - p(z_n - a_n)^+ - h(a_n - z_n)^+. \quad (6)$$

Fig. 1 illustrates the dependence of the utility function on these different parameters for one control interval. As shown in the figure, within one control interval, the maximum utility a peer can derive is $(r - \alpha c)z$, which is achieved when the amount uploaded is just enough to enable the desired download, given the system constraints. Also, a penalty p is imposed when $u < \alpha z$, and another penalty h is imposed when $u > \alpha z$. Individual peers expect to stay in the system over an extended time, and

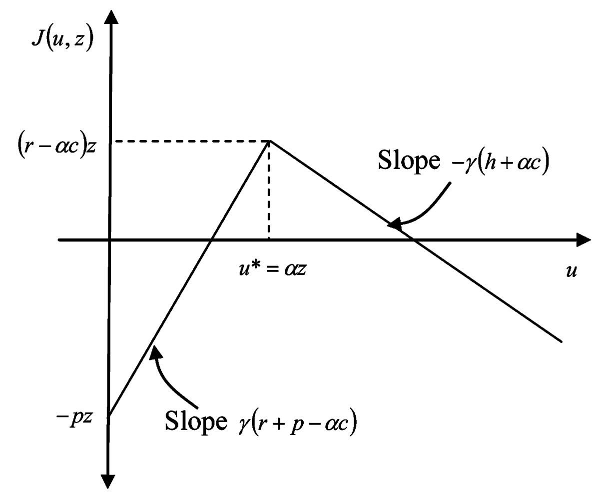

Fig. 1. Utility function given z and $s = 0$.

participate in many upload/download sessions. Hence, instead of considering an instantaneous utility function, the peer upload policy decision is based on the utility function aggregated over some number N of future intervals, thereby provisioning for possible future content downloads in intervals beyond the current one. The parameter N, labeled the planning horizon, controls the tradeoff between short and long-term behavior of the peer. Hence, we can rewrite the utility function as

$$J(N, s_1) = \sum_{n=1}^{N} \beta^{n-1} J_n(a_n, z_n, u_n) \quad (7)$$

where β (with $0 < \beta < 1$) is a discount factor that favors short term gains over long-term gains. The goal of the peer's resource contribution optimization is then to determine appropriate values of u_n to maximize the expectation $\mathrm{E}[J(N, s_1)]$ over all possible realizations of z_n.

III. Optimal Upload Strategies for the Peers

Equation (7) can be solved for u_n using dynamic programming [10], [11], [19]. However, by an algebraic transformation of the model described by equations (4)–(6), it is possible to suppress the temporal component of the model and determine one control interval optimum that is also optimal for the problem with planning horizon N. Such solutions are known as *myopic optima* in operations-research literature [11], [19]. Specifically, from (5)–(7) by substituting $u_n = \alpha(a_n - s_n)$, we get

$$J(N, s_1) = \alpha c s_1 - \sum_{n=1}^{N} \beta^{n-1} \{(\alpha c - r)a_n + p(z_n - a_n)^+ + (r + h - \beta\alpha c)(a_n - z_n)^+\} + \beta^N \alpha c s_{N+1}. \quad (8)$$

The identity $\min\{x, y\} = x - (x - y)^+$ was used while deriving (8). At the end of the planning horizon N, the peer accumulates a total saving s_{N+1}, that is finite upper bounded, i.e., $\exists \xi < \infty$ such that $s_{N+1} \leq \xi$, for any *rational* peer. Since $\beta < 1$ and $\alpha c s_{N+1}$ is upper bounded, then $\beta^N \alpha c s_{N+1} \to 0$ as $N \to \infty$. Hence, for a long enough planning horizon, it is possible to discard the last term in (8). s_1 is the state at the beginning of the planning horizon and is not affected by the upload policy.

Therefore, we can rewrite $J(s_1)$ explicitly as a function of s_1 while eliminating the dependence on N. If we define

$$w(a_n, z_n) = (\alpha c - r)a_n + p(z_n - a_n)^+ + (r + h - \beta\alpha c)(a_n - z_n)^+ \quad (9)$$

then we can obtain

$$\mathrm{E}[J(s_1)] = \alpha c s_1 - \mathrm{E}[\sum_{n=1}^{\infty} \beta^{n-1} w(a_n, z_n)]. \quad (10)$$

Under the assumption that z_n's are independent and identically distributed (i.i.d.) random variables, it can be verified that a_n and z_n are also independent. Therefore,

$$\mathrm{E}[w(a_n, z_n)] = \mathrm{E}[\mathrm{E}[w(a_n, z_n)|a_n]]. \quad (11)$$

Let

$$\begin{aligned} G(a) &= E[w(a, z) \mid a] \\ &= (\alpha c - r)a + pE[(z - a)^+] \\ &\quad + (r + h - \beta\alpha c)E[(a - z)^+]. \end{aligned} \quad (12)$$

From (10), (11), and (12), we have

$$\mathrm{E}[J(s_1)] = \alpha c s_1 - \mathrm{E}\left[\sum_{n=1}^{\infty} \beta^{n-1} G(a_n)\right]. \quad (13)$$

The existence of an optimal policy, corresponding to an optimal net purchasing power in every interval a^*, and the method to determine it are given by the following propositions from stochastic inventory theory [11], [19].

Proposition 1: There exists $a^* \geq 0$ such that $G(a^*) \leq G(a)$ for all $a \geq 0$.

Proposition 2: If $s_1 \leq a^*$, then $a_n = a^* \forall n$ is a feasible and optimal solution.

We omit the proofs here, as these are derived based on [11], [19]. We can then adapt these results from inventory management theory, to determine myopic optimization policies for our problem. Specifically, the optimal a^* can be determined as follows. We first rewrite (12) as

$$G(a) = \int_0^{\infty} (\alpha c - r)a + p(z - a)^+ + (r + h - \beta\alpha c)(a - z)^+ dF_Z(z) \quad (14)$$

where $F_Z(z)$ is the cumulative distribution function of the desired download z. We can show that $G(a)$ is convex in a when $r + p > \alpha c$. Expanding these terms, and taking derivative with respect to a, we can obtain the optimal a^* that minimizes $G(a)$ as

$$F_{\mathbf{z}}(a^*) = \frac{r + p - \alpha c}{r + p + h - \beta\alpha c} \quad (15)$$

with $r + p > \alpha c$. The optimal a^* as determined from (15) is not necessarily unique when $F_Z(z)$ is not strictly increasing with z. In that case, we pick a^* as the smallest value of z that satisfies this condition, i.e.,

$$a^* = \min\left\{z : F_{\mathbf{z}}(z) \geq \frac{r + p - \alpha c}{r + p + h - \beta\alpha c}\right\} \quad (16)$$

to obtain a unique solution.

Since $a_n = s_n + \gamma u_n$ is interpreted as the net purchasing power, Propositions 1 and 2 along with (14) determine that the optimal policy is to maintain the consumption power at the prescribed level given by (16). This is achieved by uploading the quantity u_n during each time period. Note that this result is stronger than the one sought, i.e., minimization of (13).a^* not only minimizes (13), but does so along every sample path $G(a_n)$.

Recall that s_1, the initial state is a constant and is not affected by the upload policy. The optimal strategy in the case when $s_1 > a^*$ is somewhat more involved due to the constraint $u_n \geq 0$ [10], [11], [19]. However, it can be shown that if $s_1 > a^*$, the optimal policy is $u_n = 0$ for $1 \leq n \leq M$ and $a_n = a^*$ for $n > M$, where M denotes the number of control intervals until the state falls below a^*, i.e.,

$$M = \inf\left\{n : \left(s_1 - \sum_{i=1}^{n} z_i\right) \leq a^*\right\}. \quad (17)$$

This means that no upload is needed in the periods up to and including those for which $s_n > a^*$ and the peer should follow thereafter the optimal upload policy under the case $s_1 \leq a^*$. Note that $M \in \mathbb{N}^+$ is a random variable that depends on the initial state and the desired downloads during each period. Based on this and Proposition 2, we can determine the optimal upload policy in the general case as

$$u_n^* = \max\{\alpha(a^* - s_n), 0\} \quad \forall n \in \mathbb{N}^+. \quad (18)$$

In order to illustrate the benefits of the proposed approach, we compare the optimal policy with a heuristically determined policy. The heuristic policy assigns $a_n = \mu$ for all n, where $\mu = \mathrm{E}[z_n]$ is the mean of the desired download. Since the allowed download is at most a_n, the policy $a_n = \mu$ can be considered as the first order approximation of the optimal strategy. Alternative policies can be constructed by taking into account the higher moments of the distribution of z_n. To illustrate the results of the proposed policy, we assume that the download function z_n is distributed according to

a) an exponential probability density function (pdf) with mean μ, i.e., $f_Z(z) = (1/\mu)e^{(-z/\mu)}; z \geq 0$;
b) a uniform pdf with mean μ and spread 2δ. i.e., $f_Z(z) = (1/2\delta); \mu - \delta \leq z \leq \mu + \delta$.

These two distributions correspond approximately with expected content chunk download sizes for multimedia content under different settings. For instance, when peers download and exchange personal multimedia content (content authored by users), the chunk sizes are likely to follow an exponential distribution as in [17], i.e., most chunks having small sizes (short duration, poor quality, etc.), and very few examples with large sizes. As opposed to this, when peers download and exchange more conventional content, e.g., MP3 files for popular music, the chunk sizes are likely to be uniformly distributed. For both these distributions, we have unique solutions for a^* (as $F_Z(z)$ is a strictly increasing function of z) that can be computed as

$$F_Z(a^*) = 1 - e^{\frac{-a^*}{\mu}} = \frac{r + p - \alpha c}{r + p + h - \beta\alpha c} \quad \text{or}$$
$$a^* = \mu \ln\left[\frac{r + p + h - \beta\alpha c}{h + \alpha c(1 - \beta)}\right]$$

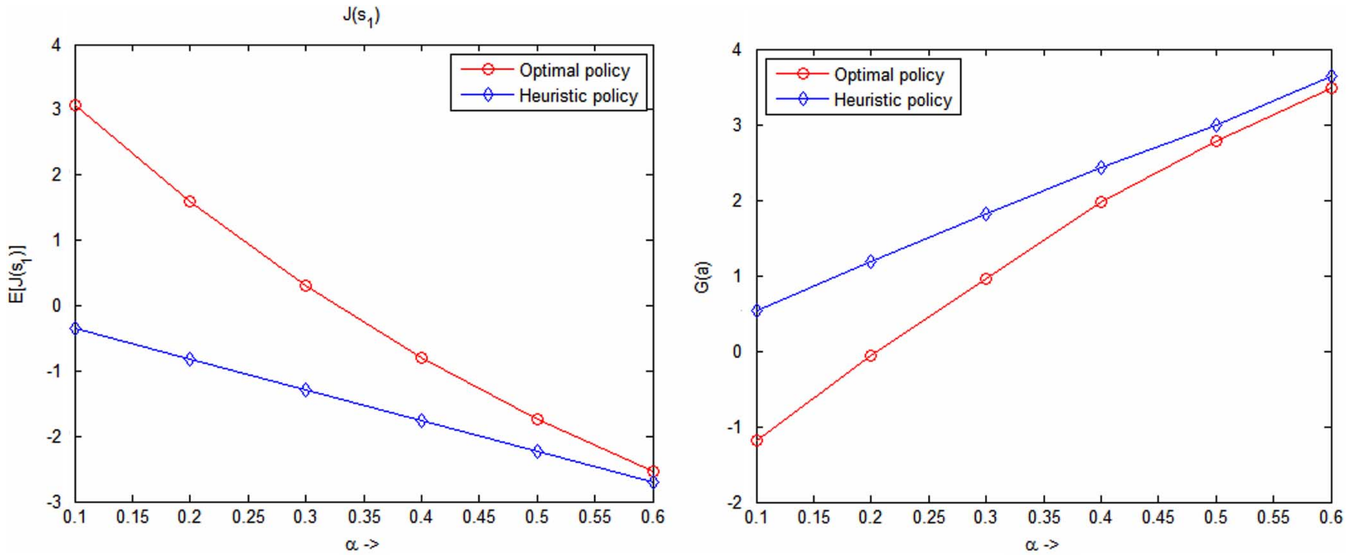

Fig. 2. Expected utility function $E[J(s_1)]$ and $G(a)$ with an exponential pdf for the desired download.

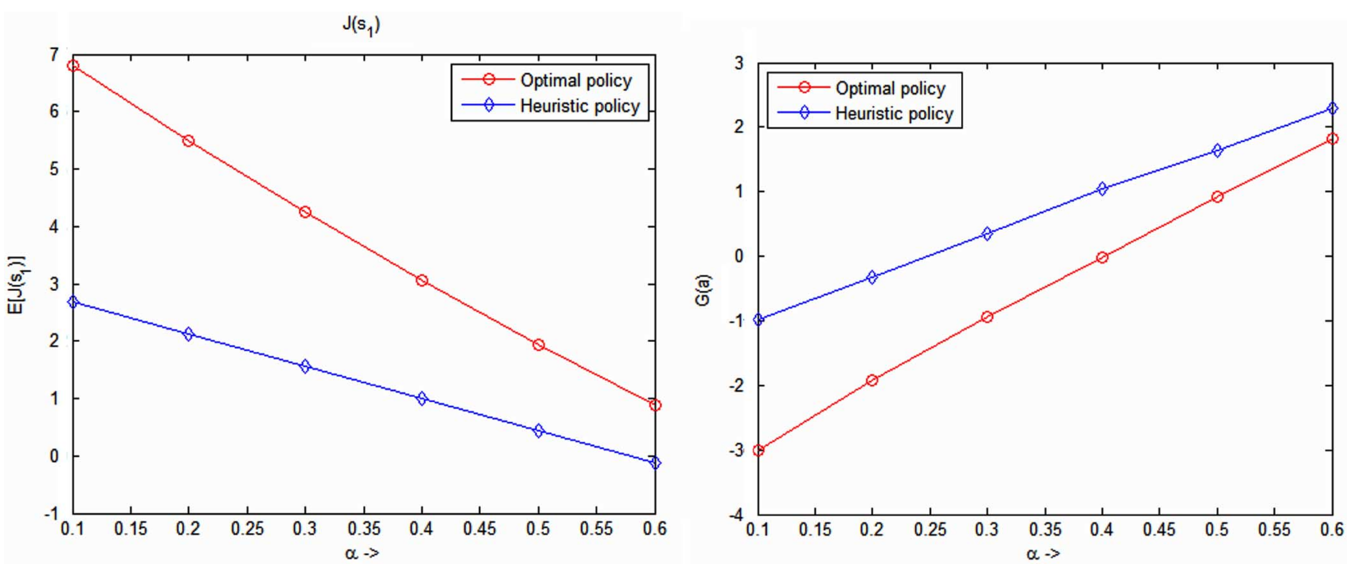

Fig. 3. Expected utility function $E[J(s_1)]$ and $G(a)$ with uniform pdf for the desired download.

for exponential distributions and

$$F_Z(a^*) = \frac{a^* - (\mu - \delta)}{2\delta} = \frac{r + p - \alpha c}{r + p + h - \beta\alpha c} \quad \text{or}$$
$$a^* = \mu - \delta + 2\delta\left[\frac{r + p - \alpha c}{r + p + h - \beta\alpha c}\right]$$

for uniform distributions.[4]

Fig. 2 depicts several results with the parameters $s_1 = 5, \beta = 0.5, r = 1, c = 1.5, h = 0.2, \mu = 5, p = 1.5$[5], and an exponential pdf for the desired download, for several different values of α. We can repeat this experiment to determine the utility function with a uniform pdf for the download function z_n, while keeping unchanged the other parameters, and the resulting functions are shown in Fig. 3.

From Figs. 2 and Figs. 3, we can observe that the selected optimal upload policy leads to significantly higher utility $E[J(s_1)]$ as opposed to the heuristic policy. Furthermore, the advantage of the optimal selection increases with decreasing α. In Section IV-D, we will investigate the impact of the upload policies on the video quality that can be derived by a participating peer.

We now revisit the model to examine the effect of changing the parameters h and p since these parameters are determined by the peer, based on its altruism, and its desired service level. Fig. 4 highlights the effect of changing h on the cumulative upload and download. As expected, with increased h, the amount

[4]With $\mu - \delta \leq a^* \leq \mu + \delta$.

[5]Note that these absolute values of the parameters are used purely as illustration. The relative values of these parameters characterize the nature of the peer, e.g., when $r/c < 1$ it costs the peer more to upload one chunk as opposed to the benefit it derives from one downloaded chunk. In the remainder of this paper we use similar illustrative values for these parameters. These may easily be mapped onto real costs and benefits.

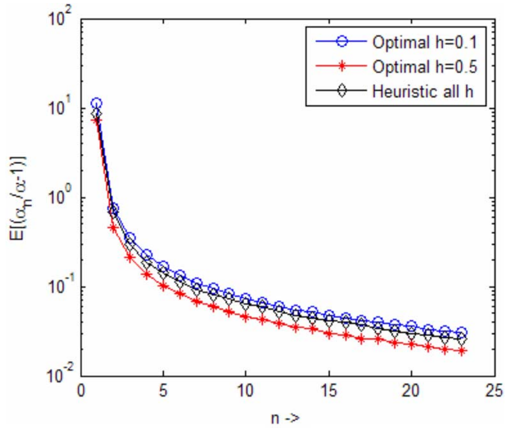

Fig. 4. Effect of changing h on the excess cumulative upload-download fraction $E[(\alpha_n/\alpha) - 1]$.

of over-provisioning reduces, i.e., the ratio of upload to download approaches the required α. Note that the heuristic policy is not affected by the selection of h.

The parameter p can be used by the peer to control the desired level of service. The service level is determined by the amount of data that the peer receives in response to its demand, i.e., whether the peer's demand is met or not, and if it is not met, how much the shortfall is (because peers are assumed to derive benefit $r\min(a^*, z_n)$ even from partial data download). The expected amount of the shortfall in peer demand per control interval, given that a shortfall has occurred, can be determined as

$$E[(z - a^*)^+] = \int_{a^*}^{\infty} (z - a^*) f_Z(z) dz \tag{19}$$

which is $\mu e^{-(a^*/\mu)}$ for an exponential distribution and $(\mu + \delta - a^*)^2/4\delta$ for a uniform distribution (when $\mu - \delta \leq a^* \leq \mu + \delta$). The corresponding probability of observing a shortfall in a specific control interval is

$$P(z > a^*) = \int_{a^*}^{\infty} f_Z(z) dz - 1 - F_Z(a^*). \tag{20}$$

Based on the above, the expected shortfall of data that the peer receives in response to its demand can be computed as $E[(z - a^*)^+]P(z > a^*)$, and we define the desired service level as a fraction

$$\kappa = \frac{E[(z - a^*)^+]P(z > a^*)}{\mu}. \tag{21}$$

Different values of parameter κ correspond to different levels of service, with $\kappa = 0$ corresponding to the highest level of service, and increasing κ corresponding to decreasing service level. Hence, given κ, the parameter p can be determined by the peer using the relationship

$$\mu\kappa = E\left[\left(z - F_Z^{-1}\left(\frac{r + p - \alpha c}{r + p + h - \beta\alpha c}\right)\right)^+\right] \times \left(1 - \frac{r + p - \alpha c}{r + p + h - \beta\alpha c}\right). \tag{22}$$

Since $F_Z(z)$ is monotonically increasing in z (for both exponential as well as uniform distributions), the parameter p can be uniquely computed by the peer. For an exponential distribution, we can rewrite (21) as

$$\kappa = \left(\frac{h + \alpha c(1 - \beta)}{r + p + h - \beta\alpha c}\right)^2 \tag{23}$$

which leads to a solution for p as

$$p = h\left(\frac{1}{\sqrt{\kappa}} - 1\right) + \frac{\alpha c}{\sqrt{\kappa}} - \beta\alpha c\left(\frac{1}{\sqrt{\kappa}} - 1\right) - r. \tag{24}$$

Similarly, for the uniform distribution, we can rewrite (21) as

$$\kappa = \frac{\delta}{\mu}\left[\frac{(h + \alpha c(1 - \beta))}{(r + p + h - \beta\alpha c)}\right]^3 \tag{25}$$

and we can solve for p appropriately.

IV. Dynamic Adaptation of Upload-Download Fraction Based on P2P System Imposed Constraints

The previous section discussed how to determine the optimal upload policy given the constraints imposed by the P2P system (the tracker or coalition of peers) in terms of the upload-download fraction α. In this section, we focus on how the optimal α can be set for different peers.[6] As mentioned earlier, we assume that the goal of the coalition is to maximize the total amount of content chunks uploaded and downloaded by the various peers. However, to be able to determine the appropriate optimal upload-download fraction for each peer, the coalition needs to be able to predict the peer behavior in terms of the amount of data it will upload and download[7] in response to any setting of α.

The peer parameters r and c (benefit and cost) are determined by the peer's connection to the network, its set of available resources, desire for requested content etc. and are known by the system. However, the peer can modify parameters h and p based on its altruistic nature as well as its desired level of service. In general, the peer may not want to explicitly expose these parameters to the coalition. In such cases, the P2P system can infer these parameters from the observed peer behavior. Specifically, the coalition is aware of the optimal upload policy a^*, as this is the net purchasing power the peer attempts to achieve with its upload in every control interval. Furthermore, given that the coalition is aware of the peer's content demand, and to what extent it is satisfied, it also can easily estimate the service level of the peer κ as well as the excess cumulative upload-download fraction $E((\alpha_n/\alpha) - 1)$. Given the estimated κ, a^*, and $E((\alpha_n/\alpha) - 1)$, and observations of peer behavior with changing α, the P2P system can solve a system of (nonlinear) equations to estimate the parameters h and p of the peer.

[6]In this section we include discussion only for peers with an exponential distribution for their demand, it is straightforward to extend this to include peers with a uniform distribution for their demand.

[7]We consider the total upload and download of a peer rather than only the total upload, as it is in the interest of the coalition that each peer is downloading sufficient content, which can subsequently be used by other interested peers. In this way, over time, each peer can become a seeder for various other peers for different types of content.

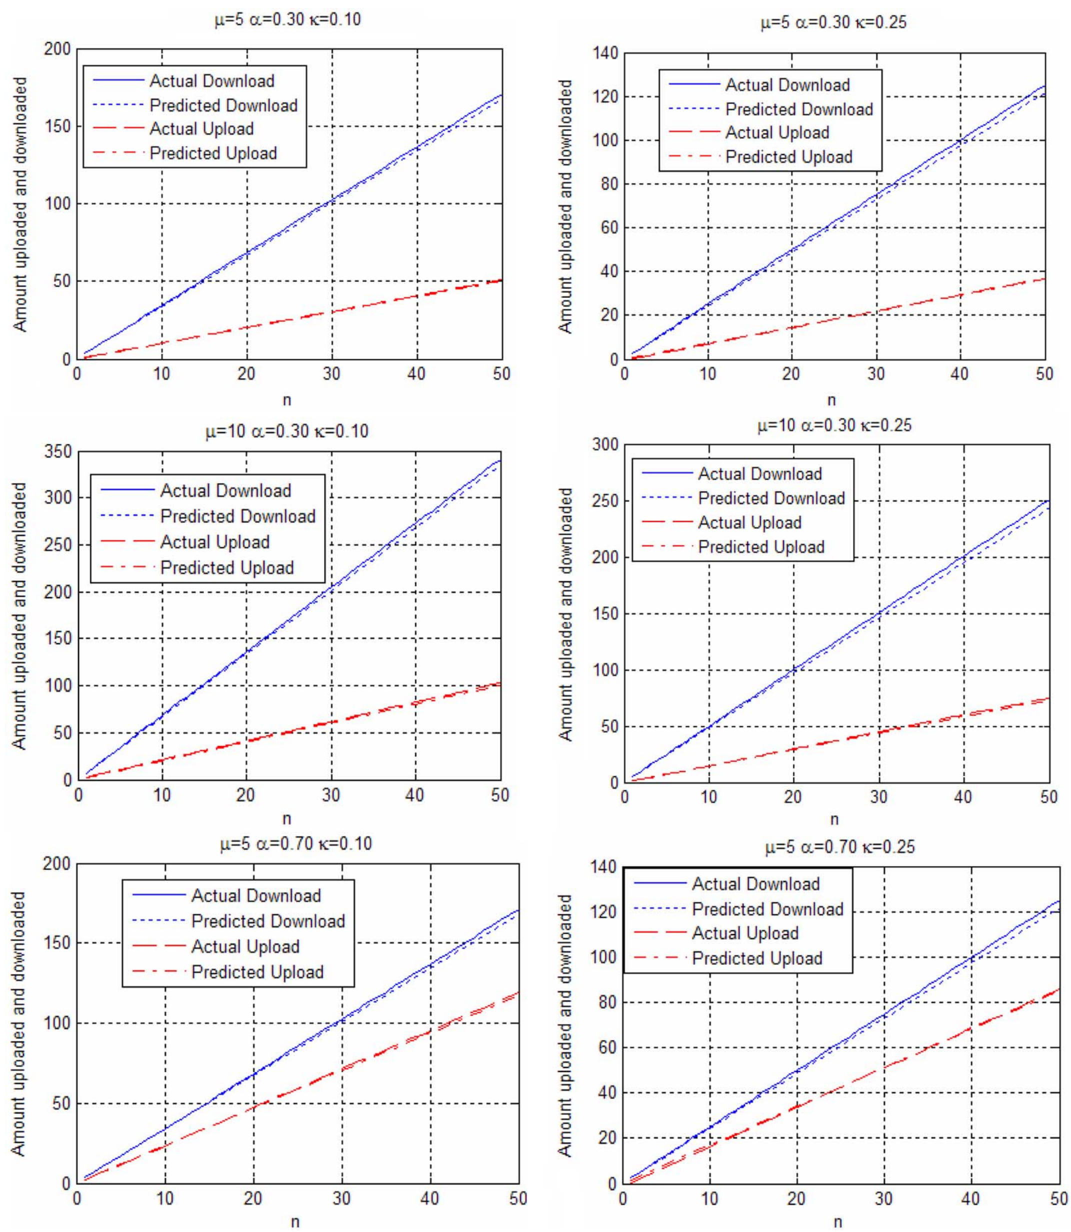

Fig. 5. Effect of mean demand, service level, and upload factor on total download and upload of that peer across various control intervals.

A. Estimation of the Peer's Download and Upload

Given these estimates of the peer's parameters, the expected amount of data downloaded $E[D]$ by the peer in any one control interval, in the steady state, can be computed as

$$\begin{aligned} E[D] &= E[D|z > a^*]P(z > a^*) \\ &\quad + E[D\,|\,z \leq a^*](1 - P(z > a^*)) \\ &= a^*P(z > a^*) + (1 - P(z > a^*)) \int_0^{a^*} z f_Z(z) dz. \end{aligned} \tag{26}$$

Hence, a simple estimate for the total amount of data transferred by the peer can be computed as $(1+\alpha)E[D]$. For an exponential distribution of the content size, $E[D]$ can be computed as

$$E[D] = a^* e^{-\frac{a^*}{\mu}} + \left(1 - e^{-\frac{a^*}{\mu}}\right)\left(\mu - (a^* + \mu)e^{-\frac{a^*}{\mu}}\right). \tag{27}$$

Examples of the impact of service level and upload factor on total download and upload, both actual as well as estimated, of a peer is shown in Fig. 5. The results in this figure show that depending on the mean demand as well as the service level parameter (κ) peers upload and download different amounts or content. As expected, peers upload and download more content with increasing μ and decreasing κ. Furthermore, it is clear that actual download and upload behavior closely follows the predicted behavior. In order for the coalition to provide incentives to peers, it needs to estimate the current demand and service level that the peer requires.

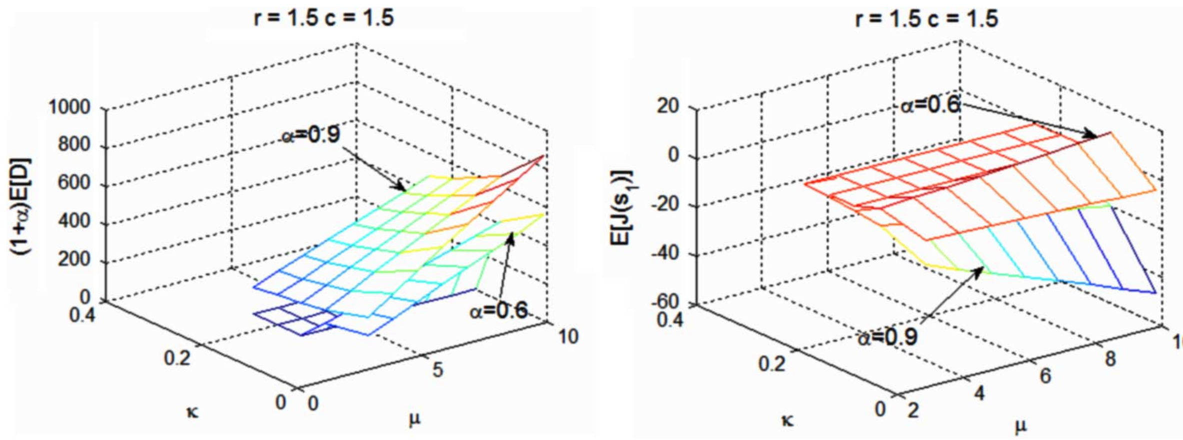

Fig. 6. Example of the coalition benefit (left) and utility function (right) for a particular peer.

B. Peer Response to Incentives

Given the imposed upload-download fraction α, peers determine how much content they are willing to download (mean demand μ) and what level of service they are willing to accept (κ and correspondingly p). Note that this is done to maximize the utility derived by the peer. In order to investigate the relationship between utility and μ and κ, we rewrite the peer's utility as

$$\begin{aligned} &E[J(s_1)] \\ &= \alpha c s_1 - \sum_{n=1}^{\infty} \beta^{n-1}((\alpha c - r)a^* + pE[(z-a^*)^+] \\ &\quad + (r+h-\beta\alpha c)E[(a^*-z)^+]) \end{aligned} \tag{28}$$

$$\begin{aligned} &\text{or, } E[J(s_1)] \\ &= \alpha c s_1 - \sum_{n=1}^{\infty} \beta^{n-1}((\alpha c - r)a^* + pE[(z-a^*)^+] \\ &\quad + (r+h-\beta\alpha c)(a^* - \mu - E[(z-a^*)^+])) \end{aligned} \tag{29}$$

since $E[(a^*-z)^+] + E[(a^*-z)^-] = a^* - \mu, E[(a^*-z)^+] = -E[(z-a^*)^-]$, and $E[(a^*-z)^-] = -E[(z-a^*)^+]$. Thus, we can rewrite $E[J(s_1)]$ in terms of α, μ, and κ as

$$\begin{aligned} &E[J(s_1, \alpha, \kappa, \mu)] \\ &= \alpha c s_1 - \sum_{n=1}^{\infty} \beta^{n-1} \left((\alpha c - r)a^* + \frac{p\mu\kappa}{P(z > a^*)} \right. \\ &\quad \left. + (r+h-\beta\alpha c)\left(a^* - \mu - \frac{\mu\kappa}{P(z>a^*)}\right)\right). \end{aligned} \tag{30}$$

For an exponential distribution, we can show that $\kappa = e^{-(2a^*/\mu)} = (P(z > a^*))^2$ and hence, we have

$$\begin{aligned} &E[J(s_1, \alpha, \kappa, \mu)] \\ &= \alpha c s_1 - \sum_{n=1}^{\infty} \beta^{n-1} \left((\alpha c - r)a^* + \frac{p\mu}{\sqrt{\kappa}} \right. \\ &\quad \left. + (r+h-\beta\alpha c)\left(a^* - \mu - \frac{p\mu}{\sqrt{\kappa}}\right)\right). \end{aligned} \tag{31}$$

Furthermore, we can also substitute for a^* and p to obtain the dependences of the utility function only on α, μ and κ. An illustrative example of the utility function and coalition benefit for a particular peer (with $r = 1, c = 1.5, h = 0.5$) is depicted in Fig. 6.

Note that the amount downloaded and uploaded tends to increase with increasing μ and level of service. Hence, for the coalition to increase its benefit, it needs to induce a peer in this direction through the appropriate selection of upload-download fraction α. However, peers change their desired level of service κ and their desired amount of content only if their expected utility function $E[J(s_1)]$ increases.

From Fig. 6, we see a general trend that the peer utility increases with increasing μ for $\alpha = 0.6$ and higher values of $\kappa(\geq 0.1)$, as both these determine the amount of content that the peer needs to upload to satisfy its demand (increased upload requirements negatively impact the utility). Furthermore, when $\alpha = 0.9$, the peer utility increases with decreasing μ (e.g., for $\kappa = 0.1, 0.01$), thereby leading to reduced overall coalition benefit in terms of the amount of data uploaded and downloaded. Hence, it is not always obvious which parameter selection α will lead to the peer increasing its upload and download.

C. Selection of Optimal Upload-Download Fraction

In general, it may be observed that increasing α makes it more difficult for peers to increase their demand or level of service, while decreasing α causes the coalition benefit to diminish. Hence, the coalition needs to tradeoff these conflicting goals and select α given the type of peers in the coalition. In general, this optimization problem can be written as

$$\begin{aligned} \alpha^{\text{opt}} &= \arg\max_{\alpha}(1+\alpha)E[D(\alpha, \kappa^{\text{opt}}, \mu^{\text{opt}})], \\ \{\kappa^{\text{opt}}, \mu^{\text{opt}}\} &= \arg\max_{\kappa,\mu} E[J(s_1, \alpha, \kappa, \mu)]. \end{aligned} \tag{32}$$

The optimization clearly consists of the two components: an optimization by the peer to maximize its utility, and an optimization by the coalition to maximize the amount of upload and download[8]. Given that these functions are nonconvex, we cannot solve this optimization analytically. Instead, we use a dynamic programming based approach. Note though that such an approach is not likely to be very expensive, as typical peers have

[8]Note that there are several other ways of combining these two contrasting objectives into one utility for optimization, that lead to different behaviors, fairness etc.

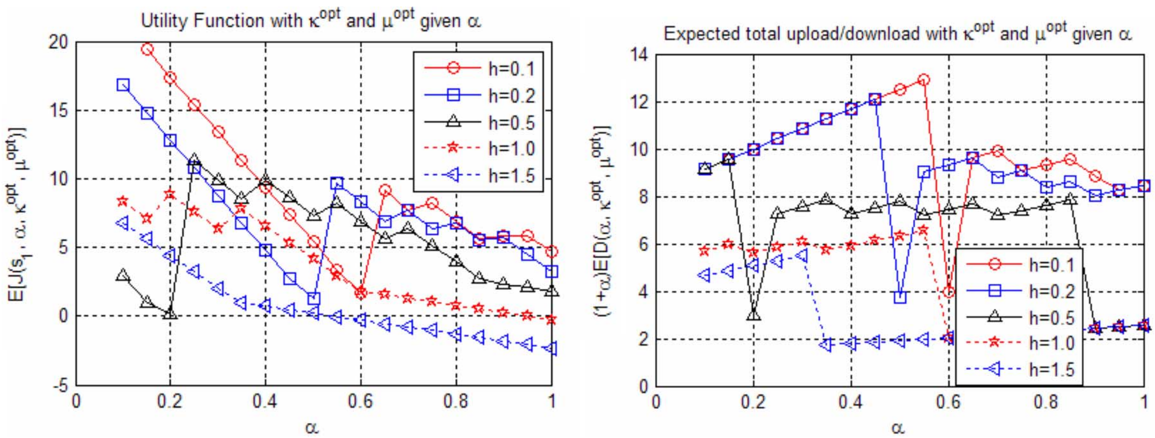

Fig. 7. Optimal utility and upload-download peer behavior given $\alpha(r = 1.5, c = 1)$.

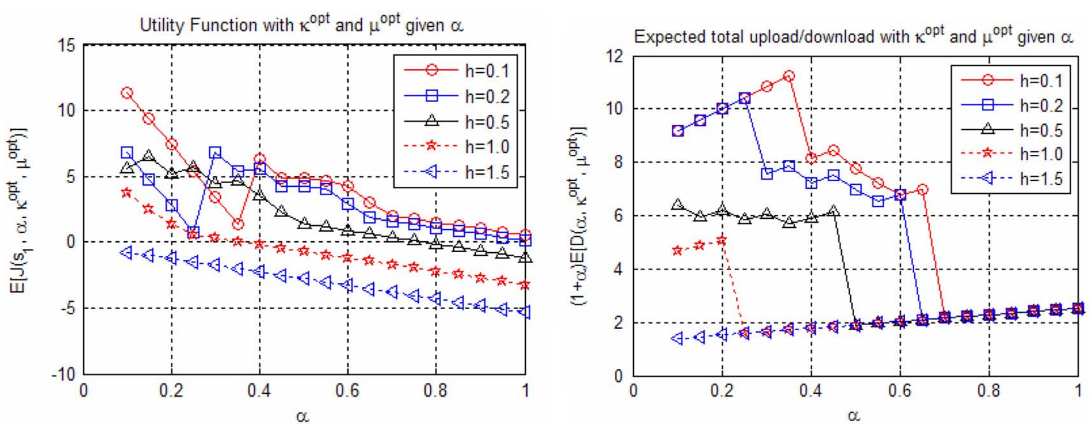

Fig. 8. Optimal utility and upload-download peer behavior given $\alpha(r = 1, c = 1)$.

only a certain discrete set of demand levels, and desired service levels, reducing the size of the search space. Importantly, the tracker (either centralized or the coalition as a whole) can infer the behavior of the peer, and can then determine the parameter $\alpha^{\rm opt}$ appropriately to increase the peer's upload and download, such the overall P2P system benefit is maximized. Examples of the utility function and expected download upload with optimally selected (by peer) κ and μ for different values of α and h are shown in Fig. 7.

In Fig. 7, we consider peers with potential greater benefit of downloading one unit of data as opposed to the cost of uploading one unit of data ($r > c$), and with different altruism parameters. Also, we consider the service level parameter to lie in the range $0.01 \leq \kappa \leq 0.25$ and the mean demand level to lie in the range $3 \leq \mu \leq 10$, with a start state $s_1 = 3$. The first thing to observe is that the optimal peer utility steadily decreases with increasing α until a certain point beyond which it increases suddenly before continuing to decrease. This sudden transition corresponds to the peer moving from one service level to another (increasing κ, thereby reducing service level). At the transition point, it is no longer optimal for the peer to maintain a high service level (small κ) as its upload demands increase beyond any gains it receives from the download, and therefore it needs to reduce its service level. Conversely, the total amount of upload and download increases with increasing α, until the same point (change of service level) beyond which the trend continues. Additionally, the exact location of the transition point changes with the altruism parameter h. As the peer becomes more altruistic $(h \rightarrow 0)$ the transition point moves further to the right,[9] i.e., the peer is more willing to accept a higher upload-download fraction α. We also repeat these experiments for peers with $r = c = 1$ and the results are presented in Fig. 8. Similar observations as for the previous case can also be made in this case. Additionally though it is clear that as the benefit for the peer decreases, its transition point moves quicker to the left, i.e., the peer can tolerate only small increases in α before changing the desired service level. In fact, when the peer altruism parameter is high, the coalition needs to compel (by selecting a high value for α high) the peers to upload a larger fraction of their download to be able to increase the total upload-download. Based on these observations, the $\alpha^{\rm opt}$ that the coalition needs to select, in order to maximize total download and upload, for peers with different behavior can be summarized in Table I. As expected, Table I shows that as it becomes more expensive for peers to upload data, they require a smaller $\alpha^{\rm opt}$ to contribute more resources to the coalition. The exact value of the optimal upload-download fraction depends on the selected ranges of the parameters, and the expected peer behavior. Given $\alpha^{\rm opt}$, peers operate with the optimal upload-download policy in

[9]Note that the first transition point for $h = 1.0$ and $h = 1.5$ are not on the plot, due to the range of selected α values.

TABLE I
SELECTED $\alpha^{\rm opt}$ FOR DIFFERENT PEER BEHAVIOR WITH $\mu^{\rm opt}$ AND $\kappa^{\rm opt}$

	$h = 0.1$	$h = 0.2$	$h = 0.5$	$h = 1.0$	$h = 1.5$
$r = 1.5, c = 1$	0.55	0.45	0.35	0.55	0.3
$r = 1, c = 1$	0.35	0.25	0.1	0.2	1.0
$r = 1, c = 1.5$	0.2	0.15	0.1	0.15	1.0

terms of maximizing their individual utilities. It is also clear that extensions of the above approach that consider a joint function of the coalition benefit and peer utility while determining $\alpha^{\rm opt}$ can also be easily developed.

V. ILLUSTRATION FOR VIDEO P2P SYSTEMS

To illustrate the potential impact of the proposed content sharing policies on multimedia P2P networks, we used a P2P set-up based on [26] (using the PlanetLab platform), which was developed using an instrumented version of the BitTorrent implementation. The deployed implementation is able to record the necessary upload/download statistics required to assess the performance of the proposed algorithms.

We deployed 30 peers (leechers) belonging to three classes corresponding to different maximum upload rates (i.e., 20 kbps, 50 kbps, and 200 kbps). These client classes, determined based on the upload bandwidth, are identical to those used in [26]. Bandwidth limitations were enforced by imposing rate constraints on the deployed PlanetLab nodes. All peers, however, have the same download rate of 1 Mbps, and deploy the optimal upload strategy as designed in this paper. Using a simple linear relationship between the upload and download bandwidths, we may then compute the cost-benefit ratios for these different client classes as 50, 20 and 5, respectively. Note that other ratios, based on the actual utility derived from content, may also be used. Peers exchange video files, at CIF (352×288) resolution and 30 frames per second, encoded in a prioritized manner using H.264 with data partitioning, with a maximum bit-rate of 1 Mbps. During each experiment, a single video file of approximate size 100 Mbits[10] was downloaded by the peers from one content seeder. The video file was partitioned into chunks that have size uniformly distributed around a mean of 100 Kbits. We implemented a simple chunk download algorithm that first downloads chunks having the highest impact on the video quality. We assume that peers do not change their available upload bandwidth, or disconnect before receiving a complete copy of the content. All peers (leechers) join the torrent at the same time, emulating a flash crowd scenario. In this practical experiment, we set the sharing fraction $\alpha = 0.5$ for all peers, and do not modify it dynamically.

We make several very interesting observations based on these practical experiments. First, we implemented a reputation-based system similar to [30] to determine whether peers fulfill their upload requirements according to their determined content sharing fraction. Based on our experiments, we see that in all cases, peers that do not fulfill their sharing fraction are always choked (not provided content) by other participating peers. Furthermore, this choking happens in less than a quarter of the time required for the total download, thereby ensuring that free-riders are not able to download the content fully, and also preventing them from decoding a reasonable video quality. Specifically, for our example video files, the last free-rider is choked within at most 258 s of the start of download (the mean time to download is $\sim$1480 s). Thus, free-riders cannot successfully manipulate the coalition of peers as in e.g., [29]. Second, similar to [26], we observe that given a set of peers with a similar altruism factor, peers cluster based on their upload capacities. This may be explained by the choking policies used by peers in this system, which encourage high peer reciprocation by favoring peers that upload a large amount of content. Hence, the unchoking of peers happens more frequently for peers with similar upload capacities because they are able to reciprocate with high enough rates [26]. Third, if the altruism factor differs for the various peers, we observe significantly different peer-clustering behavior from that observed in [26]. Specifically, peer clustering is no longer determined purely based on the peer upload capacity, but instead, it also strongly depends on the peer altruism characteristics. This effect is stronger in our system because the proposed upload strategy depends heavily on the altruism (h) of the peer. We perform an experiment where 20 peers with the same upload capabilities of 200 kbps interact with each other to download content. Of these peers, 10 have an altruism factor $h = 0.2$ (more altruistic) and the other 10 have an altruism factor $h = 1.5$ (less altruistic), respectively. We observe that the altruistic peers are able to download content almost twice as fast as the nonaltruistic peers (i.e., an average of 1219 s as opposed to 2086 s averaged across all peers over three simulation runs in which the peers fully download the entire content). A similar result has been observed in [29] for a more sophisticated P2P system called BitTyrant. However, as opposed to the BitTyrant system where the strategic behavior of peers is unpunished, in the proposed system, the coalition of peers in the swarm will choke the peers that do not fulfill the agreed-upon upload/download ratio. This imposes greater emphasis on the altruism of peers, and also provides long-term incentives for peers to be altruistic.

Finally, we also determine the amount of decoded content (rate and video quality) for the three different client classes given the parameters $\alpha = 0.5, h = 0.2$, and $r = 1$ for all clients, and costs $c^1 = 50, c^2 = 20$ and $c^3 = 5$ corresponding to their upload-download bit-rate ratio. We consider three different types of video content, represented by the well-known

[10]This corresponds to ~ 100 s of video encoded at this rate. Typical MPEG test sequences are 10 s long, and were looped 10 times to get this approximate size.

TABLE II
EXPECTED DECODED VIDEO QUALITY

	Mean Download	Decoded Video Quality in PSNR (dB) Stefan	Foreman	Coastguard
$r=1,\ c^1=50$	74.2Mbits	31.4	38.2	32.6
$r=1,\ c^2=20$	86.8Mbits	32.8	38.8	33.6
$r=1,\ c^3=5$	95.6Mbits	33.9	39.2	34.4

Stefan, Foreman and Coastguard sequences. The average downloaded content (in terms of bits) and corresponding decoded quality after 1000 control intervals (approximate number of total chunks) are presented in Table II. From these results, it is clear that the client with the lowest cost-to-benefit ratio can download the most chunks, representing not only the base-layer quality, but also the enhancement layers, thereby obtaining the highest decoded video quality. Note that the use of nonoptimal upload-download policies can lead to significantly reduced (up to 30–40% lower) downloads, thereby significantly affecting the video quality.

VI. CONCLUSION

In this paper, we analyze the problem of designing optimal content upload policies for individual peers in the presence of constraints, represented by an imposed upload-download fraction or sharing ratio. We use stochastic inventory management techniques to formulate this as an optimization over a certain planning horizon. Using an algebraic transformation, we derive a closed form solution for the optimal upload given the distribution of a desired future download. We also determine the design of the optimal sharing ratio for a client, in order to provide it incentives such that the total amount of uploaded and downloaded content in the P2P system is maximized. In the discussed P2P system, performance depends not only on the peers fulfilling their upload requirements, i.e., not being free-riders or partial free-riders, but also on the amount of data they actually upload/download. Importantly, our results show that, by imposing constraints on participating peers, we can successfully deter free-riders.

ACKNOWLEDGMENT

The authors would like to thank the associate editor and anonymous reviewers for identifying several limitations in the initial version of our manuscript and also for referring us to current literature in distributed P2P systems. We would also like to thank Prof. L. Zhang, UCLA, for discussions on practical implementations of our algorithm in BitTorrent.

REFERENCES

[1] *Akamai*, [Online]. Available: http://www.akamai.com
[2] E. Adar and B. A. Huberman, “Free riding on gnutella,” *1st Monday*, vol. 5, no. 10, 2000.
[3] B. Cohen, “Incentives build robustness in bittorrent,” in *Proc. P2P Econ. Workshop*, Berkeley, CA, 2003.
[4] C. Gkantsidis and P. Rodriguez, “Network coding for large scale content distribution,” in *Proc. IEEE/INFOCOM 2005*, Mar. 2005.
[5] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai, Distributing streaming media content using cooperative networking Tech. Rep. MSR-TR-2002-37, Apr. 2002.
[6] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, “Splitstream: High-Bandwidth multicast in a cooperative environment,” in *Proc. SOSP*, Oct. 2003.
[7] D. Qui and R. Srikant, “Modeling and performance analysis of bittorrent-like peer-to-peer networks,” in *Proc. Conf. Applications, Technologies, Architectures, Protocols For Comput. Commun.*, 2004.
[8] R. Sood and M. van der Schaar, “Optimal media sharing policies in peer-to-peer networks,” *Proc. SPIE*, 2004.
[9] K. Lai, M. Feldman, I. Stoica, and J. Chuamg, “Incentives for cooperation in peer-to-peer networks,” in *Proc. Workshop Econ. Peer-To-Peer Syst.*, Berkeley, CA, 2003.
[10] D. P. Bertsekas, *Dynamic Programming and Stochastic Control*. New York: Academic, 1976.
[11] D. P. Heyman and M. J. Sobel, *Stochastic Models in Operations Research*. Mineola, NY: Dover, 2004, vol. 2.
[12] J. Byers, M. Luby, M. Mitzenmacher, and A. Rege, “A digital fountain approach to reliable distribution of bulk data,” in *Proc. SIGCOMM*, 1998.
[13] J. Byers, J. Considine, M. Mitzenmacher, and S. Rost, “Informed content delivery across adaptive overlay networks,” in *Proc. SIGCOMM*, 2002.
[14] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, “Chord: A scalable peer-to-peer lookup service for internet applications,” in *Proc. SIGCOMM*, 2001.
[15] S. Androutsellis-Theotokis and D. Spinellis, “A survey of peer-to-peer content distribution technologies,” *ACM Comput. Surveys*, vol. 36, no. 4, pp. 335–371, Dec. 1, 2004.
[16] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge, “Incentives for sharing in peer-to-peer networks,” in *Proc. 2nd Int. Workshop Electron. Commerce, Lecture Notes Comput. Sci.*, 2001.
[17] A. Ali, A. Mathur, and H. Zhang, “Measurement of commercial peer-to-peer live video streaming,” in *Workshop Recent Adv. Peer-To-Peer Streaming*, Aug. 2006.
[18] K. Sripanidkulchai, A. Ganjam, B. Maggs, and H. Zhang, “The feasibility of peer-to-peer architectures for large-scale live streaming application,” in *Proc. ACM SIGCOMM*, 2004.
[19] W. S. Lovejoy, “Myopic policies for some inventory models with uncertain demand distributions,” *Manage. Sci.*, vol. 36, no. 6, Jun. 1990.
[20] G. Hardin, “The tragedy of the commons,” *Science*, no. 16, 1968.
[21] D. Xu, M. Hefeea, S. Hambrusch, and B. Bhargava, “On peer-to-peer media streaming,” in *Proc. Distributed Comput. Syst.*, 2002.
[22] J. Kubiatowicz *et al.*, “Oceanstore: An architecture for global-state persistent store,” in *Proc. ASPLOS 2000*, Nov. 2000.
[23] J. Li, P. A. Chou, and C. Zhang, Mutualcast: An efficient mechanism for content distribution in a peer-to-peer (P2P) network MSR-TR-2004-100, Sep. 2004.
[24] C. Buragohain, D. Agrawal, and S. Suri, “A game theoretic framework for incentives in P2P systems,” in *Proc. 3rd Int. Conf. Peer-To-Peer Comput. (P2P2003)*, Linkoping, Sweden, 2003, pp. 48–56.
[25] C. P. Fry and M. K. Reiter, Really truly trackerless bittorrent CMU Rep., Aug. 2006.
[26] A. Legout, N. Liogkas, E. Kohler, and L. Zhang, “Clustering and sharing incentives in bittorrent systems,” *ACM Sigmetrics*, Jun. 2007.
[27] , [Online]. Available: http://www-sop.inria.fr/planete/Arnaud.Legout/Projects/p2p_cd.html#software
[28] Y. Wang and J. Vassileva, “Trust-based community formation in peer-to-peer file sharing networks,” in *Proc. Web Intell.*, 2004.
[29] M. Piatek *et al.*, “Do incentives build robustness in bittorrent?,” in *Proc. Symp. Networked Syst. Design Implementation*, 2007.
[30] E. Damiani *et al.*, “A reputation-based approach for choosing reliable resources in peer-to-peer networks,” in *Proc. ACM Conf. Comput. Commun. Security*, 2002.

Mihaela van der Schaar (SM'04) is currently an Associate Professor of electrical engineering at the University of California, Los Angeles (UCLA), where she leads the Multimedia Communications and Systems Laboratory (http://medianetlab.ee.ucla.edu). Her research interests span several areas, including multimedia networking, multi-user communication, cross-layer design, resource management in communication and systems, and game theory. She holds 30 granted US patents and has several more pending. She is the editor (with P. Chou) of the book *Multimedia over IP and Wireless Networks: Compression, Networking, and Systems*.

Dr. van der Schaar has been an active participant and chair in the ISO Motion Picture Expert Group (MPEG) international standard since 1999, to which she made more than 50 contributions and for which she received three ISO recognition awards. She was also elected as a Member of the Technical Committee on Multimedia Signal Processing and of the Technical Committee on Image and Multiple Dimensional Signal Processing, both of the IEEE Signal Processing Society. She was an Associate Editor of the IEEE TRANSACTIONS ON MULTIMEDIA and of *SPIE Electronic Imaging Journal*. Currently, she is an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE SIGNAL PROCESSING LETTERS, and *IEEE Signal Processing Magazine e-Newsletter*. She received the NSF CAREER Award in 2004, IBM Faculty Award in 2005 and 2007, Okawa Foundation Award in 2006, Best IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY Paper Award in 2005, and Most Cited Paper Award from *EURASIP Journal Signal Processing: Image Communication* between 2004 and 2006.

Deepak S. Turaga (M'01) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Bombay, in 1997 and the M.S. and Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1999 and 2001, respectively.

He is currently a Research Staff Member in the Exploratory Stream Processing Department, IBM T. J. Watson Research Center, Hawthorne, NY. He was with Philips Research from 2001 to 2003 and with Sony Electronics during 2003–2004. His research interests lie primarily in statistical signal processing, multimedia coding and streaming, machine learning, and computer vision applications. In these areas, he has published over 30 journal and conference papers and one book chapter. He has also filed over 15 invention disclosures and has participated actively in MPEG standardization activities.

Dr. Turaga received the Best IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY Paper Award in 2005. Dr. Turaga is an Associate Editor of the IEEE TRANSACTIONS ON MULTIMEDIA.

Ritesh Sood received the B.Tech. degree in electrical engineering and the M.Tech. degree in in communications and signal processing from the Indian Institute of Bombay, Bombay, India, both in 2003. He is currently pursuing the M.A. degree in the Department of Mathematics and the Ph.D. degree in the Department of Electrical and Computer Engineering, both at the University of California, Davis.

His research focuses on band-limited methods for communications, signal processing, and the extraction and processing of knowledge from large databases. He is also interested in the application of nonlinear and stochastic optimization techniques to wireless communications and peer-to-peer networks.