OPTIMAL UPLOAD POLICIES FOR P2P NETWORKS WITH SYSTEM IMPOSED CONSTRAINTS

Ritesh Sood and Mihaela van der Schaar

University of California, Davis
Department of Electrical & Computer Engineering
Davis, CA 95616

ABSTRACT

One important aspect in Peer-to-Peer (P2P) networks is resource management since the available resources depend on the contributions made by the peers. In the absence of incentives, empirical data shows that a majority of the participants do not contribute resources. In order to deter such free-riders, some P2P systems impose constrains on the participating peers which allow a peer to gain benefit from the system (download content) as long as his resource contribution to it (uploaded content) exceeds a certain threshold. Under these system imposed constraints, we derive optimal upload policies for each peer given his estimated future download requirements and his previous contribution to the system. Our results show considerable improvements in the cost-benefit function by employing the optimal policy compared to heuristic upload policies. Moreover, the proposed P2P model is very general, and can be employed for investigating different applications, usage scenarios etc.

1. INTRODUCTION

Content distribution through Peer-to-Peer (P2P) networks has gained widespread popularity in the recent past. One of the most important reasons for this popularity is that such a distributed content delivery system is ideal for the dissemination of large size data such as multimedia files, CD images, etc. Deploying dedicated servers for this purpose is both inefficient, and prone to failure. Another advantage of P2P systems is their scalability, as available resources scale with demand.

While file-sharing is the predominant activity on existing P2P networks, live broadcast and streaming multimedia applications through P2P networks are also emerging [1, 2]. Moreover, P2P networks have also been used to provide distributed directory services, storage, and grid computations.

While earlier research on P2P systems has focused mainly on system design and traffic measurements, other important issues pertaining to the evolutionary dynamics of these systems have been recently brought forth [3]. Since the successful functioning of a P2P system depends on the resources contributed by the participating peers, it is essential to ensure that peers who benefit from the network also contribute to it. Central to the successful evolution and stable behavior of a P2P system is the issue of providing incentives to the participating peers. This topic has been investigated in [4, 5]. The proposed solutions range from microeconomic payment mechanisms to differentiated services.

This research was supported in part by “IT@intel”

A number of current P2P networks employ a simple version of the differentiated services model by stipulating that the amount of data that a peer is allowed to download is proportional to the amount of data that he uploads. More sophisticated differentiated service models can be constructed by taking into account computational resources, disk-space, and *value* of the content as perceived by the users. Typically, a central agency oversees the transactions between peers and keeps track of the upload/download statistics of the peers. However a peer cannot know in advance his future download requirements, other than in statistical terms. It is therefore necessary that the amount he uploads is sufficient to ensure that he is able to download from other peers. Conversely, expending resources that significantly exceed the necessary amount is wasteful since the accrued benefit is not utilized until later.

In this paper, we determine the optimal upload policies for a peer, given the estimate of his future download requirements, previous contributions, and the constraints enforced by the system. We determine the optimum by formulating a stochastic optimization problem. Under a very general set of assumptions, we arrive at a closed form solution to the optimization problem which is the optimal upload policy. Numerical experiments provide further insight into the system dynamics.

The organization of the rest of the paper is as follows: Sec. 2 describes the model and introduces the parameters that define the optimization problem. The optimal upload policy is derived in Sec. 3. Sec. 4 presents numerical results and Sec. 5 concludes the paper.

2. PEER-TO-PEER RESOURCE SHARING MODEL

Resource sharing between peers may be modeled as a non-cooperative game played by rational decision makers. Such a model enables us to study the upload and download policies of a subset of the peers in the network simultaneously, where each peer tries to maximize his own payoff by making decisions in a non-cooperative manner [4] [6]. In this paper, alternatively, we consider the sharing policies of a single peer and make certain assumptions regarding the rest of the network. The main advantage resulting from this simplification is the tractability of the latter problem. In [4] for instance, a system with multiple peers sharing multiple contents has been formulated as a Markov chain and the evolution of the system is analyzed numerically. 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. Moreover, the former model is more suited to study the system-wide properties of the network; for example, the existence and stability of equilibrium points, etc. The present work on the other hand focuses on the behavior of a single peer *vis-à-vis* the rest of the network. As we will see in the following sections, this approach enables us to arrive at optimal sharing policies in a purely analytical manner under a very general set of assumptions.

We assume that the system is analyzed in discrete time intervals. The relevant quantities are observed at each time interval and the appropriate controls are applied. For our system, the time step might be the duration of a content exchange, or any other suitable interval of time. The applicable control is the amount of upload. We assume that the control is applied at the beginning of the time period and the download happens thereafter. For example, the upload can be done during off hours to free up the upload bandwidth during working hours. We also assume that the requested upload from other peers always exceeds that prescribed by the policy, *i.e.*, there are always sufficient requests from other peers to enable as much upload as is dictated by the policy.

Denote the cumulative upload and download at the start of the n^{th} time period by U_n and D_n respectively. Let $\alpha_n = \frac{U_n}{D_n}$. The constraint imposed by the system is that the peers are required to maintain $\alpha_n \geq \alpha$ for $n = 1, 2, \ldots$.

Let u_n denote the amount uploaded by the peer during period n. Then, the maximum amount of data that the peer can download in the same period is

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

where $\gamma = \alpha^{-1}$. Let $s_n = \gamma U_n - D_n$ denote the state at the beginning of time n. s_n is interpreted as the accumulated savings by the end of the $(n-1)^{th}$ period. Then (1) can be rewritten as

$$d_n \leq s_n + \gamma u_n. \tag{2}$$

u_n is interpreted as the production, since it incurs a cost (upload bandwidth, etc.). d_n is the consumption in period n. Let

$$a_n = s_n + \gamma u_n. \tag{3}$$

From (2) and (3) $d_n \leq a_n$ which allows an interpretation of a_n as the net consumption power in the current period. Let z_n denote the peer's *desired download* in period n. z_n is a random quantity, since the peer can only estimate at n what his desired download is going to be at m when $m \geq n$. It is clear that $d_n \leq z_n$ for all n. The following parameters further characterize the model:

- r is the benefit of downloading a unit of data. The net benefit in period n is $r(\min\{a_n, z_n\})$.
- c is the per unit cost of uploading, which can be determined as a composite function of the upload-bandwidth, computational cost, etc. The cost in period n is cu_n.
- In case the desired download exceeds the maximum allowed download, p is the per unit penalty for not having sufficient resources to fulfill the desired level of downloads. The net penalty is $p(\max\{z_n - a_n, 0\})$.
- On the other hand, if at the end of the n^{th} period $\alpha_{n+1} > \alpha$; which implies that the amount uploaded exceeded the minimum required to satisfy the desired download. The decision leading to this condition is also penalized. The incurred penalty is $h(\max\{a_n - z_n, 0\})$. Since the current amount uploaded less the amount downloaded is carried over to the next time interval, the excess incurred cost is an insurance against future contingencies. Thus, different values of h yield policies that either favor the short-term benefits more, or the long-term ones.

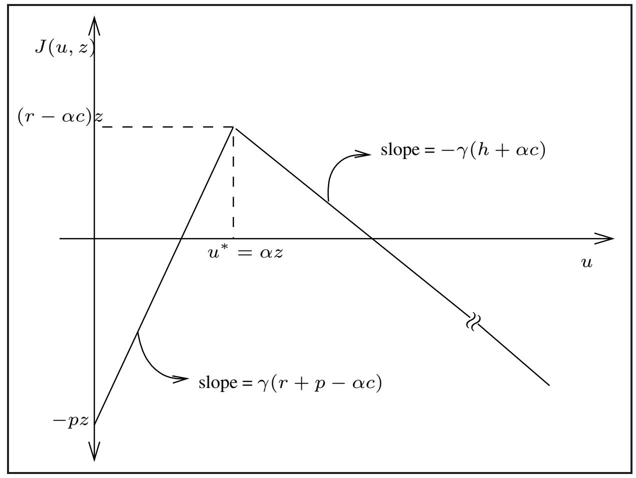

Fig. 1. Cost function for given z and $s = 0$. $u^* = \alpha z$ is the optimal upload quantity. For $u < u^*$ penalty p is imposed while for $u > u^*$, penalty h is imposed.

Denote $\max\{x - y, 0\} \triangleq (x - y)^+$. The state then evolves according to

$$\begin{aligned} s_{n+1} &= \max\{s_n + \gamma u_n - z_n, 0\} \\ &= (a_n - z_n)^+. \end{aligned} \tag{4}$$

s_1 is the initial state assigned to the peer by the system.

Based on the above model, we define the one-step cost function:

$$J_n(a_n, z_n) = r(\min\{a_n, z_n\}) - cu_n - p(z_n - a_n)^+ - h(a_n - z_n)^+. \tag{5}$$

Fig. 1 illustrates the one period cost-function. A peer expects to stay in the the system over an extended period of time, and participate in many upload/download sessions. Also, as noted before, the desired download in the future is a random quantity. Thus, the cost function that is used to determine the optimal policy is

$$J(N, s_1) = \sum_{n=1}^{N} \beta^{n-1} J_n(a_n, z_n), \tag{6}$$

where $0 < \beta \leq 1$ is the discount factor and signifies the fact that although long-term gains are important, short term gains are preferred. N is the planning horizon. The optimization problem is then to maximize $\mathbb{E}\ J(N, s_1)$ over all possible realizations of z_n by choosing appropriate values of u_n.

3. OPTIMAL UPLOAD POLICIES

Equation (6) can be solved for the optimal upload policies u_n using dynamic programming [7, 8]. However by an algebraic transformation of the model (4)-(6), and some other assumptions, it is possible to suppress the temporal component of the model and find such one-period optima which are also optimal for the problem with planning horizon N. Such solutions are known as *myopic optima* in operations-research literature [7].

From (4), (5) and (6) and substituting $u_n = \alpha(a_n - s_n)$,

$$\begin{aligned} J(N, s_1) &= \alpha c s_1 - \sum_{n=1}^{N} \beta^{n-1}\{(\alpha c - r)a_n + p(z_n - a_n)^+ \\ &\quad + (r + h - \beta\alpha c)(a_n - z_n)^+\} + \beta^N \alpha c s_{N+1} \end{aligned} \tag{7}$$

The relation $\min\{x, y\} = x - (x - y)^+$ has been used in writing (7). s_{N+1} is the savings at the end of the planning horizon. Since $\beta < 1$, and if $\alpha c s_{N+1}$ is upper bounded $\beta^N \alpha c s_{N+1} \to 0$ as $N \to \infty$. Thus, for a long enough planning horizon, it is possible to drop the last term on the right hand side of (7). s_1 is the state at the beginning of the planning horizon and is not affected by the upload policy; which only affects the terms inside the summation. Therefore, we write $J(s_1)$ explicitly as a function of s_1 while dropping the dependence on N.
Define

$$w(a_n, z_n) = (\alpha c - r)a_n + p(z_n - a_n)^+ + (r + h - \beta\alpha c)(a_n - z_n)^+. \tag{8}$$

Then

$$\mathbb{E}\ J(s_1) = \alpha c s_1 - \mathbb{E}\ \sum_{n=1}^{\infty} \beta^{n-1} w(a_n, z_n). \tag{9}$$

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

$$\mathbb{E}\ w(a_n, z_n) = \mathbb{E}\ \{\mathbb{E}\ w(a_n, z_n)|a_n\}. \tag{10}$$

Let

$$\begin{aligned} G(a) &= \mathbb{E}\ w(a, z)|a \qquad (11) \\ &= (\alpha c - r)a + p\mathbb{E}(z - a)^+ + (r + h - \beta\alpha c)\mathbb{E}(a - z)^+. \end{aligned}$$

From (9) and (10),

$$\mathbb{E}\ J(s_1) = \alpha c s_1 - \mathbb{E}\ \sum_{n=1}^{\infty} \beta^{n-1} G(a_n). \tag{12}$$

The existence of an optimal policy (a^*) and the method to determine it are given by the following propositions.

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

Proof: ref. [8]

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

Proof: ref. [8]
a^* is obtained as the solution to

$$F_{\mathbf{z}}(a^*) = \frac{r + p - \alpha c}{r + p + h - \beta\alpha c}\ , \tag{13}$$

where $F_{\mathbf{z}}(z)$ is the probability distribution function of the desired download z and $r + p > \alpha c$. a^* as determined from (13) is not necessarily unique. The non-uniqueness can be resolved in the following manner:

$$a^* = \min\{z : F_{\mathbf{z}}(z) \geq \frac{r + p - \alpha c}{r + p + h - \beta\alpha c}\}. \tag{14}$$

Since $a_n = s_n + \gamma u_n$ and is interpreted as the net purchasing power (Sec. 2), Prop. 1, Prop. 2 along with (14) say that the optimal policy is to maintain the consumption power at the prescribed level given by (14). This is achieved by uploading the necessary quantity u_n during each time period. Note that this result is stronger than the one sought, *i.e.* minimization of (12): a^* not only minimizes (12), but does so along *every sample path* $G(a_n)$.

Recall that s_1, the initial state is a constant and not affected by the upload policy. The optimal policy in the case when $s_1 > a^*$ is somewhat more involved due to the constraint $u_n \geq 0$. Due to space limitations, we only mention the result here. The reader is referred to [7, 8] for details.

Let M denote the number of periods until s_n falls below a^*, *i.e.*

$$M = \inf\left\{n : s_1 - \sum_{i=1}^{n} z_i > a^*\right\}. \tag{15}$$

Note that $M \in \{1, 2, \ldots\}$ is a random variable depending on the initial state and the desired downloads during each period.

Proposition 3 *If $s_1 > a^*$, the optimal policy is $u_n = 0$ for $n = 1, \ldots, M$ and $a_n = a^*$ for $n = M + 1, M + 2, \ldots$.*

Prop. 3 prescribes not to upload in the periods up to and including those for which $s_n > a^*$, and follow the policy under the case $s_1 \leq a^*$ thereafter.

From Prop. 2 and Prop. 3, the optimal upload policy in the general case is:

$$u_n^* = \max\{\alpha(a^* - s_n), 0\} \quad n = 1, 2, \ldots. \tag{16}$$

4. RESULTS

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 $\mu = \mathbb{E}\ z_n$ 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 to the optimal. Alternative policies can be constructed by taking into account the higher moments of the the distribution of z.

Other than the existence of a distribution function, no constraints have been imposed on z_n in arriving at (14). For the purpose of illustration, we assume the download function z_n distributed according to the (1) Exponential distribution with mean μ, and (2) Uniform distribution with upper and lower limits b_u and a_u respectively. The values of the model parameters with r as the reference are: $\frac{c}{r} = 1.5$, $\frac{p}{r} = 1.5$, $\frac{h}{r} = 0.2$ and $N = 20$. The results were obtained by averaging over 1000 runs.

Fig. 2 verifies the claim of Prop. 1 which states that a^* *minimizes* $G(a)$ along every sample path. z_n is drawn from the exponential distribution with $\mu = 5$. Correspondingly, a^* *maximizes* $J(s_1)$ for all values of α. It is seen that with increasing α, the benefit obtained by using the heuristic policy comes close to that achieved by the optimal. We therefore conclude that for these parameters, when α is high enough, a peer might save the computational effort by using the heuristic policy mentioned above. However, the benefit of using the optimal policy increases with decreasing α. Fig. 3 illustrates the optimality of a^* for z_n uniformly distributed in $[0\ 10]$. Similar conclusions can be drawn in this case as above.

The cumulative upload and download at the end of period $(n-1)$ are related by $\frac{U_n}{D_n} = \alpha_n$, under the constraint $\alpha_n \geq \alpha$. Then $\epsilon_n = \mathbb{E}\ \frac{\alpha_n}{\alpha} - 1$ denotes the relative amount uploaded in excess of the required minimum. As mentioned is Sec. 2, while a high value of ϵ implies excess cost incurred in the present, this excess cost helps meet the download requirements in the future. This trade-off

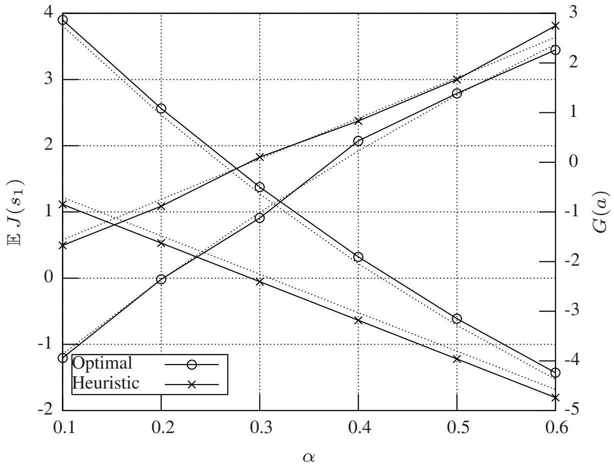

Fig. 2. The rising curve with the y-axis on the right is $G(a)$. $G(a^*) < G(a')$ for each value of α. Correspondingly $\mathbb{E}\ J(a^*) > \mathbb{E}\ J(a')$ which is the falling curve with y-axis on the left.

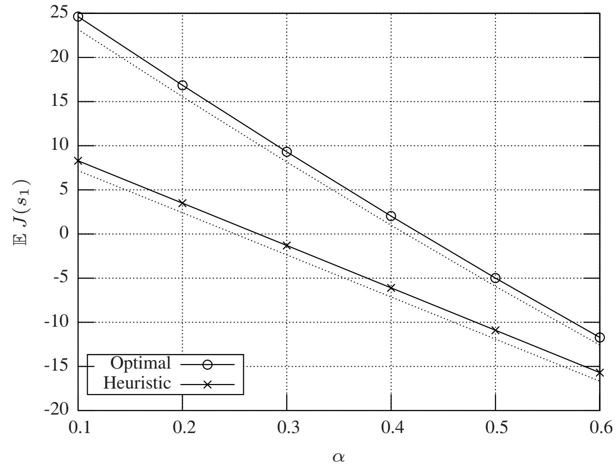

Fig. 3. $\mathbb{E}\ J(s_1)$ with uniformly distributrd z_n.

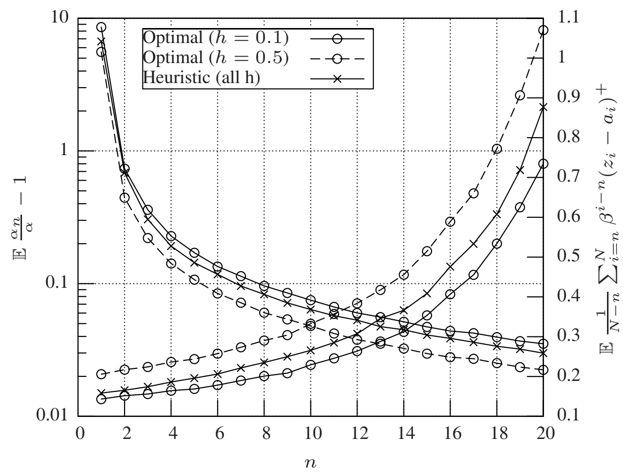

Fig. 4. The falling curve is ϵ_n with y-axis on the left. The rising curve is δ_n. As the penalty for overprovisioning increases, optimal policy dictates upload which is just enough for one period. This however results in decreased long-term benefit.

is illustrated in Fig. 4, where

$$\delta_n = \mathbb{E}\ \frac{1}{N-n} \sum_{i=n}^{N} \beta^{i-n}(z_i - a_i)^+$$

measures the expected mean unsatisfied download from current period to the end of the planning horizon. In Fig. 4 we see that ϵ_n decreases with increasing n. This is because as n approaches the end of the planning horizon, we might not make as much provision for the future as at the beginning of the planning horizon. Correspondingly, the expected mean unsatisfied demand δ_n for the remaining time intervals increases as $n \to N$. Fig. 4 also verifies that a higher value of parameter h results in a policy that favors short term benefits to long term benefits.

5. CONCLUSIONS

In this paper we have analyzed the problem of optimal content upload policies in the presence of network imposed constraints. A stochastic optimization problem over an extended planning horizon is formulated. Algebraic transformation of the original problem enable a closed form solution which yields the optimal upload policy. Numerical experiments verify the optimality of the solution and exhibit substantial improvements in the cost-benefit function as compared to a heuristic policy. Further insight is gained into the system as regards the effect of the optimization parameters on the resulting solution. By choosing appropriate values for the parameters, we arrive at policies that trade-off short term gains for long term ones, and conversely.

6. REFERENCES

[1] 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, Microsoft Research, Apr. 2002.

[2] M. Castro, P. Druschel, A-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "Splitstream: High-bandwidth multicast in a cooperative environment," in *SOSP'03*, Lake Bolton, New York, Oct. 2003.

[3] Dongyu Qui and R. Srikant, "Modeling and performance analysis of bittorrent-like peer-to-peer networks," .

[4] Ritesh Sood and Mihaela van der Schaar, "Optimal media sharing policies in peer-to-peer networks," in *IS&TSPIE Symposium on Electronic Imaging*, Aug. 2004.

[5] Richard T. B. Ma, C. M. Lee, John C. S. Lui, and K. Y. Yau, "A case of providing incentives for p2p networks," .

[6] Philippe Golle, Kevin Leyton-Brown, Ilya Mironov, and Mark Lillibridge, "Incentives for sharing in peer-to-peer networks," in *Proc. 2001 ACM Conference on Electronic Commerce*, 2001.

[7] D. P. Bertsekas, *Dynamic Programming and Stochastic Control*, Academic Press, New York, 1976.

[8] Daniel P. Heyman and Matthew J. Sobel, *Stochastic Models in Operations Research*, vol. 2, Dover Publications Inc., Mineola, NY, 2004.