

Incentivizing Upload Capacity in P2P-VoD Systems: A Game Theoretic Analysis

Weijie Wu¹, John C.S. Lui¹, and Richard T.B. Ma²

¹ The Chinese University of Hong Kong
{wjwu, cslui}@cse.cuhk.edu.hk

² Advanced Digital Science Center, UIUC and National University of Singapore
tbma@comp.nus.edu.sg

Abstract. Free riding has long been a serious problem in peer-to-peer (P2P) systems due to the selfish behavior of individual users. To conquer this problem, a key design issue of the P2P systems is to appropriately incentivize users to contribute resources. In P2P Video-on-Demand (VoD) applications, content providers need to incentivize the peers to dedicate bandwidth and upload data to one another so as to alleviate the upload workload of their content servers. In this paper, we design a simple yet practical incentive mechanism that rewards each peer based on its dedicated upload bandwidth. We use a *mean field* interaction model to characterize the distribution of number of peers in different video segments, based on which we characterize the content providers' uploading cost as a function of the peers' contribution. By using a game theoretic framework, we analyze the interaction between a content provider's rewarding strategy and the peers' contributing behaviors and derive a unique Stackelberg equilibrium. We further analyze the system efficiency in terms of *the price of anarchy*. Via extensive simulations, we validate the stability and efficiency of our incentive scheme.

Keywords: P2P-VoD, incentive, mean field, Stackelberg game.

1 Introduction

In recent years, we have witnessed the rapid growth of Peer-to-Peer (P2P) systems, many of which have large population bases, e.g., file sharing systems like BitTorrent [4], and Video-on-Demand (VoD) systems like PPLive [1] and PP-Stream [2]. The key advantage of the P2P architecture is that by utilizing the distributed resources at the peers, the system can be more scalable and fault-tolerant than traditional client-server architectures. Nevertheless, due to the selfish nature of the peers, free-riding [5] often happens where peers do not have incentives to contribute resource for other peers. Thus, designing an effective and practical incentive scheme becomes critical in encouraging the peers to contribute to the system, and thereby improving the system performance. Plenty of work has been done for systems of traditional P2P applications, for example, the tit-for-tat [4] protocol has been well-adopted for file sharing applications.

However, very limited work has been focusing on the incentive mechanisms for P2P-VoD applications.

What makes it challenging to design incentive schemes for P2P-VoD applications? Compared to file sharing, VoD applications need to satisfy more stringent temporal and spacious constraints for data delivery. To share files, peers exchange segments of files that have not been received. Segments might be received in different orders, therefore, there is hardly a temporal constraint under which a particular segment has to be received. On the contrary, when a user watches a particular video segment, this segment has to be received by the user within a short period of time, while nearby segments would not satisfy the user's instantaneous demand. Even worse, predicting the data demand is difficult because users might fast-forward and/or rewind among the video segments. The tit-for-tat scheme does not work for VoD applications, because the data demand and supply among the peers are highly volatile.

Instead of using a punishment-based scheme like tit-for-tat, we propose and analyze a *reward-based scheme* that incentivizes peers to contribute upload capacity for VoD systems. Our contributions are as follows.

- We model the stochastic operations of the peers and derive the system state by using the *mean field approximation*.
- We propose a practical reward-based incentive scheme based on the dedicated upload capacity of the peers.
- We model the interaction between the content provider and the peers by using a *Stackelberg game*. We derive the unique Stackelberg equilibrium and analyze the efficiency of the equilibrium in terms of *the price of anarchy*.
- We validate the effectiveness of our scheme and the theoretic results via extensive simulations.

Our paper is organized as follows. In Section 2, we present the system model and the reward-based incentive scheme for P2P-VoD systems. In Section 3, we study the system dynamics and characterize the content provider's cost as a function of the peers' dedicated upload bandwidth. In Section 4, we model the strategic behavior between the content provider and the peers, derive a unique system equilibrium, and analyze the efficiency of that. We evaluate the performance of our incentive scheme in Section 5. Section 6 states the related work and Section 7 concludes.

2 System Model and Reward-Based Scheme

In a P2P-VoD system, each peer can support other peers by (1) caching data that would be needed for other peers and (2) uploading data to other peers. Both aspects are equally important because a peer cannot contribute if it either does not have the needed content, or does not have upload capacity. Thus, the design space for an incentive mechanism includes both incentivizing peers to cache the right content as well as to devote upload capacity. Due to the complicated viewing operations supported by VoD systems, data demand and supply

can be volatile. Therefore, the optimal data caching policy for a peer might be difficult to predict given its local knowledge of the system. Although the content provider's global knowledge might help guide the data caching policy, collecting this global knowledge and coordinating with peers bring overhead to the system, even if the peers are willing to comply. Our study focuses on incentivizing the peers to dedicate upload capacity, while making a minimum assumption on the data caching policy. Naturally, each peer caches the video segment it recently watched/requested and can contribute to other peers that need the video segment. Under this default policy, a peer cannot contribute much either due to its little upload capacity or the low demand for its recently watched segment.

Since peers only cache the segments of a video they are watching, without loss of generality, we focus on a particular video content provided in the system. We denote N as the long-term average or steady-state number of peers that watch this video. We denote r as the required playback rate, i.e., bits per second, for serving the video. We consider homogeneous peers and explore the symmetric strategies of them. We denote u as the upload capacity of each peer, which is the maximum bandwidth a peer is willing to contribute to the system. In Section 4, the peers will choose u as their strategy to maximize utility.

2.1 Peers' Viewing Behavior

We assume that the system organizes the video as K consecutive data segments. We denote S_i as the i th data segment. We model the user behavior of the system by specifying a set of rate transition probabilities $\{p_{ij} : i, j = 0, 1, \dots, K\}$. Each p_{ij} denotes the transition probability of a typical user watching S_j after finishing segment S_i . In particular, p_{i0} denotes the probability that a user quits watching after finishing S_i . We denote λ as the aggregate external arrival rate of users that start watching the video. However, users might not start from the very first segment S_1 , because they might have watched part of the video before. Thus, p_{0i} denotes the probability that an external arrival will start with S_i . To keep consistency, we define $p_{00} = 0$ and require $\sum_{j=0}^K p_{ij} = 1$ for all $i = 0, \dots, K$. In practice, the probability p_{ij} represents a state transition where a peer performs a play, fast-forward or rewind operation when $j = i + 1$, $j > i + 1$ or $0 < j < i + 1$, respectively.

Under our default caching policy, we assume that after transitioning from segment i to segment j , a peer only provides S_i to other peers if needed when it is viewing S_j . In practice, a peer can cache and upload multiple data segments at any time. However, we will show that even by using such a restricted caching policy, one can still design a simple yet effective scheme to incentivize peers to dedicate upload capacity for other peers, which consequently reduces the workload of the content provider as well as improves the system performance. When a peer downloads data segment i , we assume that the system will direct the demand to the peers that can contribute S_i first, which upload S_i at an aggregate rate that equals the playback rate r . If the supply capacity of S_i is lower than aggregate demand rate, the content server will support the remaining data rate by using content servers' capacity.

2.2 Reward-Based Incentive Scheme

We design an incentive mechanism under which the content provider rewards the peers based on the amount of upload capacity they contribute. The reward can be in various forms, e.g., real money rebate for the service fee and virtual credits or reputation record for advanced services. Notice that any reward scheme can be represented by the currency flow from the content provider to the peers. Even for rewards in virtual currency or reputation, they imply that the P2P-VoD operator needs to invest money for developing advanced/prioritized services for users. We do not restrict the form of implementing the rewards in our paper; however, we use an abstract model to describe the reward in terms of monetary value.

We define the reward W to a peer to be a function of its dedicated maximum upload bandwidth capacity u as

$$W(u) = \int_0^u w(x)dx, \quad (1)$$

where $w(x)$ denotes the marginal reward at the contribution capacity level x . Notice that our incentive scheme is based on the maximal upload bandwidth that a peer is willing to dedicate to the system; in practice, whether a peer will upload data at the maximum capacity u depends on the demand for the data segment the peer has cached and the peer might not upload at the rate of u . We assume that the system will maximize the utilization of upload capacity of the peers whenever their data segments are requested by other peers.

A simple reward scheme is to use a linear reward function

$$W(u) = wu, \quad (2)$$

where the marginal reward $w(x) = w$ is the same for all levels of contribution x . w can also be interpreted as the reward per unit capacity dedicated by a peer. The linear reward scheme can be easily understood by the peers and implemented by the content provider in practice. In what follows, we will start with the linear reward model and extend our results for general reward functions W later.

In summary, under our reward based incentive scheme, the content provider decides the reward function W , and then each peer decides its upload capacity u dedicated to the P2P-VoD system. In Section 4, we present a game theoretic framework to analyze the interaction between the content provider and peers under this reward-based incentive scheme. Before we present this analysis, let us first investigate the distribution of peers in different video segments so as to understand the impact of peers' upload contribution on the content provider's upload cost.

3 Peers' Contribution and Content Provider's Cost

In this section, we derive the distribution of number of peers watching different segments based on the user behavior described in Section 2.1. We further characterize the content provider's upload cost as a function of both the distribution of peers and their upload contribution.

3.1 Distribution of Peers in Different Video Segments

Based on a typical user’s viewing behavior described in Section 2.1, we say that a peer is in state i when it is watching segment S_i . Thus, each peer’s viewing behavior becomes a random process. For a system with N peers, the state space of the P2P-VoD system is $[1, \dots, K]^N$. To overcome the large dimensionality of the state space, we tackle the problem from a macro perspective, i.e., instead of observing each peer’s individual state, we are only interested in the *fraction* of peers in each of the states or the distribution of the peers in the states. In particular, we use the *mean field interaction model* [3] to calculate the steady state distribution of these peers in the P2P-VoD system. We denote $q_i \in [0, 1]$ as the fraction of peers in state i . The system state can now be specified by the vector (q_1, q_2, \dots, q_K) , where $\sum_{i=1}^K q_i = 1$.

Theorem 1. *If the P2P-VoD system does not support the rewind operation, i.e., $p_{ij} = 0$ for all $1 \leq j \leq i$, then*

$$q_i = \frac{P_{0i}}{\sum_{j=1}^K P_{0j}} \quad \forall i = 1, \dots, K, \tag{3}$$

where P_{ij} denotes the aggregated probability of transitioning from state i to j , which can be defined recursively as follows.

$$P_{ij} = \begin{cases} p_{ij} & \text{if } j = i + 1, \\ \sum_{k=i+1}^{j-1} p_{ik}P_{kj} + p_{ij} & \text{otherwise.} \end{cases} \tag{4}$$

Proof. Recall λ is the external arrival rate of new peers that start watching the video. In the steady-state, we can express the external arrival rate to state i as λp_{0i} . We denote μ as the rate at which users change states, which includes the events of (1) moving to the next segment, (2) fast-forwarding or rewinding and (3) quitting the system. Because the departure rate of a particular state is proportional to the number of peers in the state, the departure rate from state i is μq_i . After finishing watching S_i , a peer has different probabilities to start watching another segment; and therefore, the transition rate from state i to state j is $\mu q_i p_{ij}$. Each system state i satisfies the following ordinary differential equation (ODE):

$$N \cdot \frac{dq_i}{dt} = \lambda p_{0i} + \sum_{j=1, j \neq i}^K \mu q_j p_{ji} - \sum_{k=1, k \neq i}^K \mu q_i p_{ik} - \mu q_i p_{i0}, \quad \forall i. \tag{5}$$

The left hand side represents the change in the number of peers in state i . The right hand side counts all the possible cases where changes can incur. The increase of number of peers in state i can occur when external arrivals start watching the video at segment S_i (with a rate of λp_{0i}), or any peer which just finished watching S_j continues (or performs fast-forward or rewind operation) to watch S_i (with a rate of $\mu q_j p_{ji}$). The decrease of number of peers in state i can

occur when a peer finishes watching S_i and starts watching another segment S_k (with rate $\mu q_i p_{ik}$), or this peer quits viewing the video (with rate $\mu q_i p_{i0}$).

In steady-state, the external arrival rate to the P2P-VoD system should be equal to the departure rate, i.e.,

$$\lambda = \sum_{i=0}^K \mu q_i p_{i0}, \tag{6}$$

and the rate of change in each state should be zero, i.e.,

$$N \cdot \frac{dq_i}{dt} = 0. \tag{7}$$

Given $p_{ij} = 0$ for all $1 \leq j \leq i$, we have $q_1 = \frac{\lambda}{\mu} p_{01}$ and $q_i = \frac{\lambda}{\mu} p_{0i} + \sum_{j=1}^{i-1} p_{ji} q_j$. By recursively solving q_i and requiring $\sum_{i=1}^K q_i = 1$, we can derive the above formula. \square

3.2 Content Provider’s Upload Cost

Since content providers are often charged by their transit providers (ISPs) based on the traffic volume going through them, we assume that the content provider’s cost is proportional to the upload capacity needed to support all peers. In steady-state, Nq_i peers watch segment S_i . Given a required playback rate of r , the aggregate required upload capacity for S_i should be $Nq_i r$. In the proof of Theorem 1, we have derived $q_i = \frac{\lambda}{\mu} p_{0i} + \sum_{j=1}^{i-1} p_{ji} q_j$, or equivalently, $Nq_i = N \frac{\lambda}{\mu} p_{0i} + \sum_{j=1}^{i-1} N p_{ji} q_j$. Among the peers watching S_i , $N \frac{\lambda}{\mu} p_{0i}$ peers are external arrivals that start watching at S_i and $N p_{ji} q_j$ peers have viewed S_j before transitioning to watch S_i . Note that we can only assure that peers performing continuous play, i.e., moving from S_i to S_{i+1} , have watched and therefore cached the whole segment of S_i . Under our simplistic caching policy, only the peers that have transitioned from S_i to S_{i+1} and currently watching S_{i+1} can upload S_i to other peers. Suppose each peer contributes u amount of capacity for uploading available video segments, the total available upload capacity for S_i would be $Nq_i p_{i,i+1} u$. If the dedicated peer contribution $Nq_i p_{i,i+1} u$ is less than the required download capacity $Nq_i r$, the content provider needs to upload segment S_i to support the difference in capacity. In particular, the playback requirement of the last segment, $Nq_K r$, must be supported by the content provider. Assume the content provider incurs a cost c_s ($c_s > 0$) per unit bandwidth capacity. Then, the content provider’s total upload cost is:

$$C_s(u) = c_s N \times \left[\sum_{i=1}^{K-1} q_i (r - p_{i,i+1} u)^+ + q_K r \right]. \tag{8}$$

Proposition 1. $C_s(u)$ is a convex and non-increasing function in u .

Proof. Because c_s , N , q_i and $p_{i,i+1}$ are all positive, the cost $C_s(u)$ is non-increasing in u . Since $(r - p_{i,i+1}u)$ is linear (and therefore, convex) in u , $(r - p_{i,i+1}u)^+ = \max(r - p_{i,i+1}u, 0)$ is convex in u . Given that the convex property keeps under summation operation, we reach the above conclusion. \square

Notice that the maximum cost is $C_s(0) = c_sNr$ when the peers do not contribute any capacity and the minimum cost is $C_s(u) = c_sNqKr$ when u is large enough. In particular, when $u \geq r$, the system might not be able to utilize all peers' upload resource, i.e., fewer peers will participate in data uploading when u increases. This implies the sub-linearity of cost saving of the content provider with respect to the increase of peers' capacity contribution u . We will show an example that validates the convexity feature in Section 5.

4 Game Theoretic Analysis on Incentive Scheme

In this section, we present a game-theoretic model to study the strategies of the content provider and the peers in a P2P-VoD system under the reward-based incentive scheme and analyze the stability and efficiency of the incentive scheme. We define w , the per capacity reward to the peers, as the strategy of the content provider and u , the amount of dedicated capacity, as the strategy of the peers. We assume that peers are homogenous and use the same u strategy in the game. We denote $[0, \bar{w}]$ and $[0, \bar{u}]$ as the strategy space of the content provider and the peers, where \bar{w} and \bar{u} are the upper-bound of the content provider's and peers' strategy respectively.

4.1 Stackelberg Game Model

From the content provider's perspective, it aims at minimizing its total cost, i.e., the cost of uploading and the cost of rewarding the peers. We define the utility of the content provider as the following:

$$\pi_s(w, u) = -C_s(u) - wuN. \quad (9)$$

Similarly, we define the utility of a peer as the reward it receives, minus its cost of upload contribution as the following:

$$\pi_p(u, w) = wu - C_p(u), \quad (10)$$

where $C_p(u)$ denotes the cost of dedicating u amount of capacity. To maximize their utilities, the content provider solves the optimization problem $\max_w \pi_s(w, u)$, and the peers solves $\max_u \pi_p(u, w)$. Here, we do not specify the form of the peer's upload cost function $C_p(u)$. Rather, we assume the cost function satisfies the following property:

- (1) $C_p(u)$ is continuous and twice differentiable in u .
- (2) $C_p(0) = 0$, $C_p'(u) > 0$, $C_p''(u) > 0$.

$C'_p(u) > 0$ means that a peer's cost increases with its dedicated capacity. $C''_p(u) > 0$ means the marginal cost also increases with the dedicated capacity. The above assumption reflects the fact that a peer's viewing performance would not be affected too much if it contributes a small amount of upload capacity; however, when a peer dedicates much upload capacity, its download rate as well as the performance of video might be substantially reduced.

Proposition 2. $\pi_p(u, w)$ is a strictly concave function in u .

Proof. Noting that $C''_p(u) > 0$ implies $-C_p(u)$ is strictly concave in u , and that wu is linear and hence concave in u , we have $\pi_p(u, w)$ strictly concave in u . \square

We consider a Stackelberg game [15] where the content provider decides w first, and after that, the peers decide u . It is natural to assume the content provider as the first-mover whereas the peers response to the reward w accordingly, because once u is determined, the content provider would have no incentives to provide any reward for the peers. To obtain the Stackelberg equilibrium of the game, we can use the backward induction [15]. In particular, the peers solve the problem $u^*(w) = \operatorname{argmax}_u \pi_p(u, w)$ given any w . By knowing the peers' best responses, the content provider solves the problem $w^* = \operatorname{argmax}_w \pi_s(w, u^*(w))$. In what follows, we analyze the existence, uniqueness and efficiency of the Stackelberg equilibrium.

4.2 Existence and Uniqueness of Stackelberg Equilibrium

We start with the following lemma, which establishes the connection between the Stackelberg equilibrium and an optimization problem:

Lemma 1. *If u^* is a solution to the following problem:*

$$\min_u C_s(u) + NuC'_p(u), \tag{11}$$

then there exists a Stackelberg equilibrium $(u^, u^*C'_p(u^*))$; further, if (u^*, w^*) is a Stackelberg equilibrium, then u^* is the solution to problem (11).*

Proof. We start by showing the first half of the statement. Denote $u^* = \operatorname{argmin}_u [C_s(u) + NuC'_p(u)]$ and $w^* = u^*C'_p(u^*)$. We show that (u^*, w^*) is a Stackelberg equilibrium. Since $\pi_p(u, w)$ is strictly concave in u , so for any given w^* , if u^* satisfies $u^*C'_p(u^*) = w^*$, then u^* maximizes the peers' utility $\pi_p(u, w^*)$. Hence, the peers do not have incentives to deviate from u^* . Suppose the content provider has an incentive to deviate from w^* and can obtain higher utility by setting $w = w_0$, where the peers' response is to set $u = u_0$ so that u_0 maximizes $\pi_p(u, w_0)$. Because of the strict concavity of $\pi_p(u, w)$, there are only three possible cases:

- (1) $C'_p(u_0) = w_0$ if $C'_p(0) \leq w_0 \leq C'_p(\bar{u})$; or
- (2) $u_0 = 0$ if $C'_p(0) > w_0$; or
- (3) $u_0 = \bar{u}$ if $C'_p(\bar{u}) < w_0$.

For any of the above cases, we have

$$\begin{aligned}
 C_s(u_0) + Nu_0C'_p(u_0) &\leq C_s(u_0) + Nu_0w_0 \\
 &< C_s(u^*) + Nu^*w^* = C_s(u^*) + Nu^*C'_p(u^*).
 \end{aligned}
 \tag{12}$$

The first inequality holds for the above three cases. The second inequality holds because we assume the content provider can have higher utility by setting $u = u_0$ instead of $u = u^*$. However, $C_s(u_0) + Nu_0C'_p(u_0) < C_s(u^*) + Nu^*C'_p(u^*)$ contradicts the fact that u^* is a solution of (11). This implies that the content provider has no incentive to deviate from w^* . Given that we have shown the peers do not have any incentive to deviate from u^* given any w^* , we conclude (u^*, w^*) is a Stackelberg equilibrium¹.

To show the second half of the statement, suppose there exists a Stackelberg equilibrium (u^*, w^*) , but u^* is not a solution to (11), i.e., there exists $u_0 \neq u^*$ such that $C_s(u_0) + Nu_0C'_p(u_0) < C_s(u^*) + NuC'_p(u^*)$. Assume the content provider sets $w_0 = u_0C'_p(u_0)$. Taking the derivative in (10) and noting the strict concavity of $\pi_p(u, w)$, we have the peers' unique best response is $u = u_0$ for given $w_0 = u_0C'_p(u_0)$. Therefore, $\pi_s(w_0, u_0) = -C_s(u_0) - Nu_0C'_p(u_0) > -C_s(u^*) - Nu^*C'_p(u^*) = \pi_s(w^*, u^*)$, which contradicts to the fact that (u^*, w^*) is a Stackelberg equilibrium. This implies u^* must be a solution to (11). \square

Theorem 2. *The Stackelberg equilibrium always exists. If $uC'_p(u)$ is strictly convex in u , then the peers' solution u^* at the Stackelberg equilibrium is unique.*

Proof. We first show the existence. The peers solve $\max_u \pi_p(u, w) = wu - C_p(u)$. For any given w , π_p is continuous and strictly concave in u over the compact set $[0, \bar{u}]$. Hence, the optimal solution $u^*(w) = \operatorname{argmax}_u \pi_p(u, w)$ exists and is unique. Substituting u by $u^*(w)$ in $\pi_s(w, u)$, the provider's utility $\pi_s(w, u^*(w))$ is continuous in w over the compact set $[0, \bar{w}]$, so $w^* = \operatorname{argmax}_w \pi_s(w, u^*(w))$ exists.

Next we show the uniqueness of u^* when $uC'_p(u)$ is strictly convex in u . Since $C_s(u)$ is convex in u (Proposition 1), and $uC'_p(u)$ is strictly convex in u , we can observe that the problem (11) is a strictly convex minimization over a compact set, which has a unique solution. According to Lemma 1, any Stackelberg equilibrium (u^*, w^*) satisfies that u^* is a solution to (11). Therefore, we conclude that the peers' solution in the Stackelberg equilibrium is unique². \square

In the proof, we assume $uC'_p(u)$ to be strictly convex. In fact, if the marginal cost $C'_p(u)$ is super-linearly increasing in u , then by multiplying a linear function $f(u) = u$, the term $uC'_p(u)$ can be guaranteed to be strictly convex.

¹ Noting the above three cases and that the content provider aims at maximizing its utility, if $u^* > 0$, then the corresponding Stackelberg equilibrium is unique where $w^* = u^*C'_p(u^*)$. If $u^* = 0$, then any (u^*, w^*) where $0 \leq w^* \leq C'_p(0)$ is a Stackelberg equilibrium.

² We do not claim the Stackelberg equilibrium is unique. The only chance of having multiple Stackelberg equilibria is $u^* = 0$, where any (u^*, w^*) with $0 \leq w^* \leq C'_p(0)$ is a Stackelberg equilibrium. When $u^* > 0$, the Stackelberg equilibrium is unique, where the content provider sets $w^* = u^*C'_p(u^*)$.

4.3 Efficiency of Stackelberg Equilibrium

Now we discuss the efficiency of the Stackelberg equilibrium. For mathematical simplicity, in this subsection, we assume $C_s(u)$ is differentiable in u .

We define the social welfare, π_w , as the sum of the content provider's and all peers' utilities:

$$\pi_w(u) = \pi_s + N\pi_p = -C_s(u) - NC_p(u). \tag{13}$$

Because of the convexity of $C_s(u)$ and the strict convexity of $C_p(u)$, we immediately have

Proposition 3. $\pi_w(u)$ is strictly concave in u .

We define the *price of anarchy* (PoA) [10] to be the ratio of the social welfare at the worst Stackelberg equilibrium to the maximal social welfare one can achieve when varying $u \in [0, \bar{u}]$. In particular, when the Stackelberg equilibrium (u^*, w^*) is unique, we have

$$\text{PoA} = \frac{\pi_w(u^*)}{\pi_w(u_w)}, \tag{14}$$

where $u_w = \operatorname{argmax}_u \pi_w(u)$ and u^* is the peers' solution at the Stackelberg equilibrium. In our model, the social welfare is non-positive, so PoA is in general no less than 1. When PoA is close to 1, it implies the system is in an efficient state. We first state the following lemma:

Lemma 2. *The peers' upload contribution at the Stackelberg equilibrium is no larger than the upload capacity that maximizes the social welfare, i.e., $u^* \leq u_w$.*

Proof. Denote $C_{SW}(u) = -\pi_w(u) = C_s(u) + NC_p(u)$, and $C_{SE}(u) = C_s(u) + uC'_p(u)$. Maximizing the social welfare is equivalent to solving $\min_u C_{SW}(u)$, $u \in [0, \bar{u}]$. According to Lemma 1, u^* can be obtained by solving $\min_u C_{SE}(u)$, $u \in [0, \bar{u}]$. Therefore, u_w and u^* are the minimizers to $C_{SW}(u)$ and $C_{SE}(u)$, respectively. By taking the first order derivative, we have

$$C'_{SW}(u) = C'_s(u) + NC'_p(u), \tag{15}$$

$$C'_{SE}(u) = C'_s(u) + NC'_p(u) + NuC''_p(u). \tag{16}$$

There are only two possible cases regarding $C'_{SW}(u)$:

(1) If $C'_{SW}(u) > 0, \forall u \in [0, \infty)$, then $u_w = 0$. Since $NuC''_p(u) \geq 0$, we have $C'_{SE}(u) = C'_{SW}(u) + NuC''_p(u) > 0, \forall u \in [0, \infty)$, so $u^* = 0 = u_w$.

(2) If there exists a $u_{SW} \in [0, \infty)$ such that $C'_{SW}(u_{SW}) = 0$, then u_{SW} must be unique due to the strict convexity of $C_{SW}(u)$. We have $u_w = \max(u_{SW}, \bar{u})$. By the concavity assumption on $C_p(u)$ and Proposition 1, $C'_s(u)$ and $NC'_p(u)$ are both non-decreasing in u and $NuC''_p(u) > 0$. Hence, for any $u > u_{SW}$, we have $C'_{SE}(u) > C'_{SW}(u) > C'_{SW}(u_{SW}) = 0$. This implies any $u > u_w = \max(u_{SW}, \bar{u})$ cannot be the minimizer of $C_{SE}(u)$, $u \in [0, \bar{u}]$. Therefore, $u^* \leq u_w$.

Combining the results in the above two cases, we have $u^* \leq u_w$. □

Theorem 3. *Denote u^* as the peers' solution at any Stackelberg equilibrium. If $u^* = 0$ or $u^* = \bar{u}$, the system obtains the maximal social welfare, i.e., $\text{PoA} = 1$.*

Proof. If $u^* = 0$, then $C'_{SE}(0) \geq 0$. Suppose $u_w \neq 0$, then by the strict convexity of $C_{SW}(u)$, we have $C'_{SW}(u_w) = 0$ and $C'_{SW}(0) < 0$. From Eq. 15 and Eq. 16, we have $C'_{SE}(0) = C'_{SW}(0) < 0$, which contradicts to $C'_{SE}(0) \geq 0$. Hence, $u_w = 0 = u^*$, so PoA=1.

If $u^* = \bar{u}$, then by Lemma 2, we have $u_w \geq u^* = \bar{u}$. In the meanwhile, $u_w \leq \bar{u}$, so $u_w = \bar{u} = u^*$, and hence PoA=1. \square

In general, $\pi_w(u_w)$ and $\pi_w(u^*)$ may not be equal. Given the concavity property in Proposition 3, $\pi_w(u)$ is strictly increasing in $[0, u_w]$. Therefore, the gap between $\pi_w(u_w)$ and $\pi_w(u^*)$ is impacted by the gap between u_w and u^* . In particular, we have $\pi_w(u_w) - \pi_w(u^*) \leq \pi'_w(u^*)(u_w - u^*)$, and $\text{PoA} \leq 1 - \frac{\pi'_w(u^*)}{\pi_w(u_w)}(u_w - u^*)$. Given a general form of $C_p(u)$, it is mathematically hard to further quantify PoA for $0 < u^* < \bar{u}$. In the next section, we will use simulations to show the efficiency of the Stackelberg equilibrium in general cases.

4.4 General Reward Scheme

The linear reward model is a simplification of the general reward model in Section 2. If we use the general model where $w(x)$ denotes the marginal reward per upload capacity x , the content provider's problem is

$$\max_{w(x)} \pi_s(w(x), u) = -C_s(u) - Nu \int_0^u w(x) dx, \quad (17)$$

and the peers' problem is

$$\max_u \pi_p(u, w(x)) = u \int_0^u w(x) dx - C_p(u). \quad (18)$$

We discuss the following question: among all possible reward models, which reward model can make the content provider obtain the maximal utility?

Theorem 4. *If the content provider can find $u^* = \operatorname{argmax}_u -C_s(u) - NC_p(u)$, then any reward scheme satisfying the following property can make the content provider's utility arbitrarily close to the maximal:*

$$W(u) = \int_0^u w(x) dx \begin{cases} \leq C_p(u) & \text{if } u \neq u^*, \\ = C_p(u) + \epsilon & \text{if } u = u^*, \end{cases} \quad (19)$$

where ϵ is a positive real number and is arbitrarily small.

Proof. For a given reward scheme $w(x)$, assume the peers set $u = u^*(w(x))$ to maximize their utility. We have $\pi_p(u^*, w(x)) = u^* \int_0^{u^*} w(x) dx - C_p(u^*) \geq 0$; otherwise, the peers can obtain $\pi_p(u, w) = 0 > \pi_p(u^*, w(x))$ by setting $u = 0$. Therefore, the content provider's utility $\pi_s(w(x), u)$ in (17) is upper-bounded by $\operatorname{argmax}_u -C_s(u) - NC_p(u)$. The content provider's utility achieved in Eq. 19 is $\operatorname{argmax}_u -C_s(u) - NC_p(u) - \epsilon$, so it can be arbitrarily close to the upper-bound when ϵ is arbitrarily small. \square

In this subsection, we relax the requirement of continuity on the reward function $W(u)$. An interesting implication is that the theorem provides us the insight in designing such reward schemes that maximize the content provider's utility. In fact, $\max_u -C_s(u) - NC_p(u)$, or $\min_u C_s(u) + NC_p(u)$ is a standard convex optimization and can be easily solved. After obtaining u^* , we can easily design a reward scheme satisfying (19). For example, we can design

$$w(x) = C'_p(x) + \epsilon\delta(u^*), \quad (20)$$

where $\delta(x)$ is the unit impulse function.

It is also interesting to note, using the general reward scheme in Theorem 4, the procedure of determining u^* is exactly maximizing the social welfare. Therefore, we have PoA=1, i.e., the social welfare is maximized when the content provider maximizes its own utility.

We briefly conclude the result in our game theoretic analysis. The interaction of the content provider and the peers can be viewed as a Stackelberg game where the content provider takes the first action and the peers follow. The existence and uniqueness of Stackelberg equilibrium shows the stability of the reward scheme, while efficiency is quantified by *price of anarchy*. We also point out the content provider's best strategy in designing reward in the general form.

5 Performance Evaluation

In this section, we validate the stability and efficiency of our incentive scheme via extensive simulations. We have the following settings:

- The system consists of $N = 10,000$ peers, one server and one video with $K = 50$ segments and playback rate $r = 500$ Kbps.
- The peers' external arrival probability to segment i is $p_{0i} = \frac{2^{i-1}}{(i-1)!}e^{-2}$ ($1 \leq i \leq K$); the probability of doing play operation is $p_{i,i+1} = 1 - 0.6e^{-0.25 \times i}$ ($1 \leq i \leq K-1$); and the probability of quit operation is $p_{i0} = 1 - p_{i,i+1}$ ($1 \leq i \leq K-1$), $p_{K0} = 1$. Assume the fast-forward and rewind operations are rare and can be omitted.
- The content provider's cost per unit capacity $c_s = 1$. The peers' upload cost $C_p(u) = c_p u^\beta$, where c_p and β are parameters we will vary in simulation.
- The content provider chooses $w \in [0, 1]$, and the peers choose $u \in [0, 1000]$ Kbps.

In Fig. 1, we plot the arrival probability p_{0i} and departure probability p_{i0} of each segment i . Observe that when a peer starts watching a video, it has probability $p_{01} = 0.135$ to start from the first segment, and $p_{02} = p_{03} = 0.271$ to start from the second or third segment, but the probability of watching from S_4 and onwards decreases rapidly. This corresponds to the reality where some people start from the beginning, but more people would like to skip the first few segments like advertisement. We can also observe that peers watching S_1 will quit the viewing course with probability 0.6, but the probability of quit operation decreases for




Fig. 1. Arrival and departure probability for each video segment




Fig. 2. Fraction of peers in each video segments




Fig. 3. Content provider's upload cost




Fig. 4. Content provider's utility

peers watching later segments of the video. In Fig. 2, we plot the fraction of peers q_i for each video segment i . We can see from the figure that there is an increasing trend of popularity from S_1 to S_4 , and a decreasing trend thereafter. This trend is due to the peers' viewing behaviors described above.

In Fig. 3, we plot the content provider's upload cost $C_s(u)$ (refer to Eq. 8) when we vary $u \in [0, 1000]$ Kbps. From this figure, we can observe that $C_s(u)$ is convex and non-increasing in u , which validates Proposition 1.

In Fig. 4, we investigate the content provider's utility $\pi_s(w, u)$ (refer to Eq. 9) when varying the unit reward w to peers. Assume the peers decide their upload capacity u to maximize their utility for given w . In this simulation, we fix $c_p = 0.15$ and $\beta = 1.2$ in the peers' upload cost function $C_p = c_p u^\beta$. We plot the content provider's utility π_s when varying $w \in [0, 1]$. When $w = 0$, peers do not contribute any upload bandwidth and thus the content provider's utility equals the negative value of the cost for supporting all peers' viewing requirement. When w increases from 0, the content provider can utilize part of peers' upload capacity so as to increase its utility. However, when w is very large, peers' decision u is also large. The content provider's utility decreases due to the huge amount of reward it has to pay to the peers. There exists an optimal unit reward $w = 0.62$ where the content provider's utility is maximized with value -3.63×10^6 .




Fig. 5. Stackelberg equilibrium

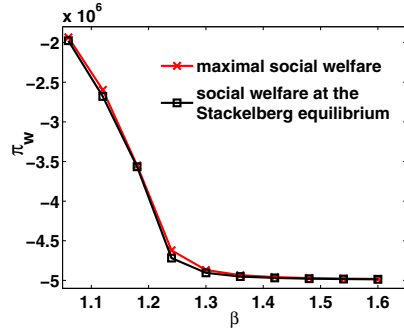


Fig. 6. Social welfare

In Fig. 5, we investigate the Stackelberg equilibrium under different parameters. In particular, we plot the value of w^* and u^* at the Stackelberg equilibrium when fixing $\beta = 1.2$ and varying $c_p \in [0.02, 0.20]$. The figure shows that when c_p increases, i.e., the upload cost of peers increases, the content provider rewards more to peers, and the peers tend to contribute less. However, we also note u^* is not monotonous in c_p , and there is a local minimum of u^* when $c_p = 0.06$.

We compare the maximal social welfare and the social welfare at Stackelberg equilibrium in Fig. 6. We fix $c_p = 0.2$ and vary $\beta \in [1.06, 1.60]$. We can observe from the figure that the social welfare at the Stackelberg equilibrium is always very near to the maximal value, i.e., $\text{PoA} \approx 1$ for any $\beta \in [1.06, 1.60]$. This simulation result illustrates the efficiency of our incentive scheme.

To summarize, all these simulation results validate our theoretic analysis and show the stability and efficiency of our incentive scheme.

6 Related Work

Incentive issue has received plenty of attentions in P2P applications. Zhao *et al.* [19] proposed a general framework to evaluate the expected performance gain and system robustness for a class of incentive protocols wherein peers can distributively learn to adapt their actions. In [16], the authors used game model to analyze the content production and sharing in P2P networks and compare the performance of different existing incentive schemes. There are also some existing works on designing particular incentive schemes. The first incentive scheme proposed for P2P system is the micropayment in [6]. Misra *et al.* [13] proposed a Shapley value approach in incentive design using a cooperative game model. Reputation [7,9] is another well-known approach where a peer's reputation represents its history of contribution in the system. Ma *et al.* [12] proposed a service differentiation approach in P2P network based on the amount of contribution each node has provided to the network community.

All these existing works are based on general P2P settings or are specifically designed for P2P file sharing systems. However, P2P-VoD systems have special features. Wu *et al.* [18] analyzed how to efficiently utilize the peers' resources,

but did not address how to incentivize peers to contribute their resources. Habib *et al.* [8] proposed a service differentiation approach for incentive scheme in P2P multimedia systems, where peers with high contribution have flexibility in peer selection so that they receive better quality of service. Mol *et al.* [14] designed a free-riding-resilient P2P-VoD system where peers favor uploading to other peers who have proven to be good uploaders. These two works are similar because these incentive schemes they proposed are both variants of the tit-for-tat mechanism in file sharing applications. Similar approaches were also proposed for live streaming systems, e.g., in [17], the authors presented a modified tit-for-tat mechanism; in [11], a multi-layered live streaming system punishes the peers with low contribution by providing them with low quality of service. Instead of using the punishment-based approach, we propose a general reward-based incentive scheme where we incentivize the peers to contribute their upload capacity. In practice, peers' upload capacity is constrained by their Internet access types and hence varies a lot. Using our approach, peers with low upload capability can still receive good quality of service provided that they accept a low level of reward.

7 Conclusion

Incentive scheme is a key design issue in P2P applications in order to encourage peers' resource contribution. However, due to the complex and stochastic nature of peers' behavior, it is challenging to design an effective incentive scheme in P2P-VoD systems. In this paper, we propose a simple yet effective reward-based incentive scheme. We model the peers' distribution in movie segments using *mean field approximation*, characterize the content provider's cost in terms of peers' contribution and develop a game-theoretic framework to analyze the interactions of the content provider and peers under our reward-based incentive mechanism. We also show the stability and efficiency of our incentive scheme via extensive simulations.

References

1. pplive, www.PPLive.com
2. ppstream, www.PPStream.com
3. Benaïm, M., Le Boudec, J.-Y.: A class of mean field interaction models for computer and communication systems. *Performance Evaluation* 65, 823–838 (2008)
4. Cohen, B.: Incentives build robustness in bittorrent. In: *Proceedings of the 1st Workshop on Economics of Peer-to-Peer Systems* (2003)
5. Feldman, M., Papadimitriou, C., Chuang, J., Stoica, I.: Free-riding and white-washing in peer-to-peer systems. *IEEE Journal on Selected Areas in Communications* 24(5), 1010–1019 (2006)
6. Golle, P., Leyton-Brown, K., Mironov, I.: Incentives for sharing in peer-to-peer networks. In: *Proceedings of the 3rd ACM conference on Electronic Commerce, EC* (2001)
7. Gupta, M., Judge, P., Ammar, M.: A reputation system for peer-to-peer networks. In: *Proceedings of the 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, pp. 144–152 (2003)

8. Habib, A., Chuang, J.: Service differentiated peer selection: an incentive mechanism for peer-to-peer media streaming. *IEEE Transactions on Multimedia* 8(3), 610–621 (2006)
9. Kamvar, S.D., Schlosser, M.T., Garcia-Molina, H.: The eigentrust algorithm for reputation management in p2p networks. In: *Proceedings of the 12th International Conference on World Wide Web (WWW)*, pp. 640–651 (2003)
10. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. *Computer Science Review* 3(2), 65–69 (2009)
11. Liu, Z., Shen, Y., Panwar, S.S., Ross, K.W., Wang, Y.: Using layered video to provide incentives in p2p live streaming. In: *Proceedings of the 2007 Workshop on Peer-to-peer Streaming and IP-TV (P2P-TV)*, pp. 311–316 (2007)
12. Ma, R.T.B., Lee, S.C.M., Lui, J.C.S., Yau, D.K.Y.: Incentive and service differentiation in p2p networks: a game theoretic approach. *IEEE/ACM Transactions on Networking* 14, 978–991 (2006)
13. Misra, V., Ioannidis, S., Chaintreau, A., Massoulié, L.: Incentivizing peer-assisted services: a fluid shapley value approach. In: *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pp. 215–226 (2010)
14. Mol, J., Pouwelse, J., Meulpolder, M., Epema, D., Sips, H.: Give-to-get: Free-riding-resilient video-on-demand in p2p systems. In: *Proceedings of SPIE, Multimedia Computing and Networking Conference, MMCN* (2008)
15. Osborne, M.J.: *An Introduction to Game Theory*. Oxford University Press, USA (2003)
16. Park, J., van der Schaar, M.: A game theoretic analysis of incentives in content production and sharing over peer-to-peer networks. *IEEE Journal of Selected Topics in Signal Processing* 4(4), 704–717 (2010)
17. Pianese, F., Perino, D., Keller, J., Biersack, E.W.: Pulse: An adaptive, incentive-based, unstructured p2p live streaming system. *IEEE Transactions on Multimedia* 9(8), 1645–1660 (2007)
18. Wu, W., Lui, J.C.S.: Exploring the optimal replication strategy in p2p-vod systems: Characterization and evaluation. In: *Proceedings of the 30th IEEE International Conference on Computer Communications, INFOCOM* (2011)
19. Zhao, B.Q., Lui, J.C.S., Chiu, D.-M.: Analysis of adaptive incentive protocols for p2p networks. In: *Proceedings of the 28th IEEE International Conference on Computer Communications, INFOCOM* (2009)