ARTICLE IN PRESS

Computer Networks xxx (2013) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation

Weijie Wu [a], John C.S. Lui [b,*], Richard T.B. Ma [c,d]

[a] School of Information Security Engineering, Shanghai Jiao Tong University, Shanghai 200240, China
[b] Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong
[c] School of Computing, National University of Singapore, 117418 Singapore, Singapore
[d] Advanced Digital Sciences Center, Illinois at Singapore, 138632 Singapore, Singapore

ARTICLE INFO

Article history:
Received 14 January 2012
Received in revised form 20 February 2013
Accepted 23 February 2013
Available online xxxx

Keywords:
P2P-VoD
Incentive
Mean-field
Stackelberg game

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 other 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* and study the long term behavior of the system under a repeated game setting. Via extensive simulations, we validate the stability and efficiency of our incentive scheme.

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 [5] and Video-on-Demand (VoD) systems like PPLive [1] and PPStream [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 [7] 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. Although plenty of work has been done for systems of traditional P2P applications, for example, the tit-for-tat [5] protocol has been well-adopted for file sharing applications, 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; and 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

* Corresponding author. Tel.: +852 3943 8407; fax: +852 2603 5024.
E-mail addresses: wuwjpku@gmail.com (W. Wu), cslui@cse.cuhk.edu.hk (J.C.S. Lui), tbma@comp.nus.edu.sg (R.T.B. Ma).

1389-1286/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2013.02.016

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 analyze the strategic interactions between the content provider and the peers using a repeated game model.
- We validate the effectiveness of our scheme and the theoretic results via extensive simulations.

Our paper is an extension of the earlier work [22] and 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. We discuss some practical issues in Section 6. Section 7 states the related work and Section 8 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. Since the complicated viewing operations supported by VoD systems, data demand and supply can be volatile and 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.

We denote K as the total number of videos in the system, and denote N as the total number of peers that watch any video in the system. We assume N is fixed but the peers can switch among different videos. 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 capacity 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 any particular video as I consecutive data segments. We denote S_{ik} as the ith data segment of video k. We model the user behavior of the system by specifying a set of rate transition probabilities $\{p_{ijk}:\ i,\ j=\ 0,1,\ldots,I;\ k=1,\ldots,K\}$. Each p_{ijk} denotes the transition probability of a typical user watching S_{jk} after finishing segment S_{ik}. In particular, p_{i0k} denotes the probability that a user quits video k after finishing S_{ik}. When this happens, this peer chooses another video to watch. Denote by ρ_l as the probability that this peer chooses video l. However, users might not start from the very first segment S_1, because they might have watched part of the video before, or they intend to skip the titles/advertisements. Thus, we let p_{0jl} denote the probability that a new arrival peer to video l starts with S_{jl}. In particular, $p_{i0k}\rho_l p_{0jl}$ denotes the probability that a peer departs from watching S_{ik} to watching S_{jl}, $k \neq l$. To keep consistency, we define $p_{00k} = 0$ and require $\sum_{j=0}^{I} p_{ijk} = 1$ for all $i = 1,\ldots,I$; $k = 1,\ldots,K$. In practice, the probability p_{ijk} 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 transiting from segment ik to segment jk, a peer only provides S_{ik} to other peers if needed when it is viewing S_{jk}. 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 ik, we assume that the system will direct the demand to the peers that can contribute S_{ik} first, which upload S_{ik} at an aggregate rate that equals the playback rate r. If the supply capacity of S_{ik} 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

$$W(u) = \int_0^u w(x)dx, \tag{1}$$

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

$$W(u) = wu, \tag{2}$$

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 ik when it is watching segment S_{ik}. Thus, each peer's viewing behavior becomes a random process. For a system with N peers and K videos (each containing I segments), the size of the state space of the system is $(IK)^N$. 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* [3,4] to calculate the steady state distribution of this peers in the P2P-VoD system. For the system with N peers, we denote $q_{ik}^N(t) \in [0,1]$ as the fraction of peers in state ik at time t. The system state at time t can be specified by the vector $[q_{ik}^N(t)], i = 1,\dots,I, k = 1,\dots,K$, where $\sum_{i=1}^{I}\sum_{k=1}^{K} q_{ik}^N(t) = 1$.

We first have the following theorem, which characterizes the stochastic system by deterministic limits when the system size is large.

Theorem 1. *For any given initial state* $[q_{ik}^N(0)] = [q_{ik}(0)]$, $i = 1,\cdots,I, k = 1,\cdots,K$, *define* $q_{ik}(t+1)$ *iteratively by the initial value* $q_{ik}^N(0)$ for $t \geqslant 0$:

$$q_{ik}(t+1) = \sum_{l=1}^{K}\sum_{j=1}^{I} q_{jl}(t)p_{j0l}\rho_k p_{0ik} + \sum_{j=1}^{I} q_{jk}(t)p_{jik}. \tag{3}$$

Then, for any fixed time t, almost surely

$$\lim_{N\to\infty} q_{ik}^N(t) = q_{ik}(t). \tag{4}$$

Proof. We will use Theorem 4.1 in [4] to prove this claim. Consider a system with N peers at time t. The transition probability for a peer from state jl to ik is

$$\mathbb{P}^N(S_{jl} \to S_{ik}) = \begin{cases} p_{jik} + p_{j0k}\rho_k p_{0ik} & \text{if } k = l, \\ p_{j0l}\rho_k p_{0ik} & \text{otherwise.} \end{cases} \tag{5}$$

Note that this transition probability does not depend on N. Hence we have

$$\lim_{N\to\infty} \mathbb{P}^N(S_{jl} \to S_{ik}) = \begin{cases} p_{jik} + p_{j0k}\rho_k p_{0ik} & \text{if } k = l, \\ p_{j0l}\rho_k p_{0ik} & \text{otherwise.} \end{cases} \tag{6}$$

Also note that the initial state $[q_{ik}^N(0)]$ are given independently of N and hence almost surely converges to $[q_{ik}(0)]$ when $N \to \infty$. According to Theorem 4.1 in [4], for any fixed time t, almost surely Eq. (4) holds. □

The above theorem indicates that when the system size is large, this stochastic system can be accurately approximated by the following deterministic difference equation:

$$q_{ik}(t+1) - q_{ik}(t) = \sum_{l=1}^{K}\sum_{j=1}^{I} q_{jl}(t)p_{j0l}\rho_k p_{0ik} + \sum_{j=1}^{I} q_{jk}(t)p_{jik} - q_{ik}(t). \tag{7}$$

If this difference equation converges, then we call $q_{ik} = \lim_{t\to\infty} q_{ik}(t)$ as the *steady state fraction* of peers in state ik. In the following theorem, we compute these fractions.

Theorem 2. *If the P2P-VoD system does not support the rewind operation, i.e. $p_{ijk} = 0$ for all $1 \leqslant j \leqslant i$. Then*

$$q_{ik} = \rho_k \frac{P_{0ik}}{\sum_{j=1}^{I} P_{0jk}} \quad \forall i = 1,\ldots,I, k = 1,\ldots,K, \tag{8}$$

where P_{ijk} denotes the aggregated probability of transiting from state ik to jk, which can be defined recursively as follows.

$$P_{ijk} = \begin{cases} p_{ijk} & \text{if } j = i+1, \\ \sum_{m=i+1}^{j-1} p_{imk} P_{mjk} + p_{ijk} & \text{otherwise.} \end{cases} \tag{9}$$

Proof. According to Theorem 1, at any given time, the fractions of peers in each state converges to

$$q_{ik}(t+1) = \sum_{l=1}^{K}\sum_{j=1}^{I} q_{jl}(t) p_{j0l} \rho_k p_{0ik} + \sum_{j=1}^{I} q_{jk}(t) p_{jik}. \tag{10}$$

In the steady state, this limiting fraction should not change over time. Hence,

$$q_{ik} = \sum_{l=1}^{K}\sum_{j=1}^{I} q_{jl} p_{j0l} \rho_k p_{0ik} + \sum_{j=1}^{I} q_{jk} p_{jik}. \tag{11}$$

Noting that $\sum_{l=1}^{K}\sum_{j=1}^{I} q_{jl} p_{j0l}$ is a constant and we denote it by D. Then, we have

$$q_{ik} = D\rho_k p_{0ik} + \sum_{j=1}^{I} q_{jk} p_{jik}. \tag{12}$$

Given $p_{ijk} = 0$ for all $1 \leqslant j \leqslant i$, we have $q_{1k} = D\rho_k\ p_{01k}$ and $q_{ik} = D\rho_k p_{0ik} + \sum_{j=1}^{i-1} p_{jik} q_{jk}$. By recursively solving q_{ik} and requiring $\sum_{k=1}^{K}\sum_{i=1}^{I} q_{ik} = 1$, we can derive the above formula. □

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_{ik} peers watch segment S_{ik}. Given a required playback rate of r, the aggregate required upload capacity for S_{ik} should be $Nq_{ik}r$. In the proof of Theorem 2, we have derived $q_{ik} = D\rho_k p_{0ik} + \sum_{j=1}^{I} q_{jk} p_{jik}$, or equivalently, $Nq_{ik} = ND\rho_k p_{0ik} + N\sum_{j=1}^{I} q_{jk} p_{jik}$. Among the peers watching S_{ik}, $ND\rho_k p_{0ik}$ peers transit to S_{ik} from watching another video, and $Np_{jik}q_{jk}$ peers have viewed S_{jk} before transiting to watch S_{ik}. Note that we can only assure that peers performing continuous play, i.e., moving from S_{ik} to $S_{i+1,k}$, have watched and therefore cached the whole segment of S_{ik}. Under our simplistic caching policy, only the peers that have transitioned from S_{ik} to $S_{i+1,k}$ and currently watching $S_{i+1,k}$ can upload S_{ik} to other peers. Suppose each peer contributes u amount of capacity for uploading available video segments, the total available upload capacity for S_{ik} would be $Nq_{ik}p_{i,i+1,k}u$. If the dedicated peer contribution $Nq_{ik}p_{i,i+1,k}u$ is less than the required download capacity $Nq_{ik}r$, the content provider needs to upload segment S_{ik} to support the difference in capacity. In particular, the playback requirement of the last segment, $Nq_{Ik}r$, 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:

$$C_s(u) = c_s N \times \sum_{k=1}^{K}\left[\sum_{i=1}^{I-1} q_{ik}(r - p_{i,i+1,k}u)^+ + q_{Ik}r\right]. \tag{13}$$

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

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

Notice that the maximum cost is $C_s(0) = c_sNr$ when the peers do not contribute any capacity and the minimum cost is $C_s(u) = c_sNr\sum_{k=1}^{K} q_{Ik}$ when u is large enough. In particular, when $u \geqslant r$, 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 homogeneous 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-bounds 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:

$$\pi_s(w,u) = -C_s(u) - wuN. \tag{14}$$

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

$$\pi_p(u,w) = wu - C_p(u), \tag{15}$$

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 solve $\max_u \pi_p(u,w)$. 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. □

We consider a Stackelberg game [17] 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 [17]. In particular, the peers solve the problem $u^*(w) = \text{argmax}_u \pi_p(u,w)$ given any w. By knowing the peers' best responses, the content provider solves the problem $w^* = \text{argmax}_w \pi_s(w,u^*(w))$. 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 3. *If u^* is a solution to the following problem:*

$$\min_u C_s(u) + NuC'_p(u), \tag{16}$$

then there exists a Stackelberg equilibrium $\left(u^*, u^*C'_p(u^*)\right)$; further, if (u^*,w^*) is a Stackelberg equilibrium, then u^* is the solution to problem (16).

Proof. We start by showing the first half of the statement. Denote $u^* = \text{argmin}_u[C_s(u) + NuC'_p(u)]$ 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 $C'_p(0) \leqslant w_0 \leqslant C'_p(\bar{u})$; 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

$$\begin{aligned} C_s(u_0) + Nu_0C'_p(u_0) &\leqslant C_s(u_0) + Nu_0w_0 \\ &< C_s(u^*) + Nu^*w^* \\ &= C_s(u^*) + Nu^*C'_p(u^*). \end{aligned} \tag{17}$$

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 (16). 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 equilibrium.[1]

To show the second half of the statement, suppose there exists a Stackelberg equilibrium (u^*,w^*), but u^* is not a solution to (16), i.e., there exists $u_0 \neq u^*$ such that $C_s(u_0) + NuC'_p(u_0) < C_s(u^*) + NuC'_p(u^*)$. Assume the content provider sets $w_0 = u_0C'_p(u_0)$. Taking the derivative in (15) 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 (16). □

Theorem 4. *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 $u^*(w) = \text{argmax}_u \pi_p(u,w)$ 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 $w^* = \text{argmax}_w \pi_s(w,u^*(w))$ 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 (16) is a strictly convex minimization over a compact set, which has a unique solution. According to Lemma 3, any Stackelberg equilibrium (u^*,w^*) satisfies that u^* is a solution to (16). Therefore, we conclude that the peers' solution in the Stackelberg equilibrium is unique.[2] □

[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 $0 \leqslant w^* \leqslant C'_p(0)$ 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 $0 \leqslant w^* \leqslant C'_p(0)$ is a Stackelberg equilibrium. When $u^* > 0$, the Stackelberg equilibrium is unique, where the content provider sets $w^* = u^*C'_p(u^*)$.

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016

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.

4.3. Efficiency of Stackelberg equilibrium

Now we discuss the efficiency of the Stackelberg equilibrium. For mathematical simplicity, in this subsection, we assume $C_s(u)$ and $C_p(u)$ are both twice differentiable in u.

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

$$\pi_w(u) = \pi_s + N\pi_p = -C_s(u) - NC_p(u). \tag{18}$$

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 $u_w = \mathrm{argmax}_u \pi_w(u)$, and u^* as the peers' solution at the Stackelberg equilibrium. We first state the following lemma:

Lemma 5. *The peers' upload contribution at the Stackelberg equilibrium is no larger than the upload capacity that maximizes the social welfare, i.e., $u^* \leqslant u_w$.*

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 3, 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

$$C'_{SW}(u) = C'_s(u) + NC'_p(u), \tag{19}$$
$$C'_{SE}(u) = C'_s(u) + NC'_p(u) + NuC''_p(u). \tag{20}$$

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 $NuC''_p(u) \geqslant 0$, 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, $u^* \leqslant u_w$.
Combining the results in the above two cases, we have $u^* \leqslant u_w$. □

We define the *bandwidth utilization* (*BU*) as the ratio of the bandwidth dedication at the Stackelberg equilibrium, to the value which maximizes the social welfare, i.e.,

$$BU = \frac{u^*}{u_w}. \tag{21}$$

According to the above lemma, we know that in general, $BU \leqslant 1$, and when BU approaches 1, it indicates that the system is in an efficient state. In particular, we have the following theorem.

Theorem 6. *If $u^* > 0$ and $C''_p(u)$ is increasing in u, then $BU \geqslant \frac{1}{2}$.*

Proof. Since $u_w \geqslant u^* > 0$, there must exist unique u_{SW} and u_{SE} such that $C'_{SW}(u_{SW}) = 0, C'_{SE}(u_{SE}) = 0$. Thus, we have $u_w = \max(u_{SW}, \bar{u})$ and $u^* = \max(u_{SE}, \bar{u})$.

Since $u_w \geqslant u^*$, we have $BU = \frac{u^*}{u_w} \geqslant \frac{u_{SE}}{u_{SW}}$. Noting the form of $C'_{SW}(u)$ and $C'_{SE}(u)$, we have

$$C'_s(u_{SW}) + NC'_p(u_{SW}) = 0, \tag{22}$$
$$C'_s(u_{SE}) + NC'_p(u_{SE}) + Nu_{SE}C''_p(u_{SE}) = 0. \tag{23}$$

Since $C'_s(u)$ and $C'_p(u)$ are both continuous functions in u, there must exist $u_1, u_2 \in [u_{SE}, u_{SW}]$ such that

$$C''_s(u_1)(u_{SW} - u_{SE}) + NC''_p(u_2)(u_{SW} - u_{SE}) - Nu_{SE}C''_p(u_{SE}) = 0. \tag{24}$$

Hence, we have

$$\frac{u_{SW} - u_{SE}}{u_{SE}} = \frac{NC''_p(u_{SE})}{C''_s(u_1) + NC''_p(u_2)} \leqslant \frac{C''_p(u_{SE})}{C''_p(u_2)}. \tag{25}$$

Since $C_p(u)$ is increasing in u and $u_{SE} \leqslant u_2$, we have

$$\frac{u_{SW}}{u_{SE}} \leqslant 1 + \frac{C''_p(u_{SE})}{C''_p(u_2)} \leqslant 2, \tag{26}$$

and thus

$$BU = \frac{u^*}{u_w} \geqslant \frac{u_{SE}}{u_{SW}} \geqslant \frac{1}{2} \quad \square. \tag{27}$$

The above theorem requires that $C''_p(u)$ increases in u. If this condition does not satisfy, it is in general difficult to characterize BU. In the following theorem, we choose a special form of the peers' cost function and derive a corresponding efficiency bound.

Theorem 7. *If $u^* > 0$ and $C_p(u) = c_p u^\beta (1 \leqslant \beta \leqslant 2)$, then $BU > \frac{\beta - 1}{\beta}$.*

Proof. Using the similar approach in the previous proof, we have

$$\frac{u_{SW}}{u_{SE}} \leqslant 1 + \frac{C''_p(u_{SE})}{C''_p(u_2)} \leqslant 1 + \frac{C''_p(u_{SE})}{C''_p(u_{SW})}, \tag{28}$$

where $u_2 \in [u_{SE}, u_{SW}]$. The second "$\leqslant$" holds because $C''_p(u)$ is a decreasing function in u and $u_2 \leqslant u_{SW}$.

Noting the form of $C_p(u)$, we have

$$\frac{u_{SW}}{u_{SE}} \leqslant 1 + \frac{u_{SE}^{\beta-2}}{u_{SW}^{\beta-2}} = 1 + \left(\frac{u_{SW}}{u_{SE}}\right)^{2-\beta}. \tag{29}$$

Define $x = \frac{u_{SW}}{u_{SE}}$ and $f(x) = x - x^{2-\beta}$. Let $g(x) = f(x) - (\beta - 1)x + \beta - 1$. Then it is easy to verify that $g'(x) > 0$, $\forall x \geqslant 1$ and $g(0) = \beta - 1 > 0$. Hence, $f(x) > (\beta - 1)x - (\beta - 1)$, $\forall x \geqslant 1$.

Hence, in order to satisfy $f(x) \leqslant 1$, we must have

$$(\beta - 1)x - (\beta - 1) < 1, \tag{30}$$

which indicates $x = \frac{u_{SW}}{u_{SE}} < \frac{\beta}{\beta-1}$. Therefore,

$$BU = \frac{u^*}{u_w} \geqslant \frac{u_{SE}}{u_{SW}} > \frac{\beta - 1}{\beta} \quad \square. \tag{31}$$

Another important measure on the efficiency of the equilibrium is the *price of anarchy* (*PoA*) [12]. In this paper, we define it 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

$$PoA = \frac{\pi_w(u^*)}{\pi_w(u_w)}, \tag{32}$$

where $u_w = \text{argmax}_u \pi_w(u)$ 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.

Theorem 8. *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., PoA = 1.*

Proof. If $u^* = 0$, then $C'_{SE}(0) \geqslant 0$. 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 Eqs. (19) and (20), we have $C'_{SE}(0) = C'_{SW}(0) < 0$, which contradicts to $C'_{SE}(0) \geqslant 0$. Hence, $u_w = 0 = u^*$, so *PoA* = 1.

If $u^* = \bar{u}$, then by Lemma 5, we have $u_w \geqslant u^* = \bar{u}$. In the meanwhile, $u_w \leqslant \bar{u}$, 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 $\pi_w(u_w) - \pi_w(u^*) \leqslant \pi'_w(u^*)(u_w - u^*)$, and PoA$\leqslant 1 - \frac{\pi'_w(u^*)}{\pi_w(u_w)}(u_w - u^*)$. It is also intuitive to investigate that a higher bandwidth utilization (i.e., larger *BU*) indicates a lower *PoA*, i.e., better system efficiency.

We would like to emphasize that the results derived in this section are only loose bounds on the bandwidth utilization. In practical systems, the performance can be very near to optimal. 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

$$\max_{w(x)} \pi_s(w(x), u) = -C_s(u) - Nu \int_0^u w(x)dx, \tag{33}$$

and the peers' problem is

$$\max_u \pi_p(u, w(x)) = u \int_0^u w(x)dx - C_p(u). \tag{34}$$

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

Theorem 9. *If the content provider can find $u^* = \text{argmax}_u - C_s(u) - NC_p(u)$, then any reward scheme satisfying the following property can make the content provider's utility arbitrarily close to the maximal:*

$$W(u) = \int_0^u w(x)dx \begin{cases} \leqslant C_p(u) & \text{if } u \neq u^*, \\ = C_p(u) + \epsilon & \text{if } u = u^*, \end{cases} \tag{35}$$

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 $\pi_p(u^*, w(x)) = u^* \int_0^{u^*} w(x)dx - C_p(u^*) \geqslant 0$; 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 (33) is upper-bounded by $\text{argmax}_u - C_s(u) - NC_p(u)$. The content provider's utility achieved in Eq. (35) is $\text{argmax}_u - C_s(u) - NC_p(u) - \epsilon$, 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 (35). For example, we can design

$$w(x) = C'_p(x) + \epsilon\delta(u^*), \tag{36}$$

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

It is also interesting to note, using the general reward scheme in Theorem 9, 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.

4.5. Repeated game model

We have analyzed the one-shot interaction of content provider and the peers using a Stackelberg game. In general, this interaction can last a long time. Will the players follow the Stackelberg equilibrium solution in the long run? In this section, we use a *repeated game model* to discuss this issue.

Assume that the game is played infinitely long. At round t, the utilities of the content provider and the peers are denoted by $\pi_s(t)$ and $\pi_p(t)$ respectively. Their utilities in the repeated game are

$$\Pi_s = \sum_{t=1}^{\infty} \delta_s^t \pi_s(t), \quad \Pi_p = \sum_{t=1}^{\infty} \delta_p^t \pi_p(t); \tag{37}$$

where δ_s and δ_p denote their discount factors.

An interesting difference of the repeated game from the one-shot game is that the peers may have incentives to deviate from the Stackelberg equilibrium: the peers may threaten to punish the content provider (e.g., contributing zero upload capacity) unless the content provider sets the reward higher than that in Stackelberg equilibrium. We would like to emphasize that, in our following repeated game model, we do not require that the peers communicate with each other; but rather, any particular peer can threaten to punish the content provider. Although we assumed homogeneous peers with the same utility function, similar analysis can apply to systems with heterogeneous peers.

Assume the peers request the content provider to set $w = \tilde{w}$, and threaten that if the content provider refuses to do so, they will set $u = 0$. We first consider what are the possible interactions in a particular round. Responding to this threat, the content provider has two possible strategies: (1) it compromises and sets $w = \tilde{w}$; or (2) it resists the threat and sets $w = w^*$ where w^* is the Stackelberg equilibrium. We denote by $\mathcal{C}$ and $\mathcal{R}$ the compromising and resisting strategy of the content provider. If the content provider plays $\mathcal{R}$, then the peers have two possible choices: (1) they punish the content provider, i.e., setting $u = 0$; or (2) they do not carry out the threat at all; rather, they accept w^* and set $u^* = \mathrm{argmax}_u \pi_p(u,w^*)$. We denote by $\mathcal{P}$ and $\mathcal{A}$ the punishing and accepting strategy of the peers. If the content provider plays $\mathcal{C}$, then the peers will surely accept this reward and set $\tilde{u} = \mathrm{argmax}_u \pi_p(u, \tilde{w})$. We still use $\mathcal{A}$ to denote this accepting strategy for $u = \tilde{u}$. Therefore, there are three cases of possible interactions in a particular round: $(\mathcal{R}, \mathcal{P}), (\mathcal{R}, \mathcal{A})$ and $(\mathcal{C}, \mathcal{A})$, where the two elements in each pair denote the content provider's and the peers' strategies, respectively. We use $\hat{\pi}_s, \pi_s^*$ and $\tilde{\pi}_s$ (resp. $\hat{\pi}_p, \pi_p^*$ and $\tilde{\pi}_p$) to denote the one-shot utility of the content provider (resp. the peers) for the above three cases respectively. In general, we have $\pi_s^* > \tilde{\pi}_s > \hat{\pi}_s$ and $\tilde{\pi}_p > \pi_p^* > \hat{\pi}_p = 0$.

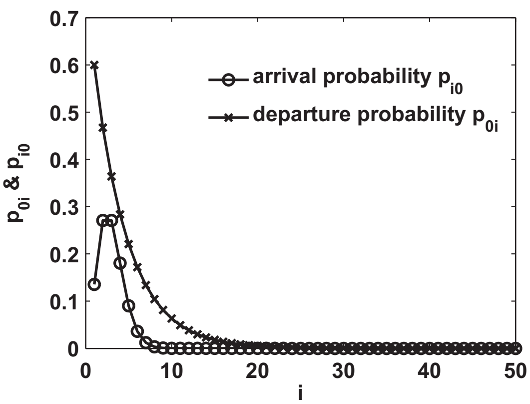

Fig. 1. Arrival and departure probability for each video segment.

We define a threat to be *credible* if, under the triggering condition, the threat-claimer would obtain no less utility by indeed carrying out the threat than not carrying out the threat. In particular, the peers' threat is credible if the peers can achieve higher (or at least equal) utility by playing $\mathcal{P}$ compared to $\mathcal{A}$ if the content provider plays $\mathcal{R}$ in an arbitrary round. Now we discuss under what condition the peers' threat is credible.

Theorem 10. *If* $\left\lfloor \log_{\delta_s} \frac{\tilde{\pi}_s - \hat{\pi}_s}{\pi_s^* - \hat{\pi}_s} \right\rfloor \leqslant \left\lfloor \log_{\delta_p} \frac{\pi_p^* - \hat{\pi}_p}{\tilde{\pi}_p - \hat{\pi}_p} \right\rfloor$, *then the peers' threat is credible; otherwise, it is not credible.*

Proof. Suppose the content provider and the peers have played $(\mathcal{C}, \mathcal{A})$ in the first $(t_0 - 1)$ rounds, and that the content provider plays $\mathcal{R}$ in round t_0. There are two possibilities regarding the peers: (1) the peers play $\mathcal{P}$, resulting that $\pi_s(t_0) = \hat{\pi}_s$ and $\pi_p(t_0) = \hat{\pi}_p$; or (2) the peers play $\mathcal{A}$, resulting $\pi_s(t_0) = \pi_s^*$ and $\pi_p(t_0) = \pi_p^*$. Whether the peers really punish the content provider depends on how many rounds they need to play $\mathcal{P}$ before the content provider plays $\mathcal{C}$. In particular, denote by t_p the number of rounds the peer can afford playing $\mathcal{P}$. After that, the content provider changes back to $\mathcal{C}$. We have:

$$\sum_{t=1}^{t_0-1} \delta_p^t \tilde{\pi}_p + \sum_{t=t_0}^{t_0+t_p-1} \delta_p^t \hat{\pi}_p + \sum_{t=t_0+t_p}^{\infty} \delta_p^t \tilde{\pi}_p \geqslant \sum_{t=1}^{t_0-1} \delta_p^t \tilde{\pi}_p + \sum_{t=t_0}^{\infty} \delta_p^t \pi_p^*. \tag{38}$$

The left side term represents a particular peer's total utility if the peers play t_p rounds of punishing strategy $\mathcal{P}$ when the content provider plays resisting strategy $\mathcal{R}$, and after that, the content provider compromises to set back $w = \tilde{w}$; the right side term represents a particular peer's total utility if the peers accept $w = w^*$ and set $u = u^*$ from t_0 onwards. If (38) holds, it is beneficial for the peers to play $\mathcal{P}$ for t_p rounds; otherwise, the peers are better off by playing $\mathcal{A}$ from t_0 onwards.

Similarly, denote by t_s the number of rounds that the content provider can afford playing $\mathcal{R}$ and being punished by the peers. After that, the peers play $\mathcal{A}$ from round $t_0 + t_s$ onwards and the content provider continues playing $\mathcal{R}$. We have:

$$\sum_{t=1}^{t_0-1} \delta_s^t \tilde{\pi}_s + \sum_{t=t_0}^{t_0+t_s-1} \delta_s^t \hat{\pi}_s + \sum_{t=t_0+t_s}^{\infty} \delta_s^t \pi_s^* \geqslant \sum_{t=1}^{\infty} \delta_s^t \tilde{\pi}_s. \tag{39}$$

If this condition holds, then the content provider has an incentive to play $\mathcal{R}$ from round t_0 onwards; otherwise, it is better for the content provider to play $\mathcal{C}$ in all time rounds.

The maximal value satisfying the above inequalities are $t_s = \left\lfloor \log_{\delta_s} \frac{\tilde{\pi}_s - \hat{\pi}_s}{\pi_s^* - \hat{\pi}_s} \right\rfloor$ and $t_p = \left\lfloor \log_{\delta_p} \frac{\pi_p^* - \hat{\pi}_p}{\tilde{\pi}_p - \hat{\pi}_p} \right\rfloor$. When $t_s > t_p$, it means the peers will give up playing $\mathcal{P}$ before the content provider changes back to $\mathcal{C}$. In this case, the threat of the peers is not credible: it is better for the peers to accept $w = w^*$ when the content provider plays $\mathcal{R}$ in round t_0. Otherwise, i.e., when $t_s \leqslant t_p$, the peers can play $\mathcal{P}$ for no larger than t_p rounds whereas the content server has to compromise and accept $w = \tilde{w}$, and therefore, the threat is credible. □

From Theorem 10, we can see that a larger gap between $\tilde{\pi}_p$ and π_p^*, or a larger discount factor δ_p, will induce a larger t_p, implying that the peers' threat is more likely to be credible. Physically, it means if the potential benefit after threatening is large, or if the peers care much about utility in future, the peers will have higher incentives to punish the content provider in a few rounds so as to force it to reward more than in the Stackelberg equilibrium.

We briefly conclude the result in our game theoretic analysis. The one-shot interaction of the content provider and the peer can be viewed as a Stackelberg game where the content provider takes the first action and the peer follows. 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. Considering the long term effect, a peer may perform "wiser" in a repeated game than in the one-shot Stackelberg game.

5. Performance evaluation

In this section, we validate the stability and efficiency of our incentive scheme via extensive simulations. Since the videos are independent in terms of peers' caching and viewing behaviors, in this simulation we focus on one particular video and omit the subscript k for video indices. The performance of a multiple video system is just a combination of the corresponding single video systems.

5.1. Simulation settings

In this section, we have the following settings:

- The system consists of $N = 10{,}000$ peers, one server and one video with $I = 50$ segments and playback rate $r = 500$ Kbps.
- The peers' external arrival probability to segment i is $p_{0i} = \frac{2^{i-1}}{(i-1)!} e^{-2} (1 \leqslant i \leqslant I)$; the probability of doing play operation is $p_{i,i+1} = 1 - 0.6e^{-0.25\times i} (1 \leqslant \ i \leqslant I - 1)$; and the probability of quit operation is $p_{i0} = 1 - p_{i,i+1}$ $(1 \leqslant i \leqslant I - 1)$, $p_{I0} = 1$. 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's strategy $w \in [0,1]$, and the peer's strategy $u \in [0,1000]$ Kbps.

In what follows, we evaluate the performance of our incentive mechanism design by showing the *Stackelberg equilibrium* and *its efficiency*. We would like to mention that we do not compare the performance of our design with the widely-used tit-for-tat mechanism since the performance measures are hardly comparable. On the one hand, if the system applies tit-for-tat, then the content server must not upload any data to users even if the required download rate is not achieved by the users; otherwise, all users can always get help from the content server so that they will set their upload bandwidth as zero and there is no actual exchange among peers. On the other hand, due to the high volatility between demand and supply of data, peers are rarely able to exchange the required data among themselves, which indicates low download rate. To summarize, tit-for-tat mechanism induces poor quality of service but zero upload consumption at the server. To summarize, our mechanism is significantly different from tit-for-tat and it is hard to find a comparable performance measure. Therefore, we do not compare the performance of our design with tit-for-tat mechanism.

5.2. Performance evaluation of our incentive scheme

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 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. (13)) 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. (14)) 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

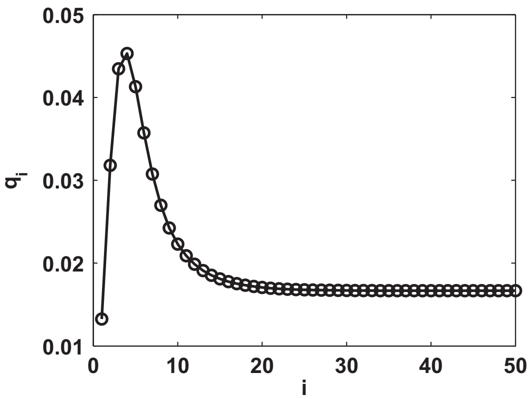

Fig. 2. Fraction of peers in each video segments.

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016

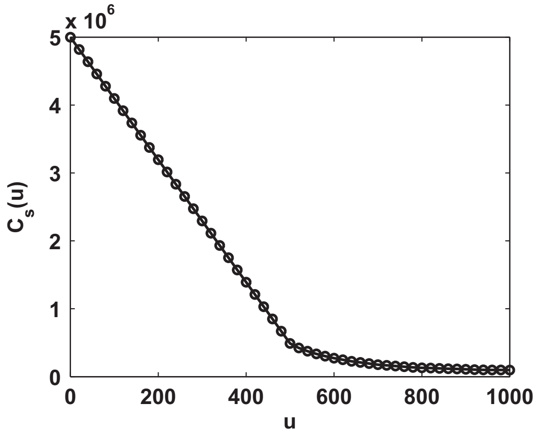

Fig. 3. Content provider's upload cost.

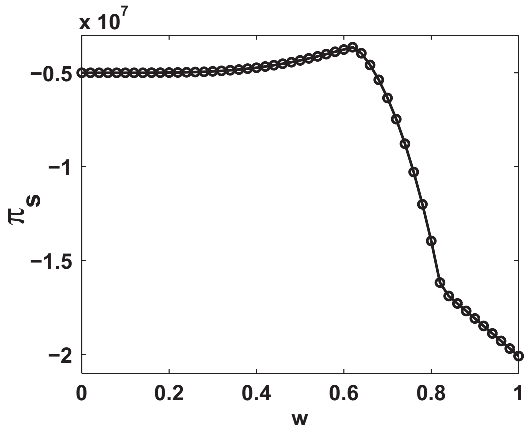

Fig. 4. Content provider's utility.

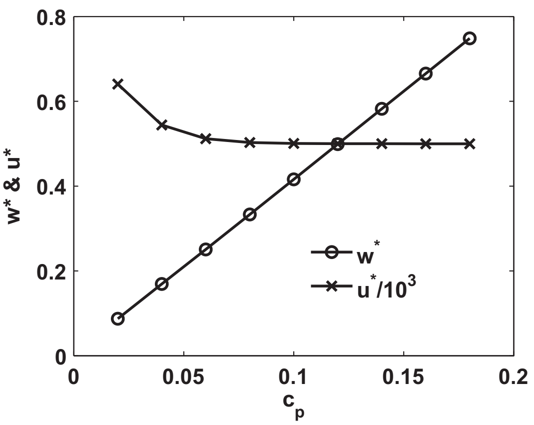

Fig. 5. Stackelberg equilibrium.

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.

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 $c_p \in [0.02, 0.18]$. 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 contribute less upload capacity at the Stackelberg equilibrium.

In Fig. 6, we evaluate the efficiency of the Stackelberg equilibrium when fixing $\beta = 1.2$ and varying $c_p \in [0.02, 0.18]$. In Fig. 7a, we compare the peers' upload capacity that maximizes the social welfare, with that at the Stackelberg equilibrium. We can observe that the two values are close to each other in general. We compare the maximal social welfare and the social welfare at Stackelberg equilibrium in Fig. 7b, which shows that the social welfare at the Stackelberg equilibrium is always very near to the maximal value, i.e., PoA ≈ 1 for any $c_p \in [0.02, 0.18]$. We do similar simulations in Fig. 7; the difference is we fix $c_p = 0.2$ and vary $\beta \in [1.06, 1.60]$. These simulation results validates the efficiency of our incentive scheme.

5.3. Impact of pre-fetching on the system performance

We have been focusing on the simple caching policy, i.e., a peer only caches the segment it is current watching. Practical systems may apply various caching policies where a peer caches some more segments in order to enlarge the possibility of successful contribution to other peers. In this subsection, we consider a natural extension of the simple caching policy, i.e., we consider pre-fetching where a peer caches the segment it is current watching, as well as one additional segment that it will watch next. For example, a peer also downloads S_{i+1} when it is watching S_i. However, we would like to point out that if a peer quits from the system, then its pre-fetched data are not useful anymore so that pre-fetching may incur wastage as well.

We still apply our linear reward scheme and study the Stackelberg equilibrium. In Fig. 8a, we show the Stackelberg equilibrium when all peers pre-fetch the next segment. Comparing to Fig. 5, we observe that we can still achieve Stackelberg equilibrium. However, we like point out that the peers may even need to set a larger upload capacity, in particular, when c_p is small. This is because the peers have a high probability of quitting the video service when they are watching the first few segments. If we let these peers pre-fetch the later segments, it increases the download requirement while with high probability, their pre-fetched data are not useful. Therefore, we modify our pre-fetching policy by letting the peers in the first five segments using the simple caching policy without pre-fetching while the rest of the peers pre-fetch the next segment. The rationale of doing so is the peers in early

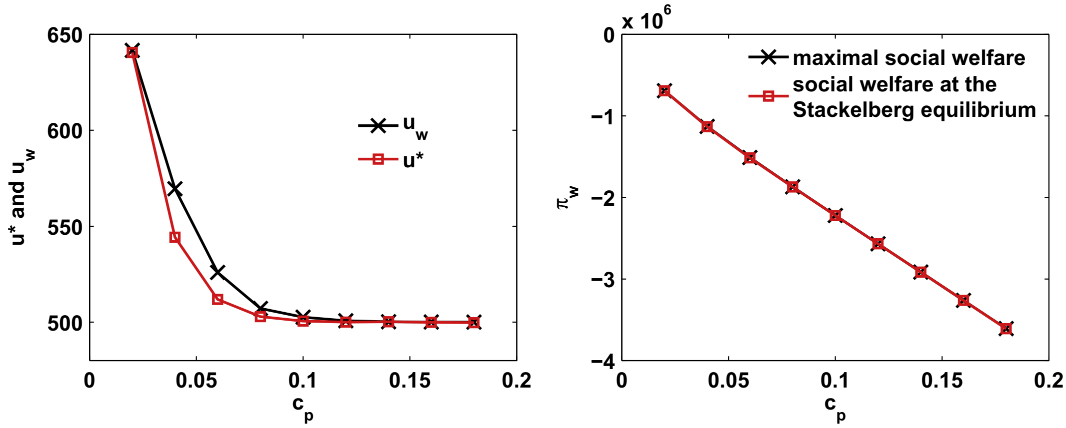

(a) comparison of peer's contribution (b) comparison of social welfare

Fig. 6. System efficiency at the Stackelberg equilibrium when varying c_p.

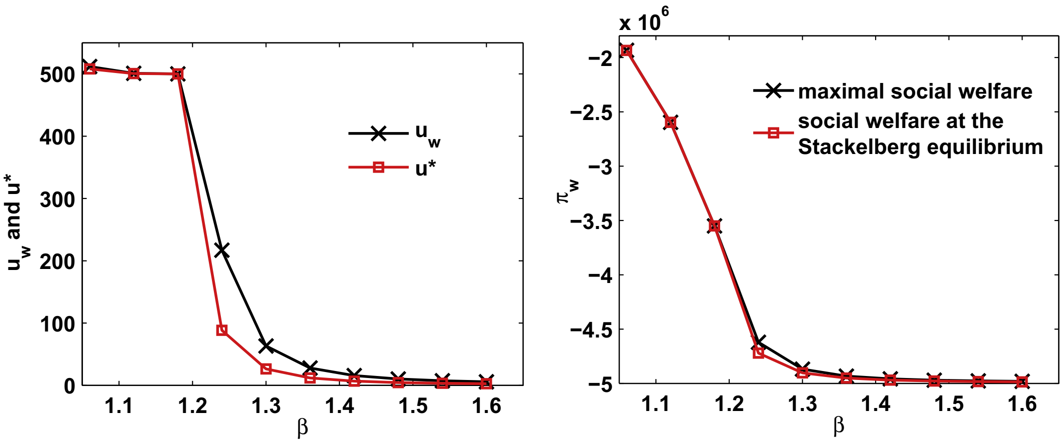

(a) comparison of peer's contribution (b) comparison of social welfare

Fig. 7. System efficiency at the Stackelberg equilibrium when varying β.

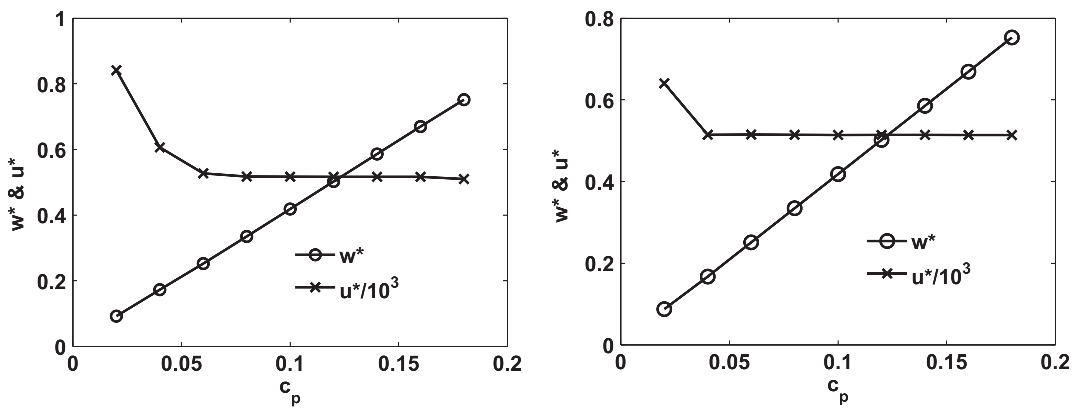

(a) Stackelberg equilibrium with full pre-fetching (b) Stackelberg equilibrium with modified pre-fetching

Fig. 8. Impact of pre-fetching.

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016

segments are more likely to quit the system. We show the corresponding Stackelberg equilibrium in Fig. 8b. We observe that the feature of Stackelberg equilibrium remains as shown in Fig. 5, while peers need to allocate a bit less upload capacity at the Stackelberg equilibrium due to the help of pre-fetching.

We would like to remark that in practical systems, prefetching may have a larger impact on reducing the upload requirement of the server. Unlike our ideal assumptions, peers in a P2P-VoD system may not be fully connected. Therefore, it is not always possible to fully utilize a peer's upload bandwidth. By caching more segments, it is more flexible to find the downloaders for a particular peer, so that the utilization of the peers' bandwidth can be more efficient.

5.4. Impact of user engagement on the system performance

In the above analysis and simulations, we assume a fixed number of peers in the system, while practical systems may be different. In often times, the content server of a P2P-VoD system has a limited upload bandwidth capacity, so it may not always be able to serve all peers due to various reasons. Therefore, if peers do not contribute enough upload bandwidth, then the system may provide service of poor performance to users (i.e., download rate is less than playback rate), and hence impacts the users' engagement, e.g., some peers may leave the system. We will discuss this practical issue in the next section.

In this section, we use simulation to show how the users' engagement may impact the system performance. For simplicity, we assume that the server has a fixed amount of upload bandwidth (and hence, a fixed amount of upload cost), but the content provider's utility changes with respect to the number of peers in the system. This is due to the fact that a large fraction of profit is advertisement income which highly depends on the market share of the VoD service. In particular, we have the following assumptions.

- The server's upload capacity is limited at 100 Mbps. The peers' download rate may be less than the video's playback rate.
- If the peers' average download rate r' is less than the playback rate r, then a fraction $\gamma = r'/r$ of peers continue watching the video, while the rest $1-\gamma$ peers depart from the system.
- The content provider incurs a lost in its utility when the number of users decreases. In particular, the content provider's utility π_s decreases by $\alpha(1-\gamma)N$ if a fraction $1-\gamma$ peers leave the system. The constant α represents the marginal reduction of advertisement income when users leave the system. In the following simulation, we set $\alpha = 1000$ to cope with the setting of c_s and r.

In Fig. 9, we plot the Stackelberg equilibrium with consideration of the impact of user engagement. We can observe from the figure that the trend of Stackelberg equilibrium is the same as Fig. 5, but the difference is the content provider proposes higher reward to incentivize the peers to contribute more bandwidth, so as to improve the video quality and attracts more peers in the system.

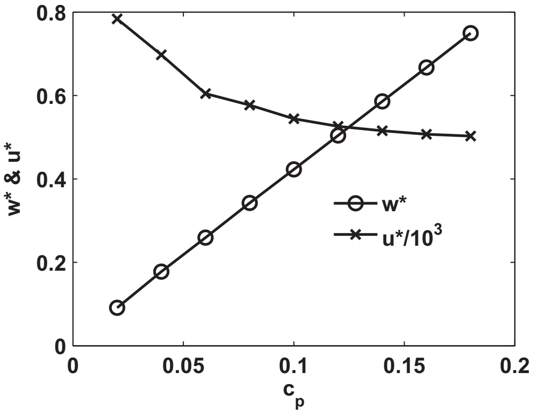

Fig. 9. Stackelberg equilibrium with uses' engagement.

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

6. Discussion on practical issues

In this section, we briefly discuss some practical implementation issues related to our incentive mechanism design. We will show that our simplistic assumptions are only for ease of mathematical expression, and can be simply extended to adapt to various practical problems.

6.1. System heterogeneity

Real systems consist of heterogeneous peers which have different sensitivities on the reward. Hence, peers may have different responses upon a given reward scheme. We want to point out that, our reward scheme design can be easily adapted to this heterogeneity. In particular, a typical peer solves

$$\max_{u} \pi_p(u, w) = wu - C_p(u) \tag{40}$$

for given unit reward w. Since peers are heterogeneous, they may have different cost functions $C_p(u)$ and hence have different optimal solutions $u^*(w)$. From the content provider's point of view, it wants to decide the unit reward w so as to maximize its own utility:

$$\max_{w} \pi_s(w, U^*(w)) = -C_s(U^*(w)) - wU^*(w), \tag{41}$$

where $U^*(w)$ denotes the total upload bandwidth dedication of all peers, which depends on the unit reward w. Further game analysis is of the content provider and all these peers. We can apply the similar approach of the previous sections and thus we omit the details here. It is also worth noting that videos can also be heterogeneous in terms of playback rate. Hence, the impact of peers' contribution can be correlated to the video it is currently watching. The content provider can also take into this consideration

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016

and propose different unit reward based which video this peer is watching.

6.2. P2P-VoD system with caching

In this paper, we apply a simplistic caching policy where a peer only caches one segment it watched in the previous time slot. Practical systems can better utilize the peers' upload bandwidth resources by letting them cache more data so as to enlarge the opportunity of contributing to other peers. These caching policies differ among systems according to the system's specific implementation and network environment. In the simulation section, we showed the applicability of our design to practical systems by illustrating some simple caching policies. In particular, we considered the pre-fetching mechanism, i.e., peers may buffer segments of the video before their playback. The rationale of doing so is to ensure that a smooth playback of the video, and at the same time, the peer may upload the pre-fetched content to other peers so as to further reduce the upload requirement of the content server.

Practical caching policy may be much more complicated, and its design is beyond the scope of this paper. How to incentivize the peers to cache the desired video data is another independent mechanism design problem, which we consider in a parallel work [23]. Nevertheless, the model in this paper provides some important insights. For example, the earlier few but not the very first segments (e.g., segment 2–6 in Fig. 2) of the video need the most replicas since a lot of peers are watching these segments, and many peers quit after watching the first few segments of the video.

In a system with more caches, the cost function of the content provider may not in the form of Eq. (13) any more, but should depend on the caching policy of the system. However, we want to emphasize that in general, it is still rational to assume that the content provider's cost function $C_p(u)$ is non-increasing and convex in the peers' contribution u. The non-increasing feature obviously holds in general. In reality, when peers contribute a small amount of upload resources, it is very easy for the system to utilize them efficiently; however, when u is large, the system may only utilize part of the resources and hence leads to the convexity feature of $C_p(u)$. Therefore, all the following analysis still applies to a general P2P-VoD system with caching.

6.3. Cheating proof guarantee

One of the rationale that we choose the upload capacity u, rather than the real upload contribution of a peer, as the criteria of reward, is for cheating proof guarantee. If peers are rewarded based on their upload contribution, then the peers may have incentives to form coalitions and upload garbage packages to each other so as to cheat for reward. In our design, a peer's reward is based on its dedicated upload capacity, and the system operator fully controls the upload resources. A peer does not have the incentive to find other peers to upload so as to increase its reward; it receives the reward even if it is not contributing. However, a peer must upload at the rate of its declared maximal value if it is assigned by the tracker to deliver data to any other peer. One can design punishment mechanism if a peer is found refusing to upload at its declared rate. Hence, we are not concerned on peers' coalition, but only investigate the upload bandwidth contribution of peers, which can be estimated by many existing tools [20].

6.4. Impact of video quality on users' viewing behaviors

In this paper, we assume the model where all peers can download the video data at the playback rate, so that peers do not need to wait for the data buffering. In real systems, the content server may not always be able to serve, so that peers may receive poor quality of service from time to time, either by receiving low resolution video, or by waiting due to data buffering. The quality of video may depend on various factors. In regarding to our design issues, the incentive scheme and caching policy we consider, decide the total bandwidth resource of the system, and hence impact the quality of service each user experiences.

It was observed in [6] that the video quality impacts the user engagement in video-over-the-Internet applications. Consequently, it in turn impacts the bandwidth resource of the system as a departed peer may not be able to contribute its upload capacity any more. Therefore, understanding the user behavior is important in designing practical incentive schemes. Authors in [6] reported some interesting results, for example, buffering ratio and join time are two important factors that impacts the engaging time and number of peers watching the video. However, up until now, we have only found measurement results but there is no explicit model to characterize these impacts. Due to the difficulty in modeling these factors, in this paper, we simply assume that the content server always upload the deficit bandwidth to all peers, such that the peers' viewing behavior is not influenced by the incentive or caching mechanisms. One future research direction is to study the interactions between user engagement, video quality and incentive mechanisms design.

7. Related work

Incentive issue has received plenty of attentions in P2P applications. Zhao et al. [24,25] 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 [18], 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 micro-payment in [8]. Misra et al. [15] proposed a Shapley value approach in incentive design using a cooperative game model. Reputation [9,11] is another well-known approach where a peer's reputation represents its history of contribution in the system. Ma et al. [14] proposed a service differentiation approach in P2P network based on the amount of contribution each node has provided to the network community.

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016

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 and Lui [21] analyzed how to efficiently utilize the peers' resources, but did not address how to incentivize peers to contribute their resources. Habib et al. [10] 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. [16] 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 [19], the authors presented a modified tit-for-tat mechanism; in [13], 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.

8. 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. Using a *mean field approximation* model, we show that the content provider's cost is a non-increasing and convex function in the peers' upload contributions. Using a Stackelberg game model, we convert the system equilibrium into an optimization problem and show the existence and uniqueness of its solution, as well as the theoretical efficiency bound of the equilibrium. Using a repeated game model, we analyze the strategic interactions of the content provider and the peers, pointing out the factors that impact the credibility of the threat. Using extensive simulations, we validate our model and show the high efficiency of our inventive mechanism.

Acknowledgement

The research of John C.S. Lui is supported in part by the Hong Kong RGC Research Funding #2150634 and #2150678.

References

[1] PPLive <www.pptv.com>.
[2] PPStream <www.ppstream.com>.
[3] M. Benaïm, J.-Y. Le Boudec, A class of mean field interaction models for computer and communication systems, Performance Evaluation 65 (2008) 823–838.
[4] J.-Y.L. Boudec, D. McDonald, J. Mundinger, A generic mean field convergence result for systems of interacting objects, in: Proceedings of the 4th International Conference on Quantitative Evaluation of Systems, Scotland, UK, 2007.
[5] B. Cohen, Incentives build robustness in bittorrent, in: Proceedings of the 1st Workshop on Economics of Peer-to-Peer Systems, 2003.
[6] F. Dobrian, A. Awan, D. Joseph, A. Ganjam, J. Zhan, V. Sekar, I. Stoica, H. Zhang, Understanding the impact of video quality on user engagement, in: Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (ACM SIGCOMM 2011), Toronto, Ontario, Canada, 2011.
[7] M. Feldman, C. Papadimitriou, J. Chuang, I. Stoica, Free-riding and whitewashing in peer-to-peer systems, IEEE Journal on Selected Areas in Communications 24 (5) (2006) 1010–1019.
[8] P. Golle, K. Leyton-Brown, I. Mironov, Incentives for sharing in peer-to-peer networks, in: Proceedings of the 3rd ACM Conference on Electronic Commerce (EC), 2001.
[9] M. Gupta, P. Judge, M. Ammar, 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), New York, NY, USA, 2003, pp. 144–152.
[10] A. Habib, J. Chuang, Service differentiated peer selection: an incentive mechanism for peer-to-peer media streaming, IEEE Transactions on Multimedia 8 (3) (2006) 610–621.
[11] S.D. Kamvar, M.T. Schlosser, H. Garcia-Molina, The eigentrust algorithm for reputation management in p2p networks, in: Proceedings of the 12th International Conference on World Wide Web (WWW), New York, NY, USA, 2003.
[12] E. Koutsoupias, C.H. Papadimitriou, Worst-case equilibria, Computer Science Review 3 (2) (2009) 65–69.
[13] Z. Liu, Y. Shen, S.S. Panwar, K.W. Ross, Y. Wang, 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), New York, NY, USA, 2007, pp. 311–316.
[14] R.T.B. Ma, S.C.M. Lee, J.C.S. Lui, D.K.Y. Yau, Incentive and service differentiation in p2p networks: a game theoretic approach, IEEE/ACM Transactions on Networking 14 (2006) 978–991.
[15] V. Misra, S. Ioannidis, A. Chaintreau, L. Massoulié, Incentivizing peer-assisted services: a fluid shapley value approach, in: Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), New York, NY, USA, 2010, pp. 215–226.
[16] J. Mol, J. Pouwelse, M. Meulpolder, D. Epema, H. Sips, Give-to-get: free-riding-resilient video-on-demand in p2p systems, in: Proceedings of SPIE, Multimedia Computing and Networking Conference (MMCN), 2008.
[17] M.J. Osborne, An Introduction to Game Theory, Oxford University Press, USA, 2003.
[18] J. Park, M. van der Schaar, 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) (2010) 704–717.
[19] F. Pianese, D. Perino, J. Keller, E.W. Biersack, Pulse: an adaptive, incentive-based, unstructured p2p live streaming system, IEEE Transactions on Multimedia 9 (8) (2007) 1645–1660.
[20] J. Strauss, D. Katabi, F. Kaashoek, A measurement study of available bandwidth estimation tools, in: Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement (IMC), 2003.
[21] W. Wu, J.C.S. Lui, Exploring the optimal replication strategy in p2p-vod systems: characterization and evaluation, IEEE Transactions on Parallel and Distributed Systems 23 (8) (2012) 1492–1503.
[22] W. Wu, J.C.S. Lui, R.T.B. Ma, Incentivizing the upload capacity in P2P-VoD systems: a game theoretic analysis, in: Proceeedings of the 3rd International Conference on Game Theory for Networks (GameNets), Shanghai, China, 2011.
[23] W. Wu, R.T.B. Ma, J.C.S. Lui, On incentivizing caching for P2P-VoD systems, in: Proceedings of the 7th Workshop on the Economics of Networks, Systems and Computation (NetEcon) (in conjunction with INFOCOM'12), Orlando, Florida, USA, 2012.
[24] B.Q. Zhao, J.C.S. Lui, D.-M. Chiu, Analysis of adaptive incentive protocols for p2p networks, in: Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM), 2009.
[25] B.Q. Zhao, J.C.S. Lui, D.-M. Chiu, A mathematical framework for analyzing adaptive incentive protocols in P2P networks, IEEE/ACM Transactions on Networking 20(2) (2012) 367–380.

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016

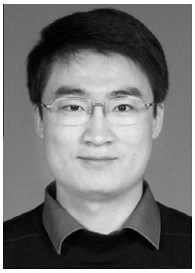

Weijie Wu received the B.Sc. degree in electronics from Peking University in 2008, and the Ph.D. degree in computer science from The Chinese University of Hong Kong in 2012. He is currently an assistant professor in School of Information Security Engineering, Shanghai Jiao Tong University. Before that, he was a postdoctoral research fellow in Department of Computer Science and Engineering at The Chinese University of Hong Kong. During his Ph.D. study, he worked as a research intern at National University of Singapore. His current research interests are in computer networks and security, including incentive mechanisms design, application of game theory on communication networks and security issues, performance evaluation and network economics.

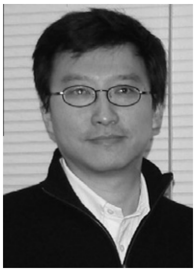

John C.S. Lui received his Ph.D. in Computer Science from UCLA. He is currently a professor in the CS&E Department at The Chinese University of Hong Kong. His research interests are in data networks, system and applied security, multimedia systems, network sciences and cloud computing. John received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award, as well as the co-recipient of the Best Student Paper Awards in the IFIP WG 7.3 Performance 2005 and the IEEE/IFIP Network Operations and Management (NOMS) Conference. He is an associate editor in the Performance Evaluation Journal, IEEE-TC, IEEE-TPDS and IEEE/ACM Transactions on Networking. John is a Fellow in ACM and IEEE, and chair of ACM SIGMETRICS.

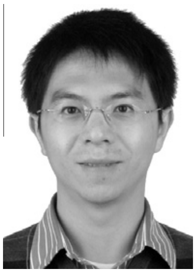

Richard T.B. Ma received the B.Sc. (first-class honors) degree in Computer Science in July 2002 and the M.Phil. degree in Computer Science and Engineering in July 2004, both from the Chinese University of Hong Kong, and the Ph.D. degree in Electrical Engineering in May 2010 from Columbia University. During his Ph.D. study, he worked as a research intern at IBM T.J. Watson Research Center in New York and Telefonica Research in Barcelona. He is currently a Research Scientist in Advanced Digital Science Center, University of Illinois at Urbana–Champaign and an Assistant Professor in School of Computing at National University of Singapore. His current research interests include distributed systems, network economics, game theory, and stochastic processes.

Please cite this article in press as: W. Wu et al., On incentivizing upload capacity in P2P-VoD systems: Design, analysis and evaluation, Comput. Netw. (2013), http://dx.doi.org/10.1016/j.comnet.2013.02.016