

Incorporating Contribution-Awareness into Mesh-based Peer-to-Peer Streaming Systems

Nazanin Magharei, Reza Rejaie and Yang Guo

Abstract

While Peer-to-Peer streaming have become increasingly popular during recent years, the proper allocation of available resources among peers in a resource constraint system remains a challenging problem. In a resource constraint system, the allocated resource and thus delivered quality to individual peers should be proportional to their contribution to the system, *i.e.*, resource allocation should be *contribution aware*. This in turn results in fairness among peers and encourages active contribution from participating peers.

In this paper, we present a contribution aware mechanism for mesh-based P2P streaming approaches. In our proposed mechanism, individual peers use a tax function to determine their number of parent peers (*i.e.*, their share of resources) based on the number of their child peers (*i.e.*, peers' contributed resources) and the aggregate available resources in the system. We examine the behavior of a commonly used tax function, and describe how the contribution aware mechanism can leverage the tax function. We conduct extensive session level simulations that properly incorporates churn and pairwise delay between peers. Our results not only demonstrate the ability of the proposed mechanism to properly allocate available resource among participating peers over a wide range of scenarios but also sheds insightful light on the dynamics of resource utilization and allocation.

I. INTRODUCTION

During recent years, Peer-to-Peer (P2P) overlays have become increasingly popular for scalable delivery of streaming content from a single source to a large number of receivers over the Internet [1]. In this approach that is generally known as *P2P streaming*, participating peers form an overlay over which individual peers contribute their outgoing bandwidth by forwarding a subset of their available content to their connected peers. P2P streaming approaches can be broadly divided into two classes: tree-based and mesh-based approach. In the tree-based approach, participating peers form one or multiple tree-shaped overlay(s) where each peer pushes a specific portion of the content (*e.g.*, a sub stream or a description) to its child peers (*e.g.*, [2]). In the mesh-based approach, participating peers often maintain a randomly connected mesh and incorporate swarming content delivery (*e.g.*, [3]).

N. Magharei and R. Rejaie are with the Department of Computer & Information Science, University of Oregon, {nazanin,reza}@cs.uoregon.edu

Yang Guo is with Thomson Lab, Yang.Guo@thomson.net

Participating peers in a P2P streaming session often have heterogeneous access link bandwidth. To accommodate bandwidth heterogeneity, a P2P streaming mechanism should ensure that delivered quality to individual peers is proportional to their incoming access link bandwidth. An elegant approach to achieve this goal is to encode the content using Multiple Description Coding (MDC) scheme. This enables a P2P streaming mechanism to deliver a proper number of descriptions to individual peers to maximize their delivered quality. However, in practice, the aggregate available resources (*i.e.*, outgoing bandwidth) in the system is often insufficient to maximize the delivered quality to individual peers due to various reasons such as asymmetric access link bandwidth, and lack of willingness to contribute resources among peers. A reasonable approach in such an environment is to ensure that delivered quality to individual peers is proportional to the amount of resources they actually contribute to the system (or their outgoing bandwidth) rather than their demand (or incoming bandwidth). In essence, the delivered quality by P2P streaming mechanisms in a resource-constraint environment should be *contribution aware*. Incorporating contribution-awareness into P2P streaming is non-trivial due to the distributed and dynamic nature of resources.

Recently, Sung et al. [4] proposed a contribution aware mechanism for tree-based P2P streaming mechanisms. In their approach, each peer uses a given tax function to determine the number of trees it is “entitled” to join based on (*i*) the aggregate available resources that are periodically collected from and reported to participating peers; and (*ii*) the number of child peers it serves (*i.e.*, the amount of contributed resources). Each peer also adaptively examines the possibility of joining more trees in order to properly use any excess resources in the system. As a parent each peer incorporate a preemption mechanism to properly allocate its resources between existing and new child peers.

This paper presents a contribution-aware mechanism for mesh-based P2P streaming mechanisms that is inspired by the proposed approach by Sung et al. [4]. We describe how such a contribution aware mechanism can be incorporated into mesh-based P2P streaming, and discuss several required modifications. We show how the proposed tax function allocates the available resources among participating peers for different values of available resources and tax rates. Furthermore, we identify the notion of “saturated region” when high bandwidth peers do not require their entitled share resources, and examine its effect on system behavior. The primary contribution of this paper is the extensive evaluations of the proposed mechanism well beyond previous studies [4] [5]. Toward this end, we use our session level P2P simulator that incorporates a realistic model for peer dynamics, to investigate the effect of key design parameters

such as tax rate and preemption policies over a wide range of scenarios. In particular, we examine how changes in aggregate available resources, distribution of contributed resources among peers, and group size affect different aspects of system performance including allocation and utilization of resources, and stability of the overlay.

Our simulations illustrate that our proposed contribution aware mechanism for mesh-based P2P streaming can properly allocate available resources among participating peers while achieving high utilization over a wide range of scenarios. We quantify the impact of each preemption policy to the stability of the overlay, and demonstrate the effect of tax rate and resource availability on the allocation of entitled and excess resources among peers. The contribution aware mechanism exhibits promising scaling behavior over examined group sizes. Finally, increasing the intervals of reporting the group state, primarily affects short-lived peers in the overlay.

The rest of this paper is organized as follows: Section II provides an overview of mesh-based P2P streaming. In Section III, we describe the proposed contribution aware mechanism. Section IV examines the behavior of the tax function. Our performance evaluation methodology along with our evaluation results and main findings are presented in Section V. Finally, we conclude the paper and sketch our future plans in Section VI.

II. MESH-BASED P2P STREAMING: BACKGROUND

To provide the required background, we present an overview of PRIME [3] as a representative mesh-based P2P streaming mechanism in this section. While the description and discussions are centered around PRIME, we believe that most of the issues and findings are generally applicable to other mesh-based P2P streaming systems. In the mesh-based P2P streaming, participating peers maintain a randomly connected and *directed* overlay, namely a *mesh*. Participating peers incorporate swarming content delivery (similar to BitTorrent [6]) where child peers pull their required packets from their parents. This approach has two advantages: *(i)* it is able to effectively utilize the outgoing bandwidth of participating peers as the group size grows, and *(ii)* it can gracefully accommodate the dynamics of peer participation (or churn).

To form an overlay, each peer maintains a certain number of parent peers from which it pulls its required content. Each peer also serves as a parent for a specific number of child peers and provides content to them. PRIME uses a central approach to parent discovery. All participating peers periodically send a heartbeat message to a well known bootstrap node that maintains the status of the entire group. When a peer requires more parents (*e.g.*, upon arrival or when the number of its parents simply drops

below a threshold), it contacts the bootstrap node to obtain a list of potential parents. In response to such a request, the bootstrap node provides the list of a random subset of those peers that can accommodate a new child peer. For each peer, we denote the number of parent and child peers as its *incoming* and *outgoing* degree, respectively.

To effectively utilize the access link bandwidth of participating peers, the incoming and outgoing degrees of participating peers should be proportional to their incoming (bw_{down}) and outgoing (bw_{up}) bandwidth [3]. Using the same ratio of incoming (or outgoing) bandwidth to incoming (or outgoing) degree for all peers implies that all connections have roughly the same average bandwidth which is called bandwidth-per-flow or bw_{pf} . bw_{pf} is a configuration parameter that is selected a priori and known by individual peers. More specifically, each peer tries to maintain its incoming and outgoing degrees at $\lfloor \frac{bw_{down}}{bw_{pf}} \rfloor$ and $\lfloor \frac{bw_{up}}{bw_{pf}} \rfloor$, respectively. Since all connections have roughly the same bandwidth, the amount of resources that a peer contributes or consumes in the overlay can be approximated with its outgoing and incoming degree, respectively.

III. CONTRIBUTION-AWARE P2P STREAMING

The primary goal of a contribution-aware mechanism is to enable individual peers to determine their share of available resources (mainly bandwidth) in the system based on the amount of resources they contribute as well as the aggregate amount of available resources in the system. Given the direct relationship between the (incoming and outgoing) bandwidth and the (incoming and outgoing) degree due to the bandwidth-degree constraint as described in Section II, the contribution-aware mechanism can be formulated as deriving the incoming degree of a peer based on its outgoing degree. More specifically, the goal of each peer is to determine its incoming degree (*i.e.*, the number of parents) based on (i) its outgoing degree (*i.e.*, the number of child peers), and (ii) the aggregate outgoing degree across all peers.

To support contribution awareness, each peer i uses a generic cost function[5] to determine its incoming degree R_i :

$$R_i = \frac{1}{t}W_i + \frac{t-1}{t} \sum_{i=1}^N \frac{W_i}{N}, \quad (1)$$

where t , N , and W_i denote the tax rate in the system, number of participating peers, and the outgoing degree that peer i is willing to contribute. In essence, R_i presents the “entitled” share of system resources for peer i and thus we refer to R_i as *entitled* degree. As shown in Eqn. (1), R_i is the sum of two terms. The first term represents the incoming degree of a peer due to its own contribution (W_i). Since the tax

rate is always greater than or equal to one ($t \geq 1$) to balance supply and demand for resource in the system, the outgoing degree of individual peers is higher than their incoming degree. This also implies that there are always extra resources in the system. The second term represents an even share of these extra resources among participating peers. This share of excess resources depends on the group state, namely group population (N) and the amount of aggregate available resources in the system ($\sum W_i$).

We assume that the tax rate t is a configuration parameter and thus known to each participating peer. If the group state information is known to individual peers, they can simply use Eqn. (1) to determine their entitled incoming degree. In subsection III-A, we describe a mechanism to collect the required group state information and to distribute them to participating peers.

In practice, the following two issues also contribute to the extra resources. First, when the aggregate incoming bandwidth of a peer reaches the maximum stream bandwidth, it does not require extra incoming degree. This implies that the incoming degree of peers is limited by $D_{max} = \frac{BW_{max}}{bwpf}$, where BW_{max} denotes stream bandwidth. We call a peer is *saturated* when its entitled degree exceeds the maximum required degree, *i.e.*, $R_i > D_{max}$. Second, the entitled incoming degree of each peer (R_i) can only take integer values. In order to avoid over-estimating the amount of allocated resources to each peer, we always use the floor of the resulting value from Eqn. (1). We can revise Eqn. (1) to address these two issues as follows:

$$R_i = \left\lfloor \min \left\{ \left(\frac{1}{t} W_i + \frac{t-1}{t} \sum_{t=1}^t \frac{W_i}{N} \right), D_{max} \right\} \right\rfloor. \quad (2)$$

To effectively utilize the excess resource in the system, the unsaturated peers can further increase their incoming degree. These extra incoming connections are referred to as *excess degree* and denoted with e_i . In summary, the total actual incoming degree of each peer (a_i) consists of two components: $a_i = R_i + e_i \leq D_{max}$. Note that it is difficult to determine the amount of aggregate excess resources in the system due to the random and dynamic nature of excess resources. In subsection III-B, we describe how individual peers determine their excess incoming degree in a distributed fashion.

Once a peer computes its entitled degree (R_i), it intends to identify D_{max} parents in the system. Towards this end, first each peer learns about a subset of participating peers through a bootstrap node. Then, it progressively contacts them to discover their ability to serve as a parent. Each peer first establishes R_i entitled connections and then explores the feasibility of establishing some excess connections as we describe in subsection III-B.

Goals & Assumptions: The contribution-aware mechanism should be able to gracefully cope with the

| Symbol | Definition |
|-----------|---|
| N | total number of peers in the system |
| W_i | the willingness of peer i , measured by degree, i.e., its bandwidth contribution to the overlay divided by bandwidth-per-flow, $bwpf$ |
| a_i | actual number of incoming degree for peer i |
| f_i | actual contribution (outgoing degree) of peer i |
| R_i | computed entitled incoming degree of peer i |
| e_i | actual excess incoming degree of peer i |
| τ | period of update |
| D_{max} | maximum required degree to get full quality live stream |

TABLE I
DEFINITION OF IMPORTANT SYMBOLS

inherent dynamics of peer participation, or churn. To achieve this goal, two issues should be addressed: (i) individual peers should periodically determine their entitled incoming degree, and adapt their overall incoming degree accordingly; (ii) each peer should implement a *preemption policy* to fairly manage the allocation of its outgoing degree among requesting child peers. In essence, the preemption policy ensures that the available resources in the system are proportionally allocated across participating peers. We describe the preemption policy at each peer in subsection III-C.

We make the following assumptions throughout this paper: First, the incoming bandwidth of each peer is larger or equal to streaming bandwidth. This implies that each peer tries to increase its overall incoming degree to its maximum value (i.e., D_{max}). This is a reasonable assumption since the bandwidth of a video stream with an acceptable quality is around 400Kbps to 600Kbps which is less than the incoming access link bandwidth for most of the today's Internet users as indicated in earlier studies [7]. Second, individual peers provide correct information about the number of child peers they can support (W_i), i.e., the amount of resources they are able and willing to contribute to the system. We believe that our proposed approach can be integrated with other existing incentive or enforcement mechanisms. Table I summarizes the notations used throughout this paper.

A. State Collection and Reporting

The state collection and reporting mechanism performs two tasks: (i) collecting the required information from individual peers and determining the group-level information such as N and $\sum W_i$; and (ii) reporting the group level information to all participating peers in the system. We consider a simple centralized approach for both state collection and reporting through a bootstrap node. When a peer joins the system,

it contacts a well-known bootstrap node and provides its willingness to contribute (W_i). During a session, each peer sends a heart-beat message to the bootstrap node once every τ seconds and reports the value of its dynamic properties including its actual outgoing degree (f_i) and incoming degree (a_i) along with its entitled degree (R_i) and the list of its parents. The bootstrap node maintains the following information for each participating peers ($W_i, f_i, a_i, R_i, \text{list of parents}$) and updates this information after receiving each heart-beat message. Each peer also sends a BYE message to the bootstrap node right before its departure. If the bootstrap node does not receive a heart-beat message from a peer for $2*\tau$ seconds, it assumes that the peer has departed and remove its record. In a nutshell, the bootstrap node has an updated state of individual peers and thus can easily determine the group-level state such as N and $\sum W_i$. Note that the state information at the bootstrap node may not be perfectly accurate since the state of each peer is likely to change between consecutive updates.

The bootstrap node reports the most recent group-level state to all participating peers once every τ seconds. When a peer receives a new report from the bootstrap node, it determines the number of its entitled connections (R_i) using Eqn. (2). If the value of R_i is larger than its current incoming degree, it continues the discovery for more parents. In contrast, if its entitled incoming degree has dropped, it increases the value of e_i accordingly. Note that peers do not explicitly disconnect their incoming connections due to the drop of R_i , rather they consider a larger number of existing connections to be excess connections. The preemption policy at parent peers disconnects a proper number of these excess connections based on the overall demand for excess connections among peers. This passive strategy for disconnecting connections reduces dynamics in the system. τ is a configuration parameter that determines the tradeoff between the freshness of state information at the bootstrap node and the signaling overhead. More specifically, increasing the value of τ reduces the signaling overhead associated with state collection and reporting at the cost of lower accuracy for the state information at the bootstrap node. The default value of τ is 10 seconds.

B. Parent Discovery

The goal of the parent discovery mechanism is to enable each peer to locate the required number of parents to establish the desired number of incoming connections. Each peer always establishes R_i entitled connections and then explores possibility for establishing excess connections (if it requires any). Note that each peer does not label its individual incoming connections as “entitle” or “excess” connection. Instead, a child peer only keeps track of its actual number of connections (a_i) and its entitled degree

R_i that is periodically updated after each report from the bootstrap node. This is feasible in mesh-based P2P streaming mechanism because all connections have the same value and thus the total number of connections determines the delivered quality not the identity of those connections¹.

To establish an entitled or excess connection, each peer first obtains the contact information for a subset of participating peers that are likely to be able to accommodate more child peers from the bootstrap node. Since the bootstrap node maintains the state of all participating peers (*i.e.*, potential parents), it can identify potential parents and report a list of random subset of them to a requesting peer. More specifically, the bootstrap node identifies a random subset of participating peers that have at least one empty slot or a child that can be preempted by the requesting peer. In essence, the bootstrap node implicitly coordinates the connections among peers. This in turn increases the probability of success during the parent discovery process. It is worth noting that despite this coordination, it is possible that a parent rejects a request due to a recent change in its status.

Given such a list of potential parents, each peer *sequentially* contacts peers in the list, provides its minimum local state (*i.e.*, W_i , a_i and R_i)² and requests the contacted peer to serve as its parent. A contacted peer determines whether to accept or deny a request to serve as a parent based on the *local preemption policy* as we describe in the following subsection. Once a child peer receives a response from a parent, it updates the number of its entitled and excess connection accordingly and provides its updated information at its next heart-beat to the bootstrap node. Each peer continues to establish connection to more parents until its incoming degree reaches its maximum value (or D_{max}). If the list of potential parents is exhausted, the peer will contact the bootstrap node to obtain a new list. When peer i 's request for connection is rejected by a potential parent, its reaction depends on its current state as follows:

- *Looking for more Entitled Connections* ($a_i < R_i$): In this case, a child peer immediately sends a request to the next potential parent in the provided list by the bootstrap node. This rather aggressive approach to discovery is reasonable because there must be sufficient resources in the system for each peer to reach its entitled incoming degree.
- *Looking for more Excess Connections* ($a_i > R_i$): In this case, a rejected request is an indication of limited availability of excess resources in the system. Therefore, the rejected peer waits for an interval t_{wait} , called *wait interval*, before it contacts another parents to establish a connection. The wait interval

¹In contrast, the contribution aware mechanism for tree-based P2P streaming [4] must specifically label each connection because each connection provides a particular description.

²All other states that a parent might need can be derived from these information.

is exponentially backoff with each rejected request for excess connection as follows [4]:

$$t_{wait} = t_{min} * K * (e_i + \beta^{ret}) \quad (3)$$

where t_{min} is the minimum backoff time, K is a random number larger than 1, β is backoff factor and ret is the number of consecutive failures. t_{min} is set to 5sec and β is 2. This approach for determining wait time adaptively adjusts the frequency of attempt for establishing excess connections by individual peers and thus the aggregate demand for excess connections without any explicit coordination among peers.

Our proposed central approach to state collection, reporting and parent discovery might exhibit a limited scalability. However, besides its simplicity, this central approach enables the bootstrap node to perform passive coordination and improve the efficiency of parent discovery. It is feasible to incorporate a distributed mechanism for state collection, reporting and parent discovery into the mesh-based P2P streaming approach. However, we do not believe that adding such elaborate mechanism significantly change the overall behavior of the system.

C. Local Preemption Policy

The local preemption policy determines how a parent peer reacts to a request for connection from a child peer. If the current number of child peers for a parent peer is less than the degree that it is willing to contribute (W_i), then a request for connection is always accepted. However, if the outgoing degree of a parent peer is fully utilized, then a new child peer A can only replace (or preempt) an existing child peer B if providing a connection to child peer A has a higher priority. The relative priority of connection to peers A and B is determined in different scenarios as follows:

- *En-Ex Policy*: If peer A is looking for entitled connection ($a_A < R_A$) and peer B already has some excess connections ($a_B > R_B$), then a request by A can always preempt an existing connection to peer B . This policy allows a new peer to easily reach its entitled incoming degree by preempting excess connections from other peers.
- *Ex-En Policy*: If peer A is looking for an excess connection ($a_A > R_A$) when peer B only has entitled connections ($a_B \leq R_B$), then a request by A can not preempt an existing connection from peer B .
- *En-En Policy*: if both peers only have entitled connections, then A can only preempts the connection from B if the normalized incoming degree of A is less than B , *i.e.*, the following condition is

| A, B | Entitled | Excess |
|----------|--|------------------------|
| Entitled | Yes if $\frac{r_A}{W_A} < \frac{r_B}{W_B} - 1$ | Yes |
| Excess | No | Yes if $e_A < e_B - 1$ |

TABLE II

LOCAL PREEMPTION POLICIES USED BY EACH PARENT IN DETERMINING IF A NEW PEER A CAN PREEMPT AN EXISTING CHILD PEER B TO USE THAT SLOT AS A CHILD FOR THIS PARENT.

satisfied: $\frac{r_A}{W_A} < \frac{r_B}{W_B} - 1$. This condition basically ensures that all peers proportionally increase their entitled incoming degrees. Note that the equation incorporates a hysteresis effect to prevent oscillating preemption between two peers.

- *Ex-Ex Policy*: if peer A is looking for excess connections ($a_A > R_A$) and peer B has some excess connections ($a_B > R_B$), A can preempt an existing connection to peer B when it has a smaller number of excess connections (*i.e.*, $e_A < e_B - 1$). This condition balances out the number of excess connections among peers. It also incorporates a hysteresis to prevent oscillating preemption between two peers.

Table II summarizes the above preemption policies by a new peer A to an existing child peer B .

Note that when a new peer joins the system or an existing peer loses its parent due to preemption, they start the parent discovery process and could in turn preempt another peer in the system. Therefore, the observed rate of change in parents among participating peers is higher than parent departure rate that occurs only due to churn. In essence, the preemption further aggravates the instability of the overlay.

D. Differences with Tree-based Approach

The contribution aware mechanism that we describe in this section is primarily inspired and shared similar components with the approach proposed by Sung et al. [4] for tree-based P2P streaming. However, there are some notable differences between these two approaches that are worth mentioning.

In the tree-based approach, a particular description of the content is delivered through each tree. Therefore, each peer should join a proper number of trees and also serve as an internal node in only one tree. This approach raises a few issues that do not exist in mesh-based streaming as follows: First, to improve received quality in a tree-based approach, each peer should find a parent in a particular tree whereas in a mesh-based approach any new peer can serve as a parent. Second, the local preemption policy for tree-based approach should distinguish between entitled connections for internal versus leaf peers. This in turn adds new scenarios that should be addressed by the policy whereas our approach does not need to deal with this issue. Third, in the tree-based approach each connection should be specifically

labeled as “entitled” or “excess”. In contrast, in mesh-based approaches the number (rather than identity) of excess connections is simply determined by the difference between the actual number of connections and the number of entitled connections for each peer (*i.e.*, $e_i = a_i - R_i$ when $a_i > R_i$).

Another important difference between our approach and the tree-based approach is the use of peer’s willingness (W_i) instead of its actual contribution (f_i) to determine its entitled incoming degree in Eqn. (1) and (2). Given that the actual contribution of each peer is likely to be less than its willingness (*i.e.*, $f_i < W_i$), using the actual contribution has two side effects: (*i*) the available resources in the system is underestimated in the second term of Eqn. (1) and (2) (*i.e.*, $\sum_{i=1}^N f_i \leq \sum_{i=1}^N W_i$). This in turn leads to a more conservative behavior by individual peers during the parent discovery process, (*ii*) the actual contribution of peer i depends on the ability and demand of other participants to use its outgoing bandwidth. (*iii*) effect of churn (*i.e.*, departure of a child peer) results in transient drop in f_i which leads to more dynamics in the system. Our examinations revealed that this approach will slow down parent discovery and is inappropriate in a dynamic environment where peer population (and thus available resources) is constantly changing.

IV. UNDERSTANDING THE TAX FUNCTION

Before evaluating the proposed contribution-aware mechanism, we examine the behavior of the tax function (*i.e.*, Eqn. (1)) as well as the impact of main parameters on its behavior (*e.g.*, W_i). Understanding the behavior of the tax function reveals how available resources are shared among participating peers across the parameter space in the absence of any dynamics in peer participation. This in turn serves as a reference to examine the performance of the contribution aware mechanism and helps us examine the behavior of our proposed mechanism over a proper portion of the parameter space.

Given a scenario with N peers and their level of willingness to contribute (*i.e.*, outgoing degree W_i), we can define the Resource Index (RI) of a scenario as the ratio of aggregate contributed resource ($\sum W_i$) to the aggregate demand for resources. Since we assume that all peers have sufficient incoming bandwidth to receive full quality stream, the aggregate demand for resources can be simply determined as $N * D_{max}$ and thus RI is $RI = \frac{\sum W_i}{N * D_{max}}$. We can derive the value of $\sum W_i$, and replace it in Eqn. (1) as follows:

$$R_i(t) = \frac{1}{t}W_i + \frac{t-1}{t}RI * D_{max}. \quad (4)$$

Eqn. (4) represents the entitled degree of a peer i as a function of tax rate t based on the following parameters: peer’s willingness (W_i), resource index in the overlay (RI) and maximum incoming degree

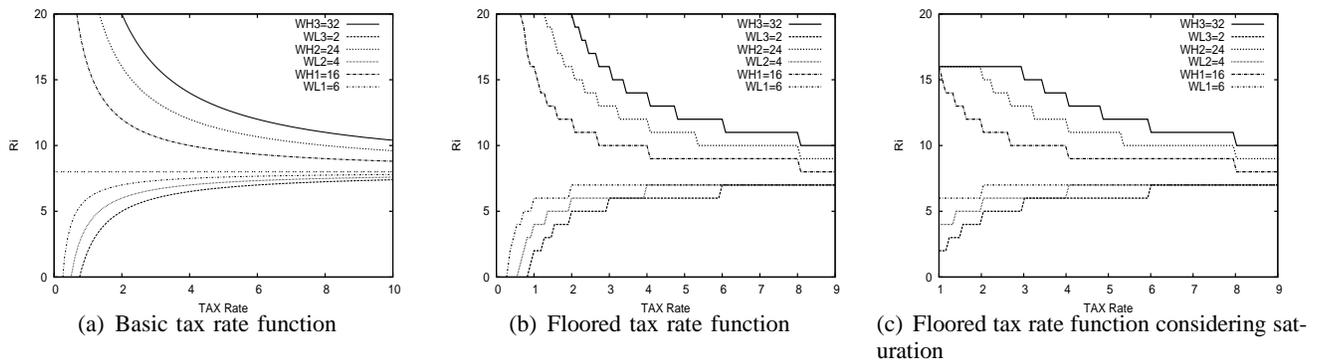


Fig. 1. Behavior of tax function with different values of W_i when RI is 0.5 and D_{max} is 16

(D_{max}).

Figure 1(a) plots $R_i(t)$ as a function of tax rate t for three different combinations of W_i when $RI=0.5$, $RI * D_{max}=8$ ³. For comparison we plot a line for $RI * D_{max}$ in the figure. This figure reveals some important properties of the tax function across the parameter space as follows: First, as the tax rate increases, the entitled degree of high bandwidth peers ($W_i > RI * D_{max}$) is gradually decreasing with tax rate whereas for low bandwidth peers ($W_i < RI * D_{max}$) the entitled degree is gradually increasing. Furthermore, the entitled degree of all peers converges towards the same value of $RI * D_{max}$ regardless of its initial value. To explain this, we note that as t increases the first term in the equation rapidly decreases and the second term converges to $RI * D_{max}$. Second, the larger the value of W_i , the faster the allocation of resources changes with tax rate. Third, the value of $RI * D_{max}$ approaches the value of the entitled degree of all peers when tax rate goes to infinite. Therefore, changing RI or D_{max} simply shifts the converging value in Figure 1(a) up or down accordingly. Fourth, as we have discussed earlier, we always use the floor value of R_i to prevent over-estimating the available resources. Figure 1(b) depicts $\text{floor}(R_i)$ (Eqn. (4)) which results in a step-like evolution of entitled degree as a function of tax rate. Fifth, as we have explained earlier, high bandwidth peers become saturated when their entitled degree is larger than the maximum degree *i.e.*, $D_{max} \leq R_i$. This implies that the actual degree of a saturated peer is limited to D_{max} . Figure 1(c) illustrates the upper limit of incoming degree for the saturated high bandwidth peers which occurs when the tax rate is low. Note that it is important to determine whether (and what fraction of) peers become saturated in a given scenario because this affects the amount of excess resources in the system which in turn determines delivered quality to non-saturated peers. We further elaborate this issue in the evaluation section.

In a nutshell, Figure 1(c) represents the behavior of tax function in a static system where the peer

³While this figure shows the tax function for positive tax rates values, in practice only tax values that are larger than 1, are of interest.

population and the available resources are fixed and known, *i.e.*, the reference static scenario. In practice, because of the dynamics of peer participation and the resulting variations in available resources, the reported group state to individual peers is not perfectly accurate. Therefore, the average behavior among participating peers could be different from the above reference case. We investigate this issue in the next section.

V. PERFORMANCE EVALUATION

As we discussed in Section II, in mesh-based P2P streaming mechanisms (such as PRIME) enforcing the bandwidth-degree ratio implies that all connections have roughly the same bandwidth. Furthermore, the swarming content delivery also implies that all connections have the same value. Therefore, main goals of the contribution aware mechanism are (i) each peer has a proper number of child peers so that its resources are effectively utilized; and (ii) each peer can identify and establish connections with a proper number of parents proportional to its share of available resources. In essence, the performance of a contribution aware mechanism for mesh-based P2P streaming should be assessed based on the ability of individual peers to keep their incoming and outgoing degrees at the proper values. Note that the delivered quality depends on both connectivity of the overlay that is managed by contribution aware mechanism, and the swarming content delivery. Therefore, to evaluate the performance of the proposed contribution aware mechanism we only examine the connectivity among peers and do not consider the content delivery mechanism and the actual delivered quality. Focusing on the overlay connectivity allows us to conduct session level (instead of packet level) simulations. This in turns enables us to simulate scenarios with a large number of peers, or high degree of dynamics in peer participation.

Toward this end, we use our P2P session-level simulator, called *psim*. *psim* is an event-driven simulator that incorporates pairwise network delay between participating peers using the King dataset[8]. Furthermore, *psim* incorporates a realistic model for churn by using a log-normal distribution (with $\mu=4.29$ and $\sigma=1.28$) for peer session time and Pareto distribution (with $a=2.52$ and $b=1.55$) to model the peer inter-arrival time as reported by prior empirical studies on deployed P2P streaming systems [9], [10].

Each simulation is run for 6000 seconds and the information is collected during the steady state when the population reaches the desired target. The reported results for each simulation are averaged across multiple runs with a different random seed. We also use the following default parameters in our simulations: on average 80% of peers are high bandwidth and the rest are low bandwidth, required incoming degree to receive full quality stream is 8, the degree of willingness for high and low bandwidth peers (*i.e.*, their

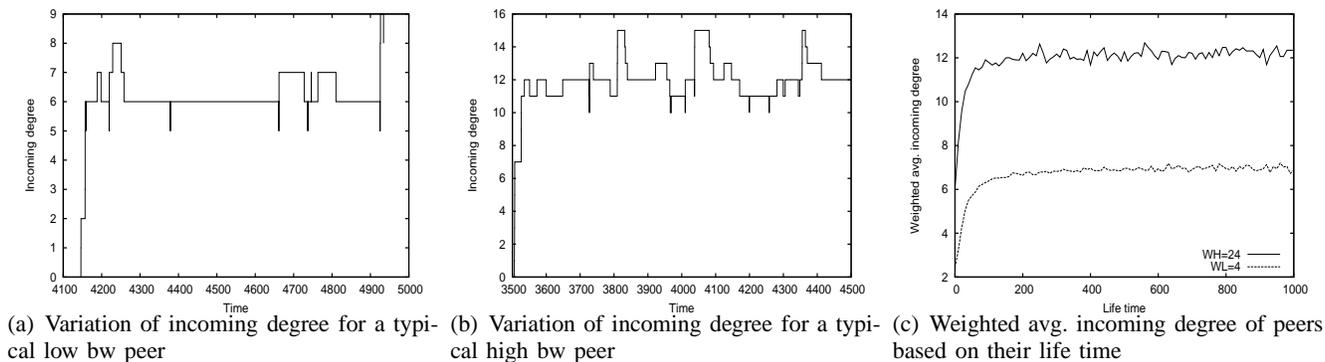


Fig. 2. Typical behavior of a high and low bandwidth peer ($t = 4$, $W_H = 24$ and $W_L = 4$). High bandwidth peers are entitled to degree of 11 while low bandwidth peers are entitled to 6.

outgoing degrees) are 24 and 4, respectively. The resource index is 0.5. The state collection and reporting is performed once every 10 seconds.

A. Dynamics of Parent Selection

We start by examining the dynamics of changes in the number of parents that are caused by the contribution aware mechanism as well as churn. Figure 2(a) and 2(b) show the typical evolution of the incoming for a low and a high bandwidth peers over time when tax rate is 4, respectively. In this scenario, the average entitled degree for high bandwidth peers is 11 and for low bandwidth peers is 6. These figures illustrate that a peer can quickly increase its incoming degree from zero to reach its entitled degree, *i.e.*, less than 20 seconds for a high bandwidth peers and 11 seconds for a low bandwidth peer. These figures also show that once the incoming degree of a peer reaches its entitled degree, its incoming degree oscillates around the entitled value due to the minor changes in available resources and the variations in the number of excess connections. Figure 2(c) presents the average incoming degree among peers whose lifetime is within the range of $[x, x+10]$ seconds. In essence, this figure shows the evolution of average incoming degrees over time and reveals that all peers reach their target incoming degree in around 60 seconds. This also implies that peers with lifetime shorter than 60 seconds, will not remain in the system sufficiently long to reach their target degree.

B. Benefits of Contribution Awareness

To examine the ability of the contribution aware mechanism to manage the incoming degree of participating peers, we present the notion of “weighted average degree”. Weighted average (incoming or outgoing) degree of a peer presents its effective average degree by weighting each degree by the interval

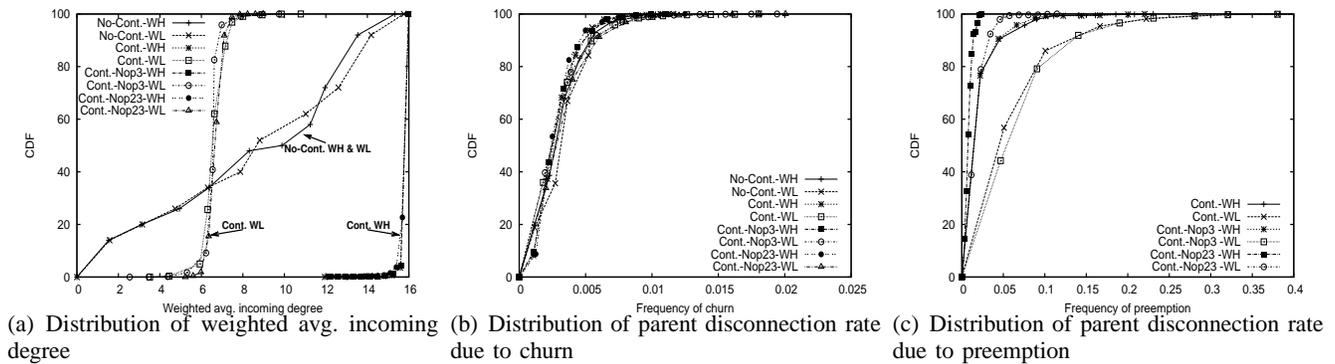


Fig. 3. Results showing the effectiveness of contribution aware mechanism with different combinations of policies- ($t = 2$, $W_H = 24$ and $W_L = 4$)

that a peer maintained at that degree. For example, if a peer has an outgoing degree of 3 for one fourth of its session and 5 for the rest of its session time, its weighted outgoing degree is 4.5. The weighted incoming and outgoing degree of each peer simply quantify the utilization and contribution of the resources during the session, respectively. We further divide the weighted average incoming degree of individual peers into weighted average entitled and excess degrees.

Figure 3(a) depicts the CDF of weighted average incoming degree among high and low bandwidth peers when tax rate is 2, with contribution-aware mechanism (labeled as Cont.*) and without it (labeled as No-Cont.*). This figure clearly shows that in the absence of the contribution-aware mechanism, the distribution of incoming degree is similar for high and low bandwidth peers, but it is rather diverse within each group, *i.e.*, the allocation of resources does not depend on the contribution of participating peers. In contrast, the distribution of incoming degree for high and low bandwidth peers are clearly separated and is very similar within each group. More specifically, all low bandwidth peers (Cont.*WL) have a degree close to 7 whereas the degree of high bandwidth peers (Cont.*WH) is very close to 16. Figure 3(a) illustrates that the contribution aware mechanism can effectively manage the allocation of resources among participating peers.

To quantify the importance of different preemption policies on the performance of the contribution aware mechanism, we present the distribution of weighted average incoming degree for high and low bandwidth peers in two other scenarios where (i) the En-En policy is off (labeled as Cont-nop3-*); and (ii) both En-En and Ex-Ex policies are off (labeled as Cont-nop23-*). Figure 3(a) indicates that eliminating Ex-Ex and En-En preemption policies does not lead to any visible change on the allocations of resources

| Scenario | all changes | Churn | Preempt. |
|-------------------------|-------------|-------|----------|
| Cont. | 1.5% | 29% | 2% |
| Cont. w/o En-En | 3.2% | 29% | 5% |
| Cont. w/o Ex-Ex & En-En | 24% | 29% | 51% |
| No-Cont. | 29% | 29% | 100% |

TABLE III
PERCENTAGE OF STABLE PEERS

among peers. In other words, the En-Ex policy appears to be sufficient to achieve good performance.⁴

Stability of the Overlay: We also quantify the stability of the overlay by measuring the parent disconnection rates for individual peers. We further divide these disconnections into two groups: disconnections that are due to parent departure versus due to preemption by other child peers. Figure 3(b) depicts the distribution of the average parent disconnection rate due to churn among both high and low bandwidth peers in all scenarios that we examined in Figure 3(a). Since the overall parent disconnection rate for each peer due to churn is directly proportional to its incoming degree, we normalize the parent disconnection rate by the incoming degree in Figure 3(b) for fair comparison. As expected, Figure 3(b) illustrates that the normalized parent disconnection rate due to churn does not change with contribution aware mechanism and does not depend on peer bandwidth (*i.e.*, peer degree). Figure 3(c) presents the distribution of the average parent disconnection rate among participating peers for high and low peers only due to preemption in all the scenarios that we examined in Figure 3(a) (except for the scenario without contribution aware since no preemption occurs in that case)⁵. Figure 3(c) shows that low bandwidth peers observe a higher rate of preemption in the base case (Cont.-WL) and even after disabling En-En preemption policy (Cont.-Nop3-WL). However, after disabling Ex-Ex and En-En, parent disconnection rate decreases significantly (Cont.-Nop23-WL). This suggests that the Ex-Ex preemption policy primarily contributes to the parent disconnection rate. Note that in this parameter setting high bandwidth peers' connections are entitled therefore they do not observe major preemption. We further examine stability in other settings in Subsection V-C.

The stability of overlay can be also characterized in a more coarse-grained fashion. Table III presents the percentage of peers whose observed time between consecutive changes in parents (regardless of their cause) is at least 600 seconds. Each row of the table represents different scenario with contribution-aware

⁴It is worth noting that En-En and Ex-Ex policies might affect the allocation of resources when RI significantly changes with time. However, constructing such a scenario requires detail information about potential dynamics of RI over time that has not been provided by previous empirical studies. We plan to further study this issue in our future work.

⁵Note that normalizing the rate of change in parents due to preemption in Figure 3(c) is not meaningful since the observed rate depends on the relative number of excess connections for each peer.

mechanism (including various combination of preemption policies) and without it. The table shows that in the absence of contribution-aware mechanism 29% of peers are stable. The percentage of stable peers with contribution aware mechanism drops to 1.5%. Disabling the En-En policy slightly improves the percentage of stable peers from 1.5% to 3.2%. However, removing the Ex-Ex policy significantly increase the percentage of stable peers to 24% which is close the observed stability without the contribution aware mechanism. Since the En-En and Ex-Ex policies significantly increase the instability of the overlay without affecting the performance of the contribution-aware mechanism, we eliminate these two policies for the remaining evaluations in this paper.

C. Effect of Tax Rate & Peer Contribution

In this section, we examine how the behavior of the contribution-aware mechanism changes with the following two key parameters that determine a particular scenario: (i) the value of tax rate (t), and (ii) the value of peer's willingness to contribute (W_i). We consider the default parameters but with three different level of contribution (*i.e.*, degree of willingness or outgoing bandwidth) for high and low bandwidth peers as follows: (i) Scenario $S1$: $WH=16$, $WL=6$, (ii) Scenario $S2$: $WH=24$, $WL=4$ and (iii) Scenario $S3$: $WH=32$, $WL=2$.

We want to keep the resource index ($RI=0.5$) and the percentage of high and low bandwidth peers (80% and 20%) fixed across these scenarios for proper comparisons. This implies that the heterogeneity of contributed resources by high and low bandwidth peers should proportionally adjusted across these scenarios so that the aggregate contributed resources remains fixed. Therefore, examining the performance of the system across these scenarios reveals how the heterogeneity of contributed resources (or W_i) among peers affect system performance.

Figure 4(b) depicts the weighted average entitled degree among high and low bandwidth peers as a function of tax rate for all three scenarios. Figure 4(a) shows that the entitled degree for high and low bandwidth peers based on Eqn. (2) in all three scenarios as a reference. Comparing these two figures indicates that the weighted average entitled degree among high and low bandwidth peers closely follows its estimated values by equation (2) despite the existing dynamics in the connectivity among peers. Figure 4(c) presents the weighted average of total incoming degree (both entitled and excess) among high and low bandwidth peers in three scenarios. This figure shows that except for very small tax values, the average values of entitled and total degrees are close.

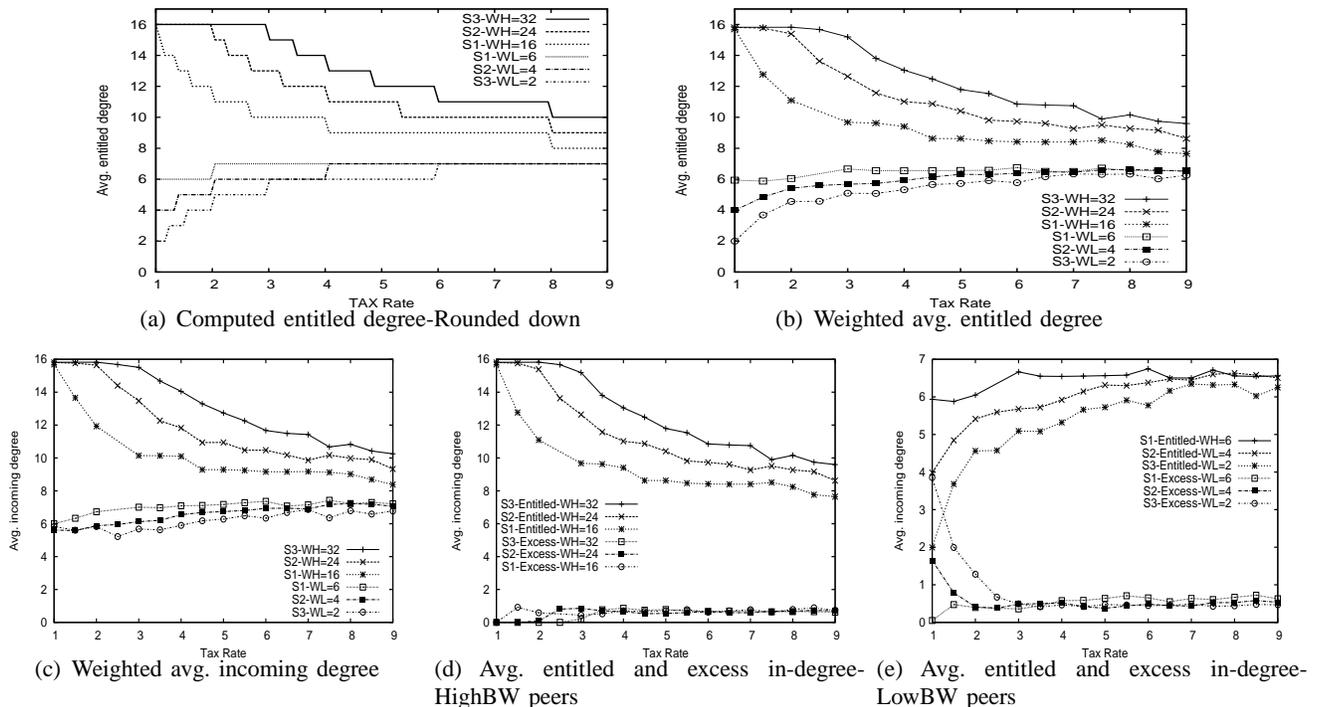


Fig. 4. Effect of tax rate and peer contribution

To further examine the changes in entitled and excess degrees in each group of peers with tax rate, Figures 4(d) depicts the weighted average value of both entitled and excess degree for high bandwidth peers in three scenarios whereas Figure 4(e) presents the same information for low bandwidth peers. These two figures illustrate the following points: First, when tax rate is small, the entitled degree of the high bandwidth peers becomes saturated and thus they do not require excess connections. Since saturated peers do not use their entitled degree, excess resources becomes available in the system, and the amount of excess resources is proportional to $(R_i - D_{max})$, where R_i is the computed entitled degree of a high bandwidth peer i . Low bandwidth peers can utilize these excess resources as excess connections as shown in Figure 4(e). The lower the entitled degree of low bandwidth peers is in these cases, the more available resources exist for excess connections. Thus low bandwidth peers can get larger number of excess connections as illustrated in Figure 4(e). Second, as long as high bandwidth peers are not saturated ($t > 4$), the average excess degree for both high and low bandwidth peers are the same and does not change with the tax rate or the distribution of peer contributions (across scenarios). The only reason for excess resources in these circumstances is the rounding of entitled degree (due to floor function). Since the amount of resulting excess resources does not change with tax rate or distribution of contribution by peers, the number of average excess degree remains fixed. This also shows that the contribution-aware mechanism evenly divide

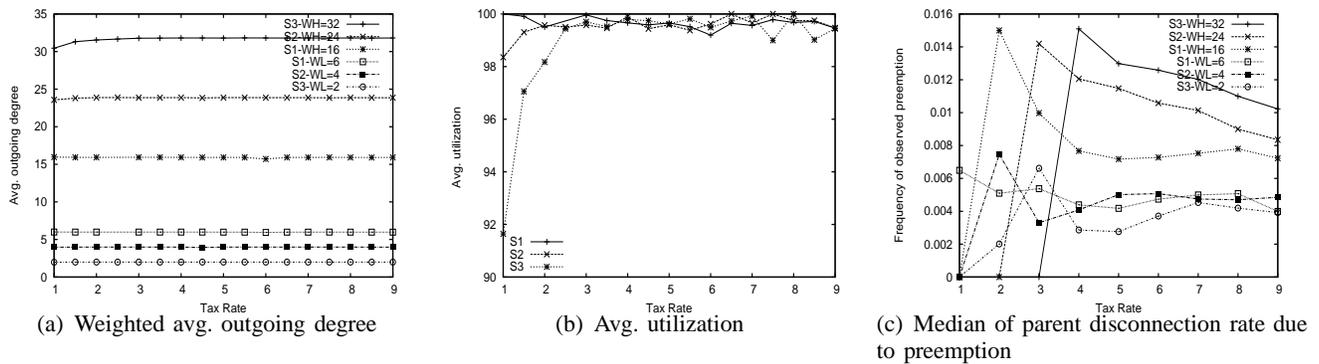


Fig. 5. Effect of tax rate and peer contribution on resource utilization & rate of preemption

excess resources among participating peers.

Utilization of Resources: To investigate the utilization of resources in the system, Figure 5(a) depicts the weighted average outgoing degree among high and low bandwidth peers for three scenarios as a function of tax rate. This figure clearly shows that the outgoing degrees of peers in all scenarios are very close to their willingness to contribute (W_i), *i.e.*, the contribution-aware mechanism can effectively utilize available resources for different distribution of resources among peers despite the dynamics of peer participation. Figure 5(b) presents the overall utilization of outgoing degree among all peers in one snapshot of the overlay. This figure shows that when high bandwidth peers are not saturated, resources are perfectly utilized. In the saturated region, the overall utilization of resources slightly drops due to the dynamics of excess connections. This is the reason for minor drop in the outgoing degree of high bandwidth peers for scenario S_3 in Figure 5(a) when tax rate is small. To explain this, we note that a relatively larger fraction of resources in the system is utilized by excess connections in the saturated region. As the fraction of excess resources and thus excess connections increases, the probability of rejected request for an excess connection grows. This in turn reduces the utilization of resources due to the backoff in adapting the *waitinterval* for retrying a rejected excess connection request.

Stability of Overlay: To quantify the stability of overlay, Figure 5(c) depicts the average parent disconnection rate due to preemption among high and low bandwidth peers across all three scenarios as a function of tax rate. Within the saturated region ($t < 4$), high bandwidth peers do not experience any preemption simply because they only establish entitled connections that can not be preempted. However, outside the saturated region, high bandwidth peers experience a fair parent disconnection rate that gradually drops with increasing tax rate. The observed rate of disconnection by low bandwidth peers is small within the saturated region since there is not much contention for resources and thus no need for preemption. Outside

| Resource Index | BW Distribution | Contribution |
|----------------|-----------------|--------------|
| 0.5 | 12%-88% | 40-4 |
| 0.8 | 23%-77% | 40-4 |
| 0.9 | 29%-71% | 40-4 |
| 1 | 34%-66% | 40-4 |

TABLE IV
PARAMETERS USED IN SIMULATIONS TO EXAMINE THE EFFECT OF RI

the saturated region, the average parent disconnection rate among low bandwidth peers does not change with tax rate across different scenarios. Moreover, while all participating peers have the same average number of excess connections outside of the saturated region, (as shown in Figures 4(d) and 4(e)), Figure 5(c) reveals that high bandwidth peers surprisingly observe a higher rate of disconnection.

The above trends in the stability of parent primarily depends on the average peer degree. More specifically, the larger the total peer degree, the higher the parent disconnection rate. To explain this issue, recall that the type of individual connections (*i.e.*, entitled vs excess) is not explicitly specified by the contribution-aware mechanism in the mesh-based P2P streaming, as we discussed in subsection III-B. Since each parent peer only uses the number of excess and entitled connections for its current children (based on their last update) in order to make preemption decisions, it is likely that two parents leverage their last update from their common child and simultaneously preempt (*i.e.*, disconnect) their connections to this child. The probability of such an event is proportional with the incoming degree of a child peer. Therefore, outside the saturated region, the change in stability as a function of tax rate is similar to the change in degree as shown in Figure 4(c).

D. Effect of Resource Index

We examine the effect of resource availability (or *RI*) on the performance of contribution aware mechanism. Toward this end, we keep the same level of heterogeneity for contributed resources where high and low bandwidth peers are willing to contribute 40 and 4 outgoing connections. However, we change the value of resource index by changing the percentage of high and low bandwidth peers as shown in Table IV. Different scenarios in Table IV are derived from reported traces by earlier empirical studies [4].

Figure 6(b) depicts the weighted average entitled degree of high and low bandwidth peers as a function of tax rate for different scenarios. Figure 6(a) shows the entitled degree of high and low bandwidth peers in the same scenarios based on equation (2) as a reference. Comparing these two figures reveals that the weighted average entitled degree of all peers generally follows their corresponding value derived from the equation. These Figures 6(a) and 6(b) clearly illustrate that as more resources become available (*i.e.*,

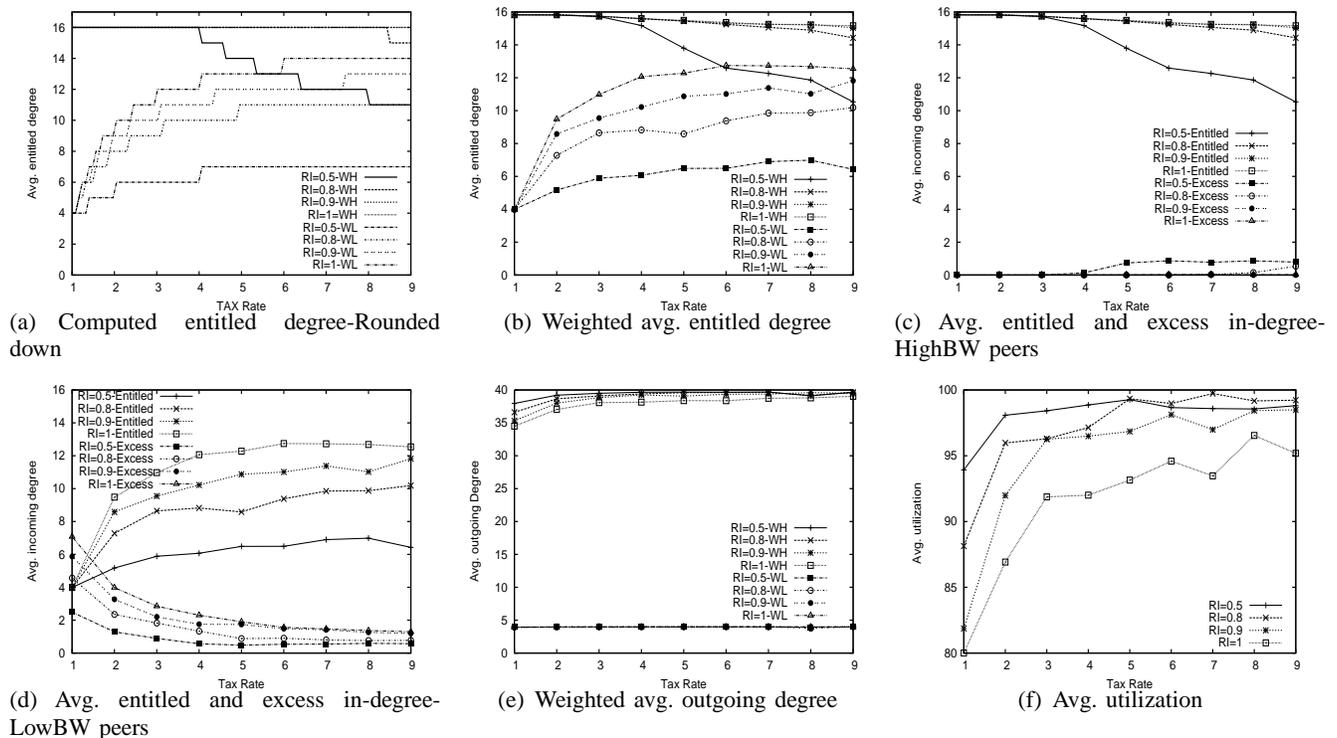


Fig. 6. Results to show the impact of resource index on effectiveness of the contribution-aware mechanism

RI increases), high bandwidth peers remain saturated for a wider range of tax rates, *i.e.*, the size of the saturated region grows. The availability of extra resources enables low bandwidth peers to establish more excess connection and changes dynamics of the overlay.

To examine the effect of RI on each group of peers, we plot the average entitled and excess degrees for high and low bandwidth peers in Figure 6(c) and 6(d), respectively. Figure 6(c) clearly illustrates the saturated region for high bandwidth peers in different scenarios where they do not have any excess connection. On the other hand, Figure 6(d) reveals that low bandwidth peers manage to utilize the excess resources by establishing a larger number of excess connections within the saturated region for each scenario.

Figure 6(e) shows the average out-degree of high and low bandwidth peers as a function of tax rate in scenarios with different RI . The figure clearly shows that across different tax rate and RI values, the average outgoing degree of high and low bandwidth peers is close to their maximum contribution. Figure 6(f) presents the utilization of resources in a single snapshot of the system. This figure indicates that the overall utilization of resources is lower within the saturated region. The lower utilization of resources for both high and low bandwidth peer over small tax rate is due to the larger fraction of excess connections in these settings that results in a larger number of failed attempts to establish connection to a parent.

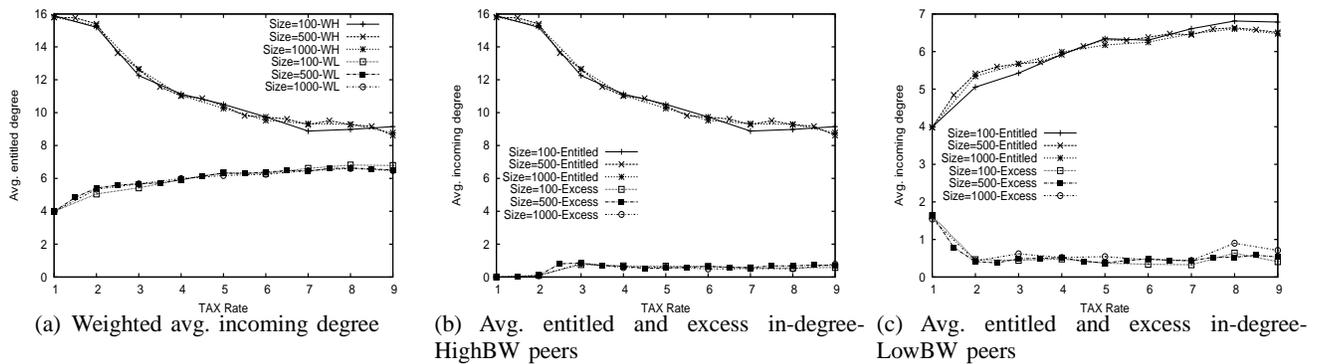


Fig. 7. Effect of group size

This in turn leads to an exponentially increasing wait time which reduces resource utilization. We note that while exponential increase of *waitinterval* adjusts the aggregate demand for excess connection with the availability of resources, there is still a possibility of improper parent selection due to imperfect information on the location of available resources which leads to improper usage of resources. We have observed this effect in the subsection V-C over small tax rate as well.

E. Effect of Group Size

We now investigate how well the contribution aware mechanism scales with the average number of concurrent peers in a session⁶. Toward this end, we change the average population from 100 to 1000 peers where $RI = 0.5$ and high and low bandwidth peers are willing to contribute up to 24 and 4 connections, respectively.

Figure 7(a) depicts the weighted average in-degree of high and low bandwidth peers as a function of tax rate for three different group sizes. Figures 7(b) and 7(c) show the average entitled and excess degrees of high and low bandwidth peers for different group sizes, respectively. These figures collectively illustrate that the average entitled and excess degree of low and high bandwidth peers are very close for different group sizes. This suggests that the contribution aware mechanism is likely to scale with the number of participating peers.

F. Effect of Update Frequency

In this subsection, we explore the effect of reporting interval on the performance of the contribution aware mechanism in a scenario where $RI=0.5$ and high and low bandwidth peers are willing to contribute

⁶Note that the total population changes with churn but psim can set the arrival rate in order to keep the average population at a desired number.

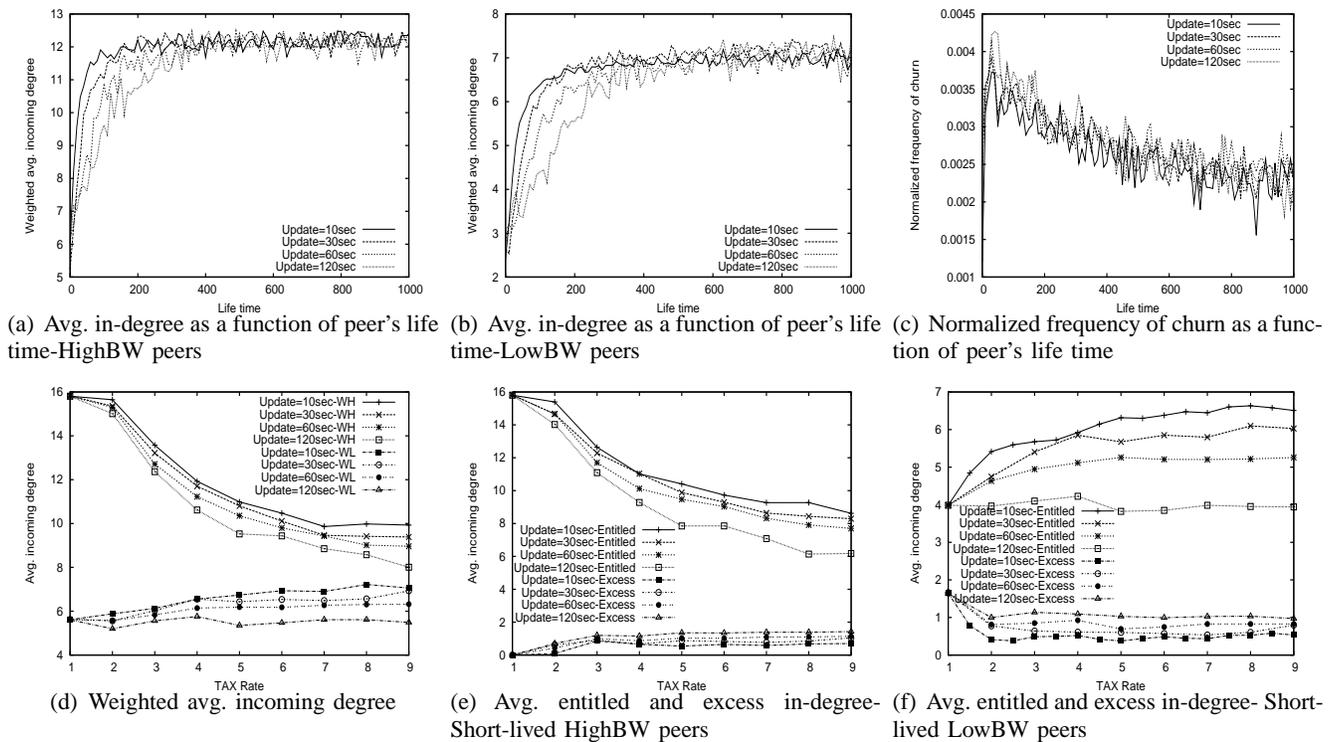


Fig. 8. Impact of update frequency on effectiveness of the contribution-aware mechanism

up to 24 and 4 connections, respectively. In general as the update interval increases, the reported group state to individual peers and thus their estimate of available resources becomes obsolete. Underestimating the available resources will lead to a lower utilization of resources whereas overestimating could result in an imbalance allocation of resources in the absence of En-En preemption policy.

We first study the effect of update interval during the startup phase for individual peers when peers try to reach their target degree after arrival. Figure 8(a) and 8(b) depicts the average incoming degree among high and low bandwidth peers with lifetime between $[x, x+10]$ seconds for different update intervals, respectively. The tax rate in these figures is 4 and the results for other tax rates exhibit similar behavior. We truncated the x-axis at 1000 seconds since the behavior remains the same for higher life time values. These figures clearly illustrate that increasing uptime primarily affects short-lived peers (with life time less than 400 seconds) that have not reached their target degrees. As the update interval increases, the effect is similar for both high and low bandwidth peers, and results in a lower incoming degree. To explain this result, we note that in our target scenarios, the group population and thus RI has a relatively small fluctuation due to churn⁷. Since the amount of aggregate resources is relatively stable, once long-lived

⁷One can generate artificial group dynamics that leads to significant and rapid changes in RI . However, such dynamics appear to be unrealistic since it is inconsistent with the reported peer arrival and peer session times in previous empirical studies.

peers establish their connections, the only change in their parents is due to churn. Therefore, increasing update interval does not have a major effect on them. However, short-lived peers are still building up their connections and are very sensitive to inaccurate information. Specifically, if a peer can not successfully identify all its entitled parents, it needs to wait until group state is updated at the bootstrap node to provide a proper list of parents. Inaccurate information could also affect ability of long lived peers to replace a departed parent. To quantify the frequency of such events, Figure 8(c) depicts the average value of normalized frequency of churn among peers whose lifetime is between $[x, x+10]$ seconds. This figure indicates that as peer's life increases, it observes the lower rate of churn among parents as well. This is simply due to the fact that a connection between long-lived parent-child remains intact as long as aggregate resources do not change.

Figures 8(e) shows the average entitled and excess incoming degree for high bandwidth peers that are short-lived (lifetime less than 400 seconds) for different update intervals. Figure 8(f) depicts the same information for short-lived, low bandwidth peers. These figures illustrate a couple of points: (i) the overall trend of change in average degree with tax rate is similar for all update intervals; (ii) increasing the update interval results in a major drop in entitled degree and a minor increase in excess degree. These changes in the entitled and excess degrees are larger for higher tax rate. These trends can be explained as follows: as the update interval increases, it becomes harder for short-lived peers to quickly identify the desired number of parents due to the higher inaccuracy in the available group state at the bootstrap node. This leads to a lower utilization of resources and allows the excess connections to dynamically utilize a small fraction of this unused resource. Obsolete information affects only the second term in Eqn. (2) (*i.e.*, $\sum_{i=1}^N f_i \leq \sum_{i=1}^N W_i$), by increasing tax rate, this term plays a more important role than the first term. This results in a larger difference in incoming degrees when update interval increases.

VI. CONCLUSION

This paper presented a contribution aware mechanism for mesh-based P2P streaming based on the notion of tax function. We examined the behavior of a commonly used tax function and describe how it can be incorporated into mesh-based P2P streaming mechanisms to ensure proper allocation of resources among well behaved peers. We conducted extensive session level simulations to illustrate the ability of the proposed mechanism in proper allocation and high utilization of resources over a wide range of scenarios.

We plan to pursue this work along the following directions. First, we would extend the notion of contribution awareness to a group of non-cooperative peers by enabling individual peers to securely

report their own contribution to the system and reliably verify the contribution by other peers. Second, we plan to incorporate a pairwise incentive mechanism (similar to BitTorrent) between connected peers in a bi-directional overlay as an alternative approach to incorporate contribution awareness into P2P streaming mechanism.

REFERENCES

- [1] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "Coolstreaming: A data-driven overlay network for live media streaming," in *IEEE INFOCOM*, 2005.
- [2] V. N. Padmanabhan, H. J. Wang, and P. A. Chou, "Resilient peer-to-peer streaming," in *IEEE ICNP*, 2003.
- [3] N. Magharei and R. Rjaie, "PRIME: Peer-to-Peer Receiver-driven MESH-based Streaming," in *IEEE INFOCOM*, 2006.
- [4] Y. Sung, M. Bishop, and S. Rao, "Enabling Contribution Awareness in an Overlay Broadcasting System," in *ACM SIGCOMM*, 2006.
- [5] Y. Chu, J. Chuang, and H. Zhang, "A case for taxation in peer-to-peer streaming broadcast," in *PINS, ACM SIGCOMM workshop on Practice and theory of incentives in networked systems*, 2004.
- [6] B. Cohen, "Bittorrent." [Online]. Available: <http://www.bittorrent.com>
- [7] Y. Chu, A. Ganjam, T. Ns, S. Rao, K. Sripanidkulchai, J. Zhan, and H. Zhang, "Early Experience with an Internet Broadcast System Based on Overlay Multicast," in *USENIX*, 2004.
- [8] T. M. Gil, F. Kaashoek, J. Li, R. Morris, and J. Stribling, "King dataset," 2004. [Online]. Available: <http://pdos.csail.mit.edu/p2psim/kingdata/>
- [9] E. Veloso, V. Almeida, W. Meira, A. Bestavros, and S. Jin, "A hierarchical characterization of a live streaming media workload," in *Internet Measurement Workshop*, 2002.
- [10] K. Sripanidkulchai, A. Ganjam, and B. Maggs, "The Feasibility of Supporting Large-Scale Live Streaming Applications with Dynamic Application End-Points," in *ACM SIGCOMM*, 2004.