Peer-to-Peer Netw. Appl.
DOI 10.1007/s12083-010-0078-y

Incorporating contribution-awareness into mesh-based Peer-to-Peer streaming systems

Nazanin Magharei · Reza Rejaie · Yang Guo

Received: 10 December 2009 / Accepted: 3 June 2010

Abstract While Peer-to-Peer streaming has become increasingly popular over the Internet during recent years, the proper allocation of available resources among peers in a resource constraint environment, remains a challenging problem. In a resource constraint environment, the allocated resources 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 which is essential for scalability of P2P systems. However, contribution-aware resource allocation is challenging due to the distributed and dynamic nature of resources in P2P systems. In this paper, we present a tax-based contribution-aware scheme for live mesh-based P2P streaming approaches. In our proposed scheme, 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 scheme can leverage the tax function. Through extensive simulations we demonstrate the ability of our proposed scheme to properly allocate available resources among participating peers over a wide range of scenarios. We show that the amount of resources (i.e., bandwidth) is divided across peers proportional to their contribution and in our default simulation setting the median delivered quality to high bandwidth peers with high contribution is improved by 100%. We believe that our results shed an insightful light on the dynamics of resource utilization and allocation in the context of live mesh-based P2P streaming.

Keywords Peer-to-Peer · Streaming · Mesh-based · Resource management

1 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 (e.g., [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* approaches. 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 substream or a description) to its child peers (e.g., [2–5]). In the mesh-based approach, participating peers often maintain a randomly connected mesh and incorporate

N. Magharei (✉) · R. Rejaie
Department of Computer and Information Science,
University of Oregon, Eugene, OR 97403-1202, USA
e-mail: nazanin@cs.uoregon.edu

R. Rejaie
e-mail: reza@cs.uoregon.edu

Y. Guo
Research and Innovation, Technicolor,
2 Independence Way, Princeton, NJ 08540, USA
e-mail: Yang.Guo@technicolor.com

swarming content delivery (e.g., [1, 6–10]). Mesh-based P2P streaming approach enables participating peers to contribute their resources (i.e., outgoing bandwidth) more effectively which in turn improves the utilization of available resources among peers and leads to a better scaling property for mesh-based approach compared to the traditional tree-based approach [11].

The feasibility of P2P streaming primarily depends on the scalability of contributed (or available) resources, namely outgoing bandwidth, with the number of participating peers. In practice, the resources are often insufficient to maximize the delivered quality to individual peers [5]. Such a resource constraint scenario occurs when a subset of participating peers are unwilling or unable to contribute as much resources as they demand, and this deficit can not be compensated by the excess resources from other peers. Since allocated resources to individual peers determine their maximum delivered quality, in such a "resource constraint" setting, a fair scheme should allocate resources to individual peers proportional to their contributed resources (or their outgoing bandwidth) rather than their demand (or incoming bandwidth). Moreover, excess resources of high bandwidth peers should be divided to all peers proportional to their contribution. Allocating resources in a *contribution-aware* fashion provides incentives among participating peers to actively contribute their resources in order to improve their own observed quality. Incorporating contribution-awareness into P2P streaming is in essence, a distributed resource management problem which is challenging due to the distributed, heterogeneous and dynamic nature of available resources as peers join and leave the system. Because of these challenges, a majority of studies on P2P streaming have not directly address the issue of contribution-aware resource allocation.

In this paper, we propose a tax-based [12, 13] contribution-aware scheme for live mesh-based P2P streaming approach. Such a contribution-aware scheme is a promising approach to effectively manage distributed and dynamic resources in P2P streaming by controlling the connectivity (i.e., number of parents) of individual peers. The effectiveness of incorporating tax-based contribution-aware scheme in the context of tree-based approach has been investigated by Sung et al. [12]. We focus on the mesh-based approach due to its superior performance and better scaling properties compared to the tree-based approach [11]. Towards this end, we make three main contributions as follows: First, we describe how a tax-based contribution-aware scheme can be seamlessly incorporated into the mesh-based P2P streaming. Second, we examine the behavior of the proposed tax function in allocating the available resources among participating peers for different values of aggregate resources and tax rates. We identify the so-called "saturated region" where high bandwidth peers do not require their allocated share of resources, and examine the ability of the proposed scheme to effectively utilize these excess resources. Third, we perform extensive evaluations of the proposed contribution-aware scheme.

Towards this end, we show that the connectivity of individual peers directly determines their observed performance in the mesh-based P2P streaming. Then, to incorporate the contribution aware scheme in the context of mesh-based P2P streaming approach, each peer uses a given tax function to determine the number of parents it is "entitled" to (or its share of resources, i.e., its incoming degree) based on the aggregate available resources in the system, and the amount of its contributed resources. To effectively utilize the excess resources in the system, the unsaturated peers can further increase their incoming degree by adaptively examining the possibility of increasing their number of parents. Each peer as a parent, incorporates a preemption policy to properly allocate its resources between existing and new child peers.

We investigate the effect of key design parameters (e.g., tax rate and preemption policies) over a wide range of scenarios using a realistic model for peer dynamics and pairwise delay. In particular, we examine how changes in aggregate available resources, distribution of contributed resources among peers, and group size affect the allocation and overall utilization of resources, as well as the stability of the overlay. Our main findings are summarized as follows:

(1) The behavior of the proposed contribution-aware scheme for mesh-based P2P streaming closely follows the theoretical model for allocated resources across different tax rates and resource indices.
(2) The performance of high bandwidth peers can be maximized when the value of tax rate is small. In fact, in our default simulation settings, the 10th, 50th and 90th percentile of delivered quality to high bandwidth peers are improved by 800%, 100% and 12%, respectively.
(3) Increasing the tax rate has an opposite effect on the weighted average entitled and excess incoming degree for the high and low bandwidth peers.
(4) Comparing the effect of various preemption policies, reveals that some of the policies significantly increases the instability of the overlay. In our default simulation setting, the percentage of stable peers in a non-contribution-aware setting is 29% which can be dropped to just 1.5% by some of the

preemption policies. In our proposed scheme, the percentage of stable peers reaches to 24% which is comparable to the observed stability without the contribution-aware scheme.

(5) Increasing the intervals of reporting the group state, primarily affects short-lived peers in the overlay as they are more sensitive to the obsolete group state information. Moreover, increasing the reporting interval decreases the utilization of resources.

The rest of this paper is organized as follows: Section 2 provides an overview of mesh-based P2P streaming. In Section 3, we describe the proposed contribution-aware scheme. Section 4 examines the behavior of the tax function. Our performance evaluation methodology along with our evaluation results and main findings are presented in Section 5. In Section 6, we discuss the related work and finally, conclude the paper and sketch our future plans in Section 7.

2 Mesh-based P2P streaming: background

To provide the required background, we present an overview of PRIME [14] 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 overlay, namely a *mesh* over which they incorporate swarming content delivery. Connections between peers could be either directed where there is a parent-child relation between peers or undirected where they are neighbors. To consider a general case, we assume that peers form a directed overlay where there is a parent-child relationship between connected peers.[1] In swarming content delivery, each peer progressively reports its available content to its child peers. The packet scheduling scheme at each peer periodically determines a subset of packets that should be requested (i.e., pulled) from each parent. In this context, with a well-designed packet scheduling and enough buffer any parent can provide useful content to its child peers.

To accommodate the bandwidth heterogeneity among peers and adapt the delivered quality accordingly, the video content can be encoded using either layered coding (LC) or multiple description coding (MDC). LC encodes the streaming content into one base and multiple enhancement layers. The base layer can be independently decoded, while the enhancement layers can only be decoded cumulatively. Sequential layers can be added/dropped to adapt the streaming quality [15–17]. LC is sensitive to data losses of lower layers (i.e., without a base layer the video content cannot be decoded). On the other hand, MDC organizes the streaming content into several sub-streams where each sub-stream can be independently decoded. The delivered quality is proportional to the number of received sub-streams. The choice of various coding schemes clearly impacts the system design where MDC allows for more design flexibility, and LC is better at coding efficiency. We believe that flexibility of MDC outweighs the efficiency of LC because it allows utilization of outgoing bandwidth of any parent regardless of which layer it has. Thus, in this paper, we focus on MDC that allows each peer to receive the appropriate number of sub-streams that are delivered through its access link bandwidth and facilitates the incorporation of contribution awareness into PRIME.[2]

To form an overlay, individual peers learn about a random subset of participating peers by leveraging a central (e.g., bootstrapping) or distributed (e.g., gossip) peer discovery mechanism. Each peer tries to maintain a certain number of parent peers from which it pulls its required content while limiting the number of child peers. 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 peers, participating peers try to maintain their incoming and outgoing degrees proportional to their incoming (bw_{down}) and outgoing (bw_{up}) bandwidth [14]. 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 $bwpf$. $bwpf$ 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 $\lfloor \frac{bw_{down}}{bwpf} \rfloor$ and $\lfloor \frac{bw_{up}}{bwpf} \rfloor$, respectively.

A well-designed packet scheduling scheme delivers the content to all peers with a short delay. Towards this end, we adopt the best performing packet scheduling scheme that achieves a good performance even in resource-constrained and highly dynamic environments [6, 19]. Such a packet scheduling scheme

---

[1] An undirected overlay is a special case of directed overlay, and thus most of our discussion and findings are still valid.

[2] Note that, many previous studies have used this intuition to use MDC coding in their approach [12, 13, 18].

periodically requests the latest available packets with highest timestamps among parents. In addition, it requests a random subset of other missing packets from parents to fully utilize the incoming bandwidth. Given the average outgoing degree of peers ($OutDeg$), outgoing degree of source (Deg_{src}) and the total peer population of N, the minimum required intervals for delivery of all packets to individual peers (buffer size) through this well-designed packet scheduling scheme has been studied in great depth and shown to be ($\log_{OutDeg}(N/Deg_{src}) + 3$) intervals [1, 6].

Using peer connectivity to estimate allocated resources Since all connections have roughly the same bandwidth, the amount of resources (i.e., bandwidth) that a peer contributes or consumes in the overlay can be approximated by its outgoing and incoming degree, respectively. More specifically, when the appropriate packet scheduling and adequate buffer size at each peer are used, the delivered stream quality to each peer would be proportional to its incoming bandwidth [20]. To demonstrate this point, we conduct ns simulations where peers with heterogeneous and asymmetric access link bandwidth form a directed and randomly connected mesh. The MDC encoded streaming content has ten descriptions and all descriptions have the same rate of 160 Kbps. The incoming access link bandwidth of all peers is set to 1.6 Mbps so that each peer can receive the full quality stream. The outgoing access link bandwidth of 20% of peers (high contributors) is set to 2.4 Mbps while for the rest of the peers (low contributors) is set to 400 Kbps. In this setting, the ratio of demand to supply for resource (or resource index) is 0.5. We evaluate the effect of $bwpf$ by setting the maximum incoming degree of each peer to be 12, 16 and 24 which results in $bwpf$ of 66, 100 and 134 Kbps, respectively. To satisfy the bandwidth-degree ratio for the above settings, the outgoing degree of high and low contributors are 18 and 3 ($bwpf$ of 66 Kbps), 24 and 4 ($bwpf$ of 100 Kbps) and finally 36 and 6 ($bwpf$ of 134 Kbps), respectively. The peer population is 200 and source bandwidth is equal to the stream bandwidth as 1.6 Mbps. The buffer size is computed as discussed above ($\log_{OutDeg}(N/Deg_{src}) + 3$) and set to 6 intervals (each interval is 6 sec in this set of simulations). The simulation is run for 6,000 sec and we model peer inter-arrival and session time based on prior empirical studies on deployed P2P streaming systems [18, 21]. Figure 1 shows the 10th, 50th and 90th percentile of average delivered quality among participating peers as a function of their incoming degree. Each line represents a different $bwpf$. This figure clearly demonstrates that for a fixed $bwpf$ the delivered quality to individual peers is directly proportional to their incoming degree, regardless of the parameter choice of $bwpf$. Moreover, by increasing $bwpf$ with a smaller incoming degree, peers can receive higher quality. This result implies that the packet-level dynamics on content delivery has minimal impact on the delivered quality when a well-designed packet scheduling algorithm and sufficient buffer size are deployed at each peer. Therefore, for a fixed $bwpf$ the incoming degree of each peer is a good estimator of its observed quality (or its allocated share of resources).[3]

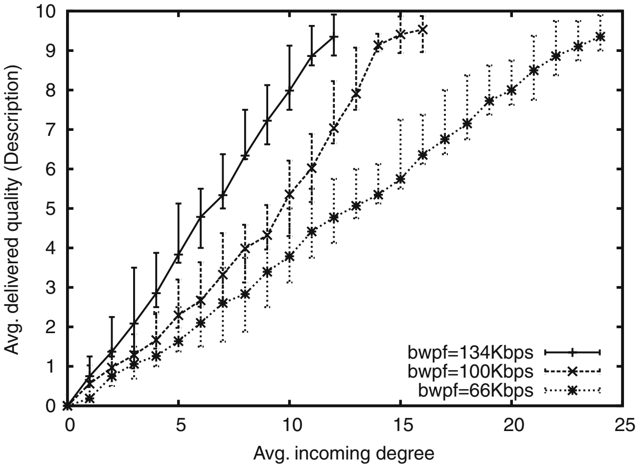

Fig. 1 Relationship between incoming degree and average delivered quality of individual peers

3 Contribution-aware P2P streaming

The primary goal of a contribution-aware scheme is to enable individual peers to determine their share of available resources (i.e., 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 2, the contribution-aware scheme in the context of live P2P mesh-based streaming 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 (1) its outgoing degree (i.e., the number of child peers), and (2) the aggregate outgoing degree across all peers.

---

[3] Note that, the choice of $bwpf$ is limited by the peer with minimum bandwidth contribution and it can be set to a fraction of it.

To support the contribution awareness, each peer i uses a generic cost function [13] to determine its incoming degree R_i:

$$R_i = \frac{1}{t}W_i + \frac{t-1}{t}\sum_{j=1}^{N}\frac{W_j}{N}, \tag{1}$$

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, respectively. 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 (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). The tax rate is always greater than or equal to one ($t \geq 1$) to balance supply and demand for 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 ($\sum W_j$).

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 use (1) to determine their entitled incoming degree. In Section 3.1, 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 as *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 (1). We can revise (1) to address these two issues as follows:

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

To effectively utilize the excess resources in the system, the unsaturated peers can further increase their incoming degree. These extra incoming connections are referred to as *excess degree* and denoted as 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 Section 3.2, 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 Section 3.2. The contribution-aware scheme should be able to gracefully cope with the inherent dynamics of peer participation, or churn. To achieve this goal, two issues should be addressed: (1) individual peers should periodically determine their entitled incoming degree, and adapt their overall incoming degree accordingly; (2) 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 Section 3.3. Table 1 summarizes the notations used throughout this paper.

Goals & assumptions 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 400 to 600 Kbps which is less than the incoming access link bandwidth for most of the today's Internet users as indicated in earlier studies [22]. Second, we assume that peers are well-behaved and provide correct

Table 1 Notations used throughout the paper

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

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. This simplifying assumption allows us to focus on the dynamics of the contribution-aware scheme rather than the security side-effects of uncooperative peers.[4]

3.1 State collection and reporting

The state collection and reporting mechanism performs two tasks: (1) collecting the required information from individual peers and determining the group-level information such as N and $\sum W_i$; and (2) 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, *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 $2*\tau$ sec, 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 sec. When a peer receives a new report from the bootstrap node, it determines the number of its entitled connections (R_i) using (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 sec.

3.2 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 an “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 and not the identity of those connections.[5]

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.

---

[4] Assuming cooperative users is not unrealistic since one can use incentive schemes [23–25] to ensure contribution of resources or deploy a P2P streaming system in a closed setting (e.g., within setup boxes) to achieve the same goal.

[5] In contrast, the contribution aware scheme for tree-based P2P streaming [12] must specifically label each connection because each connection provides a particular description.

Given such a list of potential parents, each peer *sequentially* contacts peers in the list, provides its local state (i.e., W_i, a_i and R_i)[6] and requests the contacted peer to serve as its parent. A contacted peer determines whether to accept or deny a parent request 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 connections accordingly and provides its updated information at its next heart-beat to the bootstrap node. Each peer continues to establish connections 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 a 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 is exponentially backoff with each rejected request for excess connections as follows [12]:

  $$t_{\text{wait}} = t_{\min} * K * \left(e_i + \beta^{ret}\right) \tag{3}$$

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

We note that state collection, reporting, and parent discovery can be performed in a distributed fashion (e.g., [12]). For example, similar to the multiple-tree-based P2P streaming approach, a peer can traverse the mesh in a systematic fashion (starting from the source) and examine each peer to find a proper number of parents with desired type. In the similar fashion, peers' information can be collected and then propagated through the overlay. While this distributed approach does not require a central coordination point which might affect the scalability of the scheme, it can introduce a heavy signaling overhead to those participating peers that are located closer to source and add new dynamics (or introduce other side effects) that can affect the overall behavior of the contribution-aware scheme. We believe that a simple central approach, enables the bootstrap node to perform passive coordination and improve the efficiency of parent discovery. Moreover, it can properly represent a contribution aware scheme in mesh-based P2P streaming and can be used in practice as well.

3.3 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 the child peer A has a higher priority. The relative priority of a connection to peers A and B is determined in different scenarios as follows:

- *En-Ex policy* If peer A is looking for an 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 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 the oscillating preemption between two peers.

---

[6] All other states that a parent might need can be derived from these information.

Table 2 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

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

- *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 2 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.

4 Understanding the tax function

Before evaluating the proposed contribution-aware scheme, we examine the behavior of the tax function (i.e., (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 scheme and helps us examine the behavior of our proposed scheme 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 resources ($\sum W_i$) to the aggregate demands 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 $\frac{\sum W_i}{N * D_{\max}}$. We can derive the value of $\sum W_i$, and replace it in (1) as follows:

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

Equation 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 ($D_{\max}$).

Figure 2a plots $R_i(t)$ as a function of tax rate t for three different combinations of W_i when $RI = 0.5$, and $RI * D_{\max} = 8$.[7] 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 infinity. Therefore, changing RI or $D_{\max}$ simply shifts the converging value in Fig. 2a up or down accordingly. Fourth, as we have discussed earlier, we always use the floor value of R_i to prevent overestimating the available resources. Figure 2b depicts $floor(R_i)$ (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 2c 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, Fig. 2c represents the behavior of the tax function in a static system where the peer

---

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

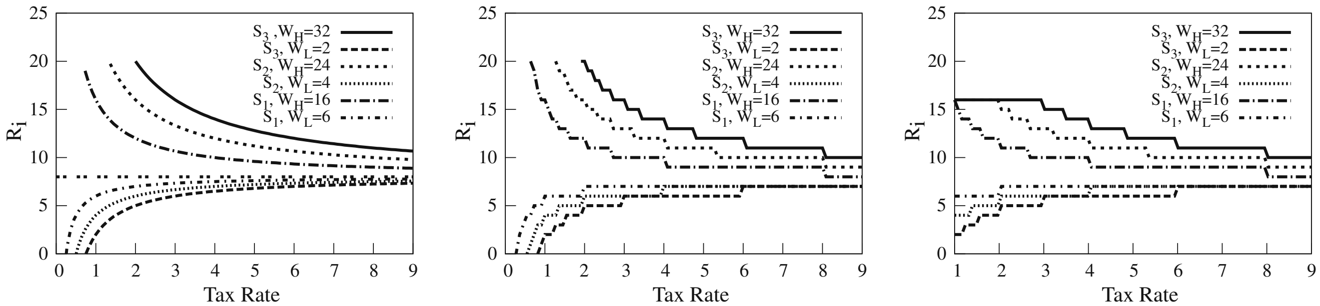

(a) Basic tax rate function

(b) Floored tax rate function

(c) Floored tax rate function considering saturation

Fig. 2 Behavior of tax function with different values of W_i when RI is 0.5 and $D_{\max}$ is 16

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.

5 Performance evaluation

As we discussed in Section 2, 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 scheme are (1) each peer has a proper number of child peers so that its resources are effectively utilized; and (2) 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 scheme 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 scheme, and the swarming content delivery. However, as we have shown in Section 2, there is a direct relationship between peer incoming degree and quality. Therefore, to evaluate the performance of the proposed contribution aware scheme we only examine the connectivity among peers and do not consider the content delivery mechanism and the actual delivered quality.

Toward this end, we use our P2P session-level simulator, called *psim*.[8] *psim* is an event-driven simulator that incorporates pairwise network delay between participating peers using the King dataset [26]. 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 interarrival time as reported by prior empirical studies on deployed P2P streaming systems [18, 21].

In our evaluations, we examine the impact of each one of the following factors on the performance of the tax-based contribution-aware mechanism for the live mesh-based P2P streaming approach: (1) Dynamics of parent selection, (2) Benefits of contribution-aware mechanism, (3) Effect of tax rate and peer contribution, (4) Resource Index (RI) in the system, (5) Scalability with group size, and (6) Effect of update frequency.[9]

Each simulation is run for 6,000 sec and the information is collected during the steady state when the

---

[8] Note that, real world experiments and packet-level simulations are often useful to evaluate the protocol in a realistic setting such as realistic packet level dynamics (and background traffic), and bandwidth and RTT heterogeneity. However, we focus on session-level simulations as follows: the contribution-aware mechanism assumes all connections have the same value and primarily controls resource allocation by adjusting the incoming degree of the overlay. Therefore, this mechanism is not affected by packet level dynamics, bandwidth or RTT variations.

[9] One can compare the performance of tax-based contribution-awareness in both tree- and mesh-based approaches. However, due to the inherent differences between these two approaches [11], any observed differences in the performance of contribution-aware mechanism in tree and mesh-based will be related to major differences between them.

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 low bandwidth and the rest are high bandwidth, required incoming degree to receive full quality stream is 16, the degree of willingness for high and low bandwidth peers (i.e., their outgoing degrees) are 24 and 4, respectively. The resource index is 0.5. The state collection and reporting is performed once every 10 sec.

5.1 Dynamics of parent selection

We start by examining the dynamics of changes in the number of parents that are caused by the contribution aware scheme as well as churn. Figure 3a and b show the typical evolution of the incoming degree 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 sec for a high bandwidth peers and 11 sec 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 3c presents the average incoming degree among peers whose lifetime is within the range of [x, x+10] sec. 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 sec. This also implies that peers with lifetime shorter than 60 sec, will not remain in the system sufficiently long to reach their target degree.

5.2 Benefits of contribution awareness

To examine the ability of the contribution aware scheme 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 that a peer maintained at that degree. For example, if a peer has an outgoing degree of 3 for one forth 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 4a depicts the CDF of weighted average incoming degree among high and low bandwidth peers when tax rate is 2, with contribution-aware scheme and without it (labeled as No-Cont.*). This figure clearly shows that in the absence of the contribution-aware scheme, 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 (* W_L) have a degree close to 7 whereas the degree of high bandwidth peers (* W_H) is very close to 16. Figure 4a illustrates that the contribution aware scheme 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 scheme, we present the distribution of weighted

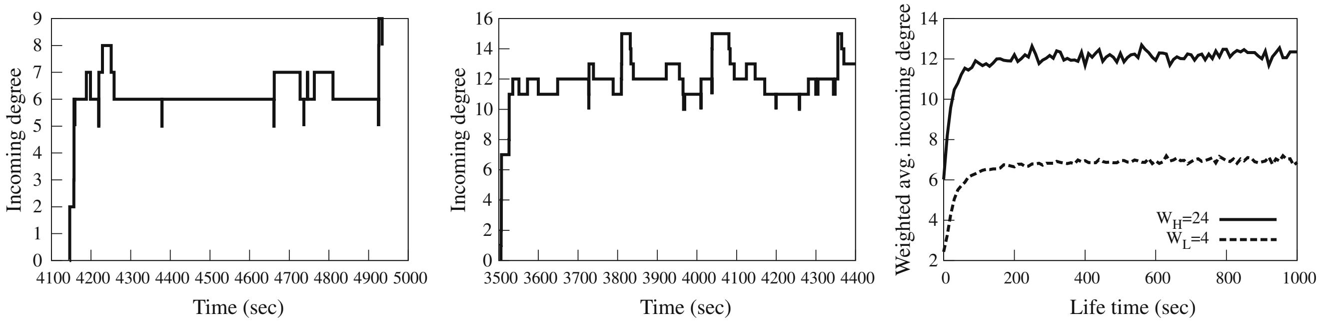

(a) Variation of incoming degree for a typical low bw peer

(b) Variation of incoming degree for a typical high bw peer

(c) Weighted avg. incoming degree of peers based on their lifetime

Fig. 3 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

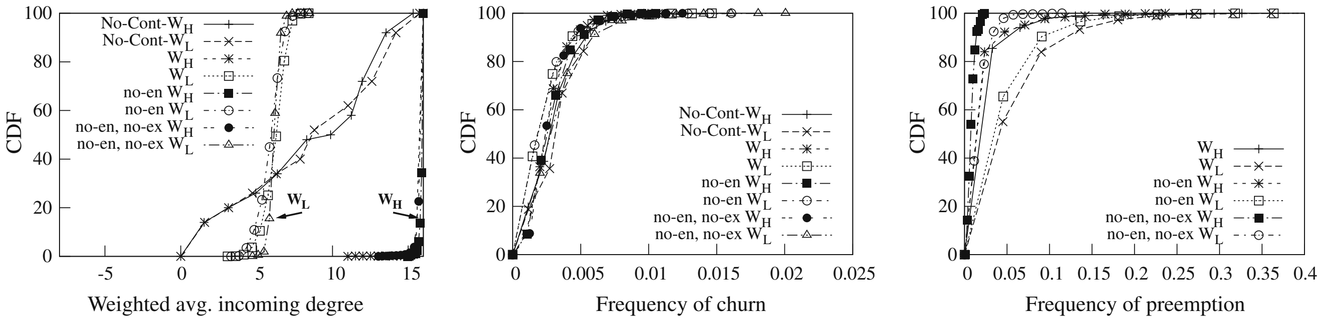

(a) Distribution of weighted avg. incoming degree

(b) Distribution of parent disconnection rate due to churn

(c) Distribution of parent disconnection rate due to preemption

Fig. 4 The effectiveness of contribution aware scheme in terms of distribution of resources (i.e., degree) among peers and parent disconnections due to churn and preemptions across different combinations of policies-($t = 2$, $W_H = 24$ and $W_L = 4$)

average incoming degree for high and low bandwidth peers in two other scenarios where (1) the En-En policy is off (labeled as no-en *); and (2) both En-En and Ex-Ex policies are off (labeled as no-en, no-ex *). Figure 4a indicates that eliminating Ex-Ex and En-En preemption policies does not lead to any visible change on the allocations of resources among peers. In other words, the En-Ex policy appears to be sufficient to achieve good performance.[10]

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 4b 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 Fig. 4a. 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 Fig. 4b for fair comparison. As expected, Fig. 4b illustrates that the normalized parent disconnection rate due to churn does not change with contribution aware scheme and does not depend on peer bandwidth (i.e., peer degree). Figure 4c presents the distribution of the average parent disconnection rate among participating peers for high and low bandwidth peers only due to preemption in all the scenarios that we examined in Fig. 4a (except for the scenario without contribution aware since no preemption occurs in that case).[11] Figure 4c shows that low bandwidth peers observe a higher rate of preemption in the base case (W$_L$) and even after disabling En-En preemption policy (no-en W$_L$). However, after disabling Ex-Ex and En-En, parent disconnection rate decreases significantly (no-en, no-ex W$_L$). 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 Section 5.3.

The stability of overlay can be also characterized in a more coarse-grained fashion. Table 3 presents the percentage of peers whose observed time between consecutive changes in parents (regardless of their cause) is at least 600 sec. Each row of the table represents a different scenario with contribution-aware scheme (including various combination of preemption policies) and without it. The table shows that in the absence of contribution-aware scheme 29% of peers are stable. The percentage of stable peers with contribution aware scheme 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 increases the percentage of stable peers to 24% which is close the observed stability without the contribution aware scheme. Since the En-En and Ex-Ex policies significantly increase the instability of the overlay without affecting the performance of the

---

[10] It 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.

[11] Note that normalizing the rate of change in parents due to preemption in Fig. 4c is not meaningful since the observed rate depends on the relative number of excess connections for each peer.

Table 3 Percentage of stable peers for scenarios with different combinations of policies with and without contribution-aware scheme

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%

contribution-aware scheme, we eliminate these two policies for the remaining evaluations in this paper.

5.3 Effect of tax rate & peer contribution

In this section, we examine how the behavior of the contribution-aware scheme changes with the following two key parameters that determine a particular scenario: (1) the value of tax rate (t), and (2) the value of peer's willingness to contribute (W_i). We consider the default parameters but with three different levels of contribution (i.e., degree of willingness or outgoing bandwidth) for high and low bandwidth peers as follows: (1) Scenario S_1: $W_H = 16$, $W_L = 6$, (2) Scenario S_2: $W_H = 24$, $W_L = 4$ and (3) Scenario S_3: $W_H = 32$, $W_L = 2$.

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 5b depicts the weighted average entitled degree among high and low bandwidth peers as a function of tax rate for all three scenarios. Figure 5a shows that the entitled degree for high and low bandwidth peers based on (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 (2) despite the existing dynamics in the connectivity among peers. Figure 5c 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.

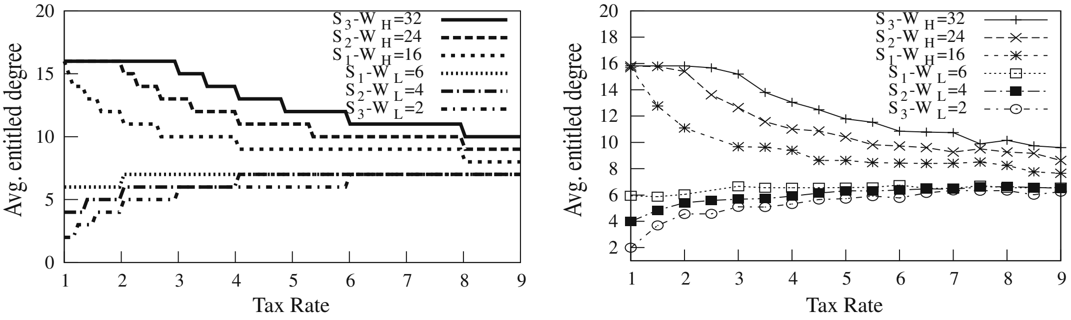

(a) Computed entitled degree-Rounded down

(b) Weighted avg. entitled degree

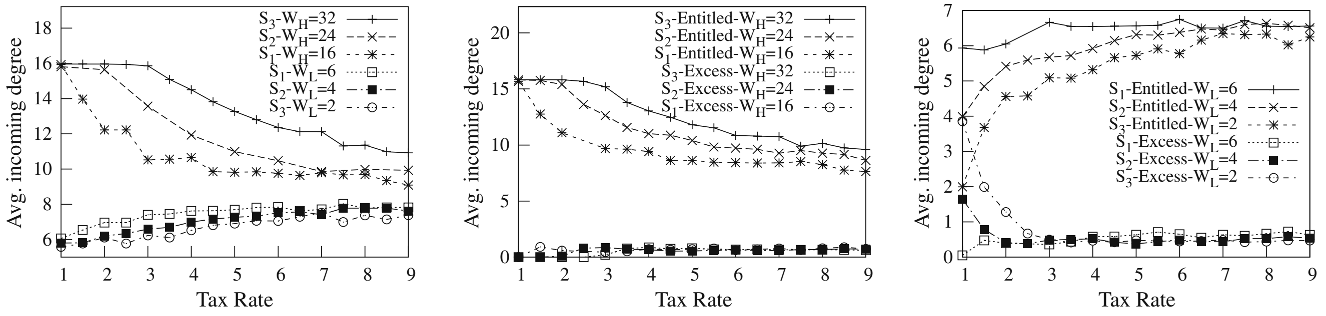

(c) Weighted avg. incoming degree

(d) Avg. entitled and excess in-degree-High BW peers

(e) Avg. entitled and excess in-degree-LowBWpeers

Fig. 5 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, Fig. 5d depicts the weighted average value of both entitled and excess degree for high bandwidth peers in three scenarios whereas Fig. 5e 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 become 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 Fig. 5e. 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 Fig. 5e. 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 scheme evenly divide excess resources among participating peers.

Utilization of resources To investigate the utilization of resources in the system, Fig. 6a 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 degree of peers in all scenarios are very close to their willingness to contribute (W_i), i.e., the contribution-aware scheme can effectively utilize available resources for different distribution of resources among peers despite the dynamics of peer participation. Figure 6b 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 a minor drop in the outgoing degree of high bandwidth peers for scenario S_3 in Fig. 6a 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 requests for an excess connection grows. This in turn reduces the utilization of resources due to the backoff in adapting the *wait interval* for retrying a rejected excess connection request.

Stability of overlay To quantify the stability of overlay, Fig. 6c 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 the saturated region, the average parent disconnection rate among low

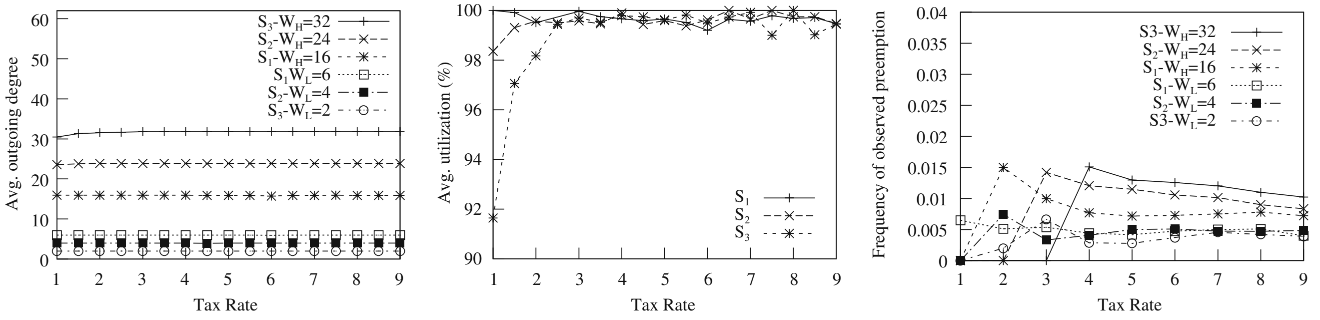

(a) Weighted avg. outgoing degree

(b) Avg. utilization

(c) Median of parent disconnection rate due to preemption

Fig. 6 Effect of tax rate and peer contribution on resource utilization and rate of preemptions 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 Fig. 5d and e), Fig. 6c reveals that high bandwidth peers surprisingly observe a higher rate of disconnection.

The above trends in the stability of parents, 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 scheme in the mesh-based P2P streaming, as we discussed in Section 3.2. 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 Fig. 5c.

5.4 Effect of resource index

We examine the effect of resource availability (or RI) on the performance of contribution aware scheme. 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 4. Different scenarios in Table 4 are derived from reported traces by earlier empirical studies [12].

Figure 7b depicts the weighted average entitled degree of high and low bandwidth peers as a function of tax rate for different scenarios. Figure 7a shows the entitled degree of high and low bandwidth peers in the same scenarios based on (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. Figure 7a and b clearly illustrate that as more resources become available (i.e., 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 connections and changes dynamics of the overlay.

Table 4 Parameters used in simulations to examine the effect of RI

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

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 Fig. 7c and d, respectively. Figure 7c clearly illustrates the saturated region for high bandwidth peers in different scenarios where they do not have any excess connection. On the other hand, Fig. 7d 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 7e shows the average outgoing degree of high and low bandwidth peers as a function of tax rate in scenarios with different RIs. 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 7f 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 peers 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 a connection to a parent. This in turns lead to an exponentially increasing wait time which reduces resource utilization. We note that while exponential increase of *wait interval* adjusts the aggregate demand for excess connections 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 Section 5.3 over small tax rates as well.

5.5 Effect of group size

We now investigate how well the contribution aware scheme scales with the average number of concurrent peers in a session.[12] Toward this end, we change the

---

[12] Note 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.

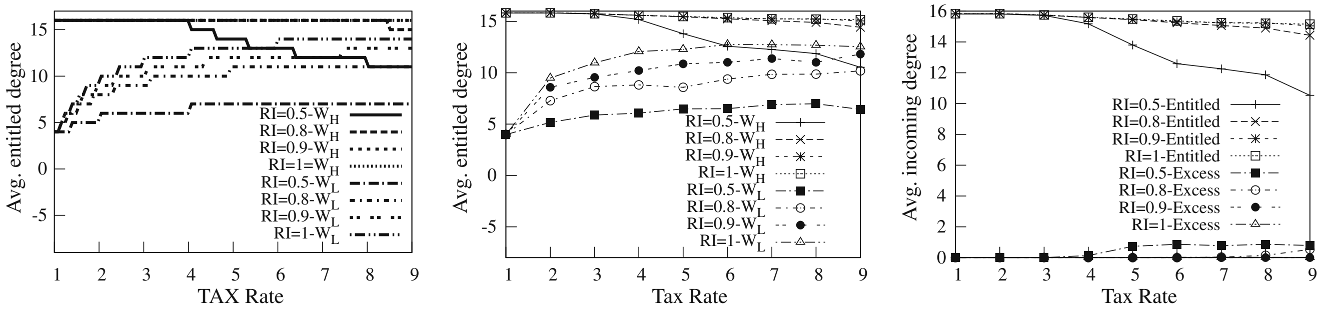

(a) Computed entitled degree-Rounded down

(b) Weighted avg. entitled degree

(c) Avg. entitled and excess in-degree-High BW peers

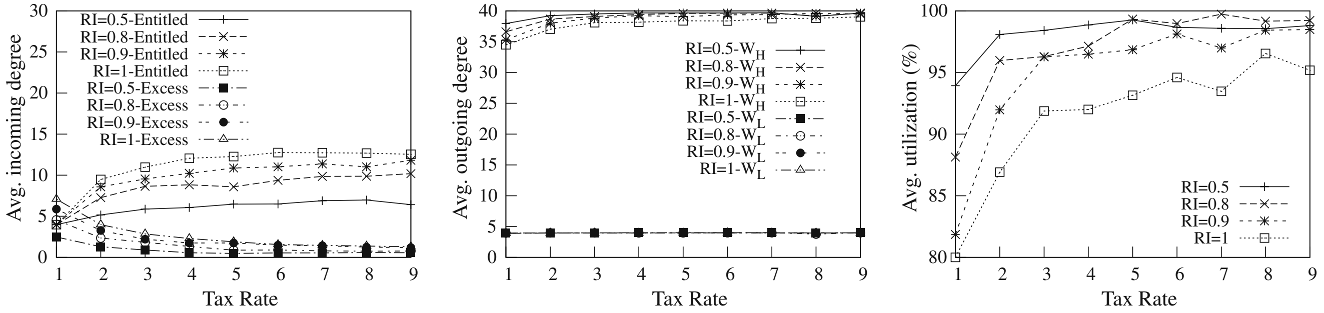

(d) Avg. entitled and excess in-degree-Low BW peers

(e) Weighted avg. outgoing degree

(f) Avg. utilization

Fig. 7 The impact of resource index on effectiveness of the contribution-aware scheme

average population from 100 to 1,000 peers where $RI = 0.5$ and high and low bandwidth peers are willing to contribute up to 24 and 4 connections, respectively.

Figure 8a depicts the weighted average in-degree of high and low bandwidth peers as a function of tax rate for three different group sizes. Figure 8b and 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 scheme is likely to scale with the number of participating peers.

5.6 Effect of update frequency

In this subsection, we explore the effect of reporting interval on the performance of the contribution aware scheme in a scenario where $RI = 0.5$ and high and low bandwidth peers are willing to contribute up to 24 and 4 connections, respectively. In general, as the update interval increases, the reported group state to

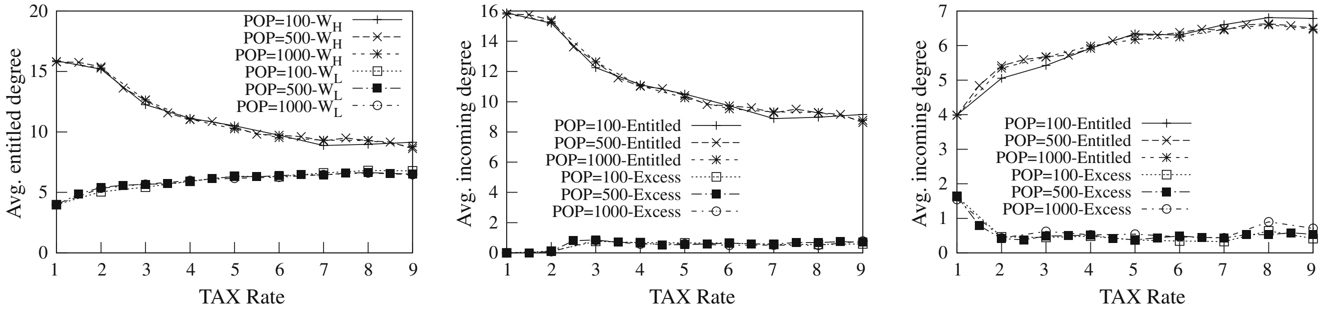

(a) Weighted avg. incoming degree

(b) Avg. entitled and excess in-degree-High BW peers

(c) Avg. entitled and excess in-degree-Low BW peers

Fig. 8 Effect of group size

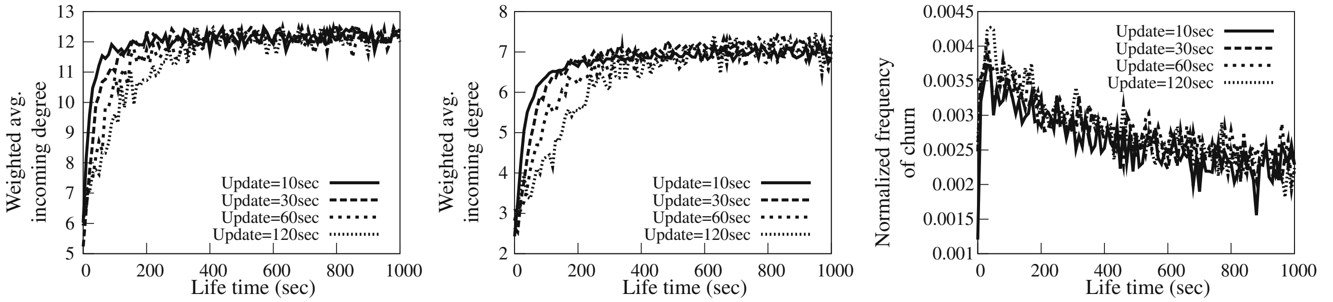

(a) Avg. in-degree as a function of peer's lifetime-High BW peers

(b) Avg. in-degree as a function of peer's lifetime-Low BW peers

(c) Normalized frequency of churn as a function of peer's lifetime

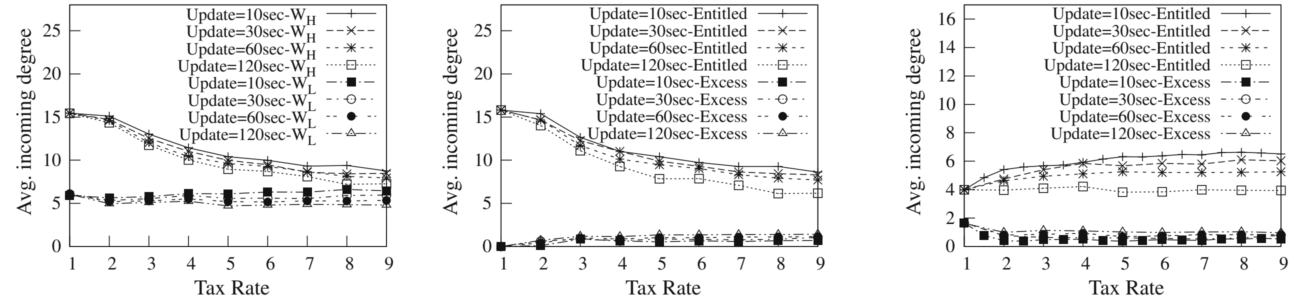

(d) Weighted avg. incoming degree

(e) Avg. entitled and excess in-degree for short-lived High BW peers

(f) Avg. entitled and excess in-degree for short-lived Low BW peers

Fig. 9 Impact of update frequency on effectiveness of the contribution-aware scheme

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 9a and b depicts the average incoming degree among high and low bandwidth peers with lifetime between [x, x+10] sec 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 1,000 sec 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 sec) 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 churn.[13] Since the amount of aggregate resources is relatively stable, once long-lived 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, Fig. 9c depicts the average value of normalized frequency of churn among peers whose lifetime is between [x, x+10] sec. 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

---

[13] One 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.

a connection between long-lived parent-child remains intact as long as aggregate resources do not change.

Figure 9e shows the average entitled and excess incoming degree for high bandwidth peers that are short-lived (lifetime less than 400 sec) for different update intervals. Figure 9f depicts the same information for short-lived, low bandwidth peers. These figures illustrate a couple of points: (1) the overall trend of change in average degree with tax rate is similar for all update intervals; (2) 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 affects the ability of 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 (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.

6 Related work

Incentive mechanism has been studied extensively in the context of P2P file sharing applications [27–33]. In general, incentive mechanisms can be divided into three categories of: payment-based, reputation-based and instantaneous methods. In the payment-based methods [28], the service recipients pay the service providers for consumed resources. Buragohain et al. [29] proposes a game theoretical model to study the potential benefits of incorporating payment-based incentive method into P2P file sharing applications. Payment-based methods typically require an infrastructure for accounting and micro-payment which leads to a limited scalability in practice.

Reputation-based methods rely on the history of a peer's contribution to the P2P network [24, 34, 35]. Reputation of peers is proportional to their overall resource contribution, and peers with higher reputation are rewarded with better performance. Habib and Chuang [34] and Tan and Jarvis [36] propose a reputation-based method where peers with higher reputation are awarded with preferential treatment in parent selection. In their proposed approach, reputation is accumulated over time across multiple streaming sessions. In general, reputation-based methods are suitable for asynchronous systems such as VoD and file-sharing applications where contribution and reward do not need to happen simultaneously and peers stay in the system long enough to build adequate reputation. In the context of live P2P streaming, empirical measurement studies have shown that the median session time of peers is very short (i.e. 25% of peers are in the system for less than 2 min) [18]. In such a dynamic system, the instantaneous contribution and demand have to be considered for a fair distribution of resources. We believe that designing an incentive mechanism that computes the instantaneous contribution of peers and allocates proportional resources to individual peers accordingly is more effective.

Instantaneous methods for incentive mechanism relax the need for maintaining long-term state information, in the form of reputation. Such methods are based on direct or indirect reciprocity. BitTorrent, a P2P file-sharing application, is a good example of direct reciprocity approach by adopting a tit-for-tat strategy [27]. In such an approach, peers upload to peers from whom they are able to download at a higher. In the context of live P2P streaming, [37] proposes an extension of BitTorrent's tit-for-tat strategy for parent selection based on local information of available bandwidth and streaming content among neighbors. Similarly, [38] and [39] extending the tit-for-tat strategy, leverage the layered encoded streaming to accommodate heterogeneity of bandwidth and enable video quality adaptation. As such direct reciprocity approaches focus on peers local information, the aggregate excess resources are randomly distributed (instead of proportionally) among peers depending on their neighbors. Furthermore, such approaches require bidirectional connections between peers to enforce the direct tit-for-tat strategy which have some implications on the properties of the overlay structure and adversely affect the performance of the P2P streaming applications [6].

In this paper, we present a form of instantaneous indirect reciprocity approach that addresses the resource management issue in the context of live mesh-based P2P streaming over a wide range of scenarios such as highly heterogeneous and asymmetric peers bandwidth and realistic churn model. To the best of our knowledge, none of the previous approaches proposes an incentive mechanism that addresses all the issues of fair distribution of excess resources, bandwidth heterogeneity and peer connectivity in the context of mesh-based live P2P streaming over highly dynamic and resource-constraint environments.

6.1 Differences with tree-based approach

The contribution aware scheme that we propose for mesh-based P2P streaming, is primarily inspired by the approach proposed by Sung et al. [12] for tree-based P2P streaming. Sung et. al. proposed a distributed mechanism based on an existing multiple-tree overlay to monitor resources in the system. Leveraging information on peers' actual contribution and peer population each peer decides the number of trees it can connect to based on a tax-based function ($\frac{1}{t} f_i + \frac{t-1}{t} \sum_{j=1}^{N} \frac{f_j}{N}$) [12, 13]. Besides incorporating the tax-based approach into different classes of P2P streaming applications (mesh vs. tree-based P2P streaming), extensive evaluating and studying the behavior of the tax function, there are some notable differences in the design level between our approach and [12] that are worth mentioning.

An important difference between our design 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 (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 several side effects: (1) the available resources in the system, is underestimated in the second term of (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, (2) the actual contribution of peer i depends on the ability and demand of other participants to use its outgoing bandwidth. (3) 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.

The other differences are due to specific requirements for mesh and tree-based approach. 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 the mesh-based streaming as follows: First, to improve received quality in the tree-based approach, each peer should find a parent in a particular tree whereas in the 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 the mesh-based approach, 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 $a_i > R_i$).

7 Conclusion

This paper presented a contribution aware mechanism for live mesh-based P2P streaming based on the notion of tax function. We examined the behavior of a commonly used tax function and described how it can be incorporated into mesh-based P2P streaming mechanisms to ensure proper allocation of resources among well behaved peers. We conducted extensive 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 and compare this approach with the tax-based contribution-aware approach presented in this paper.

Acknowledgements This material is based upon work supported in part by the NSF CAREER Award CNS-0448639. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

1. Zhang X, Liu J, Li B, Yum TSP (2005) Coolstreaming: a data-driven overlay network for live media streaming. In: IEEE INFOCOM, Miami, FL, USA, March 2005
2. Padmanabhan VN, Wang HJ, Chou PA (2003) Resilient Peer-to-Peer streaming. In: IEEE international conference on network protocols, Georgia, USA, Nov 2003
3. Castro M, Druschel P, Kermarrec A-M, Nandi ARA, Singh A (2003) SplitStream: high-bandwidth multicast in cooperative environments. In: ACM symposium on operating systems principles, Lake George, NY, USA, Oct 2003

4. Venkataraman V, Yoshida K, Francis P (2006) Chunkyspread: heterogeneous unstructured end system multicast. In: IEEE international conference on network protocols, Santa Barbara, CA, USA, Nov 2006
5. Li B, Xie S, Qu Y, Keung G, Liu J, Lin C, Zhang X (2008) Inside the new coolstreaming: principles, measurements and performance implications. In: IEEE INFOCOM, Phoenix, AZ, USA, Apr 2008
6. Magharei N, Rejaie R (2009) PRIME: Peer-to-Peer receiver-drIven MEsh-based streaming. In: IEEE/ACM transactions on networking, Aug 2009
7. Pai V, Kumar K, Tamilmani K, Sambamurthy V, Mohr A (2005) Chainsaw: eliminating trees from overlay multicast. In: IEEE international workshop on Peer-to-Peer systems, Ithaca, NY, USA, Feb 2005
8. Vlavianos A, Iliofotou M, Faloutsos M (2006) Bitos; enhancing BitTorrent for supporting streaming applications. In: IEEE global internet workshop, Barcelona, Spain, Apr 2006
9. Pianese F, Keller J, Biersack EW (2006) Pulse, a flexible P2P live streaming system. In: IEEE global internet workshop, Barcelona, Spain, Apr 2006
10. Shah P, Paris J (2007) Peer-to-Peer multimedia streaming using BitTorrent. In: IEEE international performance computing and communications conference, New Orleans, LA, USA, Apr 2007
11. Magharei N, Rejaie R, Guo Y (2006) Mesh or multiple-tree: a comparative study of P2P live streaming services. In: IEEE INFOCOM, Anchorage, AK, USA, May 2006
12. Sung Y, Bishop M, Rao S (2006) Enabling contribution awareness in an overlay broadcasting system. In: ACM SIGCOMM, Pisa, Italy, Sep 2006
13. Chu Y, Chuang J, Zhang H (2004) A case for taxation in Peer-to-Peer streaming broadcast. In: ACM SIGCOMM workshop on practice and theory of incentives in networked systems, Portland, OR, USA, Aug 2004
14. Magharei N, Rejaie R (2006) PRIME: Peer-to-Peer receiver-drIven MEsh-based streaming. In: IEEE INFOCOM, Anchorage, AK, USA, May 2006
15. Cui Y, Nahrstedt K (2003) Layered Peer-to-Peer streaming. In: International workshop on network and operating systems support for digital audio and video, Monterey, CA, USA, June 2003
16. Magharei N, Rejaie R (2005) Adaptive receiver-driven streaming from multiple senders. In: ACM/SPIE Multimedia Systems Journal, Apr 2005
17. Liu J, Li B, Zhang QQ (2003) Adaptive video multicast over the internet. IEEE Trans Multimedia 10:22–33
18. Sripanidkulchai K, Ganjam A, Maggs B (2004) The feasibility of supporting large-scale live streaming applications with dynamic application end-points. In: ACM SIGCOMM, Portland, OR, USA, Aug 2004
19. Bonald T, Massouli L, Mathieu F, Perino D, Twigg A (2008) Epidemic live streaming: optimal performance trade-offs. In: ACM SIGMETRICS, Annapolis, MD, USA, June 2008
20. Magharei N, Rejaie R (2007) Dissecting the performance of Live Mesh-based P2P Streaming. Tech Rep CIS-TR-07-05. Available at: http://mirage.cs.uoregon.edu/pub/tr07-05.pdf. Accessed 20 June 2008
21. Veloso E, Almeida V, Meira W, Bestavros A, Jin S (2002) A hierarchical characterization of a live streaming media workload. In: Internet measurement workshop, Marseille, France, Nov 2002
22. Chu Y, Ganjam A, Ns T, Rao S, Sripanidkulchai K, Zhan J, Zhang H (2004) Early experience with and internet broadcast system based on overlay multicast. In: Annual conference on USENIX, Boston, MA, USA, June 2004
23. Ngan T, Wallach DS, Druschel P (2004) Incentives-compatible Peer-to-Peer multicast. In: Workshop on the economics of Peer-to-Peer systems, Cambridge, MA, USA, June 2004
24. Habib A, Chuang J (2004) Incentive mechanism for Peer-to-Peer media streaming. In: IEEE international workshop on quality of service, Montreal, Canada, June 2004
25. Ma R, Lee SCM, Lui JCS, Yau D (2004) A game theoretic approach to provide incentive and service differentiation in P2P networks. In: ACM sigmetrics, New York, NY, USA, June 2004
26. Gil TM, Kaashoek F, Li J, Morris R, Stribling J (2004) King dataset. Available at: http://pdos.csail.mit.edu/p2psim/kingdata/. Accessed 15 May 2007
27. Cohen B (2003) Incentives build robustness in BitTorrent. In: Workshop on the economics of Peer-to-Peer systems, Berkeley, CA, USA, June 2003
28. Golle P, Leyton-Brown K, Mironov I (2001) Incentives for sharing in Peer-to-Peer networks. In: ACM conference on electronic commerce, Tampa, FL, USA, Oct 2001
29. Buragohain C, Agrawal D, Suri S (2003) A game theoretic framework for incentives in P2P systems. In: International conference on Peer-to-Peer computing, Linkoping, Sweden, Sep 2003
30. Turner D, Ross K (2003) The lightweight currency protocol. Internet Draft, Sep 2003
31. Feldman M, Lai K, Stoica I, Chuang J (2004) Robust incentive techniques for peer-to-peer networks. In: ACM conference on electronic commerce, New York, NY, USA, May 2004
32. Ranganathan K, Ripeanu M, Sarin A, Foster I (2003) To share or not to share: an analysis of incentives to contribute in collaborative file sharing environments. In: Workshop on the economics of Peer-to-Peer systems, Berkeley, CA, USA, June 2003
33. Ngan T, Wallach D, Druschel P (2003) Enforcing fair sharing of Peer-to-Peer resources. In: IEEE international workshop on Peer-to-Peer systems, Berkeley, CA, USA, Feb 2003
34. Habib A, Chuang J (2006) Service differentiated peer selection: an incentive mechanism for Peer-to-Peer media streaming. IEEE Trans Multimedia 8:610–621
35. Gupta M, Judge P, Ammar M (2003) A reputation system for Peer-to-Peer networks. In: International workshop on network and operating systems support for digital audio and video, Monterey, CA, USA, June 2003
36. Tan G, Jarvis SA (2006) A payment-based incentive and service differentiation mechanism for Peer-to-Peer streaming broadcast. In: IEEE international workshop on quality of service, New Haven, CT, USA, June 2006
37. Pianese F, Perino D, Keller J, Biersack EW (2007) Pulse: an adaptive, incentive-based, unstructured P2P live streaming system. IEEE Trans Multimedia 9:1–15
38. Liu Z, Shen Y, Panwar SS, Ross KW, Wang Y (2007) Using layered video to provide incentives in P2P live streaming. In: Workshop on Peer-to-Peer streaming and IP-TV, Kyoto, Japan, Aug 2007
39. Liu Z, Shen Y, Ross KW, Panwar SS, Wang Y (2008) Substream trading: towards an open P2P live streaming system. In: IEEE international conference on network protocols, Orlando, FL, USA, Oct 2008

Nazanin Magharei is currently a Ph.D. candidate in the Computer Science Department at the University of Oregon. She received her BSc degree in Electrical Engineering from Sharif University of Technology, Iran in 2002. Her research interests include Peer-to-Peer streaming and Peer-to-Peer overlay characterization and modeling.

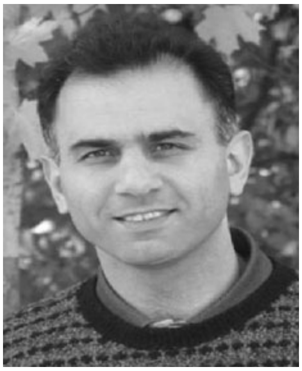

Reza Rejaie is currently an Associate Professor at the University of Oregon. From 1999 to 2002, he was a Senior Technical Staff member at AT&T Labs—Research in Menlo Park, California. He received a NSF CAREER Award for his work on Peer-to-Peer streaming in 2005. Reza received his M.S. and Ph.D. degrees from the University of Southern California in 1996 and 1999, and his B.S. degree from the Sharif University of Technology in 1991. Reza is a senior member of both the ACM and IEEE.

Yang Guo is a Principal Scientist at Research and Innovation, Technicolor at Princeton, NJ. Prior to the current position, he had a short stint at The MathWorks, and was a senior Postdoc research associate at Computer Networks Research Group at University of Massachusetts at Amherst. Yang obtained his Ph.D. degree from University of Massachusetts at Amherst in 2000, and his B.S. and M.S. degrees from Shanghai JiaoTong University in 1991 and 1994, respectively. His research interests lie in multimedia system and networking in general, and peer-to-peer, online social networking, content distribution, and media streaming in particular. He has experiences on both research and industry R&D oriented projects.