

Substream Trading: Towards an Open P2P Live Streaming System

Zhengye Liu[†], Yanming Shen[†], Keith W. Ross[‡], Shivendra S. Panwar[†], Yao Wang[†]

[†]Department of Electrical and Computer Engineering

[‡]Department of Computer and Information System

Polytechnic Institute of New York University

Abstract—We consider the design of an open P2P live-video streaming system. When designing a live video system that is both open and P2P, the system must include mechanisms that incentivize peers to contribute upload capacity. We advocate an incentive principle for live P2P streaming: a peer’s video quality is commensurate with its upload rate.

We propose Substream Trading, a new P2P streaming design which not only enables differentiated video quality commensurate with a peer’s upload contribution but can also accommodate different video coding schemes, including single-layer coding, layered coding, and multiple description coding. Extensive trace-driven simulations show that substream trading has high efficiency, provides differentiated service, low start-up latency, synergies among peers with different Internet access rates, and protection against free-riders.

I. INTRODUCTION

BitTorrent is a remarkably popular file-distribution technology, with millions of users sharing content in hundreds of thousands of torrents on a daily basis. Fundamental to BitTorrent’s success is its *openness* – the BitTorrent protocol is published in the Internet, and the source code of the baseline implementation is made widely available. This openness has enabled developers to create over 50 independent BitTorrent client implementations [1], dozens of independent tracker implementations [2], and a multitude of torrent search sites. The openness of the protocol has further fostered open discussions in both the online developer and research communities, leading to further performance and security improvements. It has also fostered innovations in the broader BitTorrent ecosystem, including recent deployments of distributed trackers, using DHTs and gossiping, in many popular BitTorrent clients.

A second key element in BitTorrent’s success is, of course, its P2P design. Since BitTorrent peers assist in file distribution, a file can be distributed to an unlimited number of peers with modest initial seeding capacity.

But with an open P2P design, it becomes necessary to incorporate an incentive mechanism to encourage peers to contribute upload bandwidth. Without such an incentive in an open protocol, clients can easily be written to free-ride (that is download without uploading) or be configured to upload at low rates. Bram Cohen, the inventor of the original BitTorrent system, recognized the need of building into the system a simple, but effective incentive mechanism [3]. Fundamentally,

BitTorrent’s incentive principle is as follows: a peer will get the file faster if it contributes more upload bandwidth to the torrent. This incentivizes users to upgrade their ISP access and/or increase the maximum upload rates (typically configurable) in their BitTorrent clients. BitTorrent provides this basic incentive using the celebrated tit-for-tat algorithm [3], in which peers trade blocks of content with each other. (Although several recent studies have shown that the tit-for-tat algorithm is not sufficient for preventing free-riders or fully incentivizing users, e.g. [4], the algorithm has nevertheless been very successful in practice.) Tit-for-tat effectively creates a differentiated service at the application layer, providing high-speed uploaders with short download times and low-speed uploaders with high download times.

In this paper, inspired by BitTorrent’s open and P2P philosophies, we consider the design of an *open P2P live-video streaming system*. Ideally, such a design would lead to an open protocol and numerous independent client, seed, and tracker implementations. Such a live P2P streaming system would also allow arbitrary users to seed live video channels – including live user-generated content – in much the same way that BitTorrent allows arbitrary users to seed files. Eventually, we expect much of the live content to emanate from handheld wireless devices, and may include such diverse sources as professors’ lectures, Little League baseball games, and political demonstrations.

Recently, several P2P live video systems have been successfully deployed. They have reported phenomenal success on their Web sites, claiming tens of thousands of simultaneous users in a single channel, with stream rates between 300 kbps to 1 Mbps. These systems include CoolStreaming [5], PPLive [6], [7], PPStream [8], UUSee [9], [10] and many more. The success of these systems shows the potential of broadcasting live video over P2P networks. However, all of these systems are closed and proprietary: the protocols are not published; independent client, seed, and tracker implementations are not possible without reverse engineering; there is no forum for discussion and criticism of the various designs; and the companies fully determine what content is distributed over their systems.

As with BitTorrent, when designing a live video system that is both open and P2P, the system must include mechanisms that incentivize peers to contribute upload capacity. But unlike BitTorrent, the incentive cannot be “download the file faster” as there is no notion of faster downloads in live streaming. We advocate a new incentive principle for live P2P streaming,

This work is supported in part by the National Science Foundation under Grant CNS-0435228, and also in part by the New York State Center for Advanced Technology in Telecommunications (CATT).

namely, *peers that upload more see higher quality video*. Ideally, peers that free-ride will receive at best poor quality; peers that upload at a high average rate can receive maximal quality; and peers that upload at more modest rates receive moderate quality. Implicit in this incentive principle is that the system will make available different video qualities. With different video coding schemes, video quality can be defined with different criteria.

In this paper we propose **Substream Trading**, a new P2P streaming design which enables differentiated video quality that is commensurate with a peer’s upload contribution. Importantly, the substream trading framework can accommodate different video coding schemes, e.g., single-layer coding, layered coding, and multiple description coding (MDC). Moreover, the design provides a framework for an open P2P live video standard. Substream trading has the following key characteristics:

- *Substream trading*: The design uses a tit-for-tat mechanism based on substream trading. In the baseline design, peers exchange substreams on an even parity basis: if Alice gives Bob exactly n substreams, Bob will reciprocate with exactly n substreams. If peers receive more substreams and correspondingly more useful bits, they can obtain a better video quality. Thus, the more substreams a peer uploads the more substreams it receives and the better the quality. This is the basic mechanism that incentivizes users to upload more to obtain better video quality. Our final design also allows for altruistic behavior, permitting Alice to over reciprocate to Bob when she has spare upload capacity.
- *Mesh design*: Peers self-organize into a mesh as a function of their available bandwidth and content. The mesh overlay is robust and easy to manage in the highly dynamic P2P environment.
- *Substream rather than chunk focus*: Peers notify, select, request and deliver video in basic content units of substreams instead of chunks. As compared to a chunk-based pull-scheme (as in PPLive), the substream design achieves a smaller playback lag with less signaling overhead.
- *Flexibility in video coding*: The design can accommodate different video coding schemes. With single-layer video, the substreams can be generated by time division of the encoded video; while with layered coding or MDC, the substreams are the video layers or descriptions, respectively.

To thoroughly investigate substream trading, we apply it on a single-layer video system and a layered video system. For both systems, the substream trading scheme can provide differentiated video quality and a high overall system performance. In this paper, we make the following contributions:

- We first make the simple, but critical observation that an open P2P live streaming system needs an incentive mechanism, and that the appropriate incentive for streaming is not “download faster” but to get better quality.
- We propose a new P2P streaming design, substream trading, based on a dynamic mesh network (as opposed to a tree design), and trading substreams (rather than trading

chunks). To our knowledge, this is the only P2P live streaming framework that supports different video coding schemes and explicitly addresses the incentive issue. The design can be used as a framework for an open P2P live streaming system.

- We examine the possible integration of the proposed substream trading scheme with different video coding schemes: single-layer coding with and without simulcast, layered coding, and MDC.
- Using traces for peer dynamics from a real-world P2P live streaming system, we evaluate the performance of substream trading using both single-layer video and layered video. We show that it is self-scaling, has high efficiency, provides differentiated service, low start-up latency, synergies among peers with different Internet access rates, and protection against free-riders.

The remainder of this paper is structured as follows. Section II discusses our fundamental design decisions. Section III presents the design of the substream trading system. We consider integration of the substream trading mechanism with different video coding schemes in Section IV, and delineate pros and cons of each. In Section V, we evaluate the performance of single-layer video and layered video systems based on trace-driven simulations. We discuss related work in Section VI and conclude in Section VII.

II. FUNDAMENTAL DESIGN CHOICES

For a live P2P video distribution, there is a source (analogous to the initial seed in BitTorrent) and a group of peers watching the video. We refer to the source and the group of peers as a *torrent*. The source encodes and divides the captured video into chunks, and disseminates the video chunks into the P2P torrent. Each peer receives chunks from the source or from other peers (or from both the source and peers). The video source may only have modest upstream Internet access (perhaps less than 1 Mbps), using cable modem, DSL, Wi-Fi, or 3G wireless networks.

A. Tree or Mesh?

There is a debate in the literature on which architecture (tree vs. mesh) is more suitable in P2P live streaming. With a tree approach, peers are organized into a single tree or multiple trees [11]–[14]. The source pumps video chunks through the trees. The tree structure can be optimized to efficiently disseminate video chunks. However, performance with trees can suffer from peer churn, and trees can be difficult to manage in a highly dynamic P2P environment.

With a mesh approach, peers self-organize into a mesh as a function of their available bandwidth and content. If there is an overlay link between two peers in the mesh, those two peers are said to be partners. In a decentralized fashion, peers form and update partnerships, and explicitly exchange content availability information with their partners. Based on this information, peers select what content to request from their partners. Currently all of the large-scale industrial deployments (PPLive, PPStream, UUSee, and so on) use a mesh design. The most important feature of the mesh approach is that the

dynamic overlay is very robust to peer churn, due to the loose relationship among peers. It has further been reported that mesh overlay outperforms tree or multiple tree from several perspectives [15]. We therefore adopt a mesh overlay in our design.

B. Chunk or Substream?

With a mesh overlay, a key design consideration is what is the basic content unit for notifying, selecting, requesting and delivering. One widely adopted approach is to divide a video into chunks, with each chunk consisting of 1-4 seconds of video. Such a chunk-based design reduces the dependency of a peer on a particular partner: a peer can request a chunk from any of its partners who has this chunk. This flexibility further increases the robustness of a system to peer churn. Additionally, it allows a peer to use its upload bandwidth with fine granularity. A peer with low upload bandwidth can serve a chunk to its partner, even though it may take a relatively long time. However, this chunk-based design has a playback lag and overhead trade-off [16], [17]. To reduce the playback lag, a peer has to send data availability notifications frequently; otherwise, the lag will normally be large.

Recently, substream-based approaches have been proposed to mitigate this problem [16], [17]. In these proposals, the video is first divided into multiple substreams by simply time dividing a single layer video. For example, assume there are a total of S substreams, substream s will contain chunks $s, s + S, s + 2S, \dots$, from the coded video chunk stream. A supplier informs its partners of the substreams it has available. A receiver then determines which substreams should be obtained from which suppliers. When a supplier is assigned to send a receiver a particular substream, it forwards any received chunk belonging to this substream to the receiver immediately, without explicit chunk request notifications. This significantly reduces the playback lag. Additionally, signaling overhead is reduced by batching the notifications of chunks into substreams. As with a chunk-based mesh design, this design is robust to peer churn. The substreams can be very thin so as to efficiently use the upload bandwidth of peers. *In many ways, the substream design provides the best features of trees (which have small playback lags) and chunk-based meshes (which are robust to high churn rates). In this work, we adopt the substream-based mesh approach.*

III. SYSTEM DESIGN

A. Mesh Overlay with Substreams

In a substream-based mesh-pull system, the source encodes a video into S substreams, with the rate of each substream denoted by r . The substreams can be generated by time dividing a single-layer video, by layered coding, by MDC or by some other scheme. Each substream is further divided into chunks of Δ seconds. At any given instant of time, a peer participating in a torrent will receive a subset of the S substreams from one or more other peers in the torrent (including possibly the source); this same peer will redistribute zero or more of the substreams it receives to other peers. In order for a peer to request substreams from other peers, it needs a mechanism

for discovering other peers in the torrent (a peer discovery service) and a mechanism for determining which substreams these discovered peers have. Figure 1 shows a simple mesh-substream system with one source, two substreams, and four peers.




Fig. 1. A simple illustration of a mesh-substream system.

After a newly arriving peer P obtains a list of other peers participating in a torrent from the peer discovery service (e.g., by tracker, DHT, or gossiping), it contacts peers on the list, searching for *partners*, with whom peer P establishes overlay links. Typically, a peer selects its partners based on some policy, which is referred as a *partner selection problem*.

After having found a sufficient number of partners (on the order of the number of substreams), peer P selects substreams from its partners and the partners sequentially push the video chunks of their selected substreams to peer P . Peer P requests *substream maps* from its partners periodically, indicating which substreams they currently have. For peer P , each of its partners may have more than one substream, and each substream may be available at more than one partner. Given the partners and their substream availabilities, peer P needs to determine which substreams should be obtained from which partners. We refer this problem as a *substream selection problem*.

From time to time, peer P may have to drop partners that do not have sufficient upload bandwidth or video content, based on some policy. This is referred as a *partnership maintenance problem*. After peer P drops partners, it will try to find replacement partners from the peer list.

B. Substream Trading

The essence of our scheme is *substream trading*: two partners exchange substreams with each other in a tit-for-tat fashion. In pure tit-for-tat substream trading, peer P sends n substreams to Q if and only if Q sends n different substreams to P . One simple observation is that if all peers use pure tit-for-tat, then each peer receives a number of substreams that is exactly equal to the number of substreams it uploads to other peers; thus, a peer with higher upload contribution is more likely to trade more substreams and eventually obtain better quality.

Peer P and peer Q may trade only one substream or more than one substream with each other. If peer P trades more than one substream, e.g., two substreams, with peer Q , then for peer P , peer Q can be considered to be two *virtual* partners, $Q1$ and $Q2$. Peer P trades only one substream with each of the two virtual partners. For this reason, in the following sections, we assume two peers only trade one substream with each other.

1) *Partner selection* : We denote the configured maximum upload rate of peer P as bandwidth C , which varies from

peer to peer. To simplify notation, assume throughout that C is a multiple of r . Thus, the peer is able to trade up to $T = \min(C/r, S)$ substreams. When a peer is trading less than T substreams and has spare bandwidth, it will search for additional partners for trading. When peer P meets another peer Q that also has spare bandwidth, the two peers should decide whether to establish the partnership.

A peer can randomly select peers from the active peers in the system, and gradually wash out the unsuitable peers by partner adaptation. To accelerate this process, a peer can select a partner from a pre-screened set of candidates. We propose to use substream maps of peers for the pre-screening. Two peers (say P and Q), who intend to establish the partnership, first exchange substream maps with each other. Based on the substream maps, peer P knows the number of substreams peer Q currently has, and vice versa. Assume peer P and peer Q have s_P and s_Q substreams, respectively. Without loss of generality, we assume $s_P \geq s_Q$. In our design, since peer P has more content (and potentially has higher upload bandwidth) than peer Q, peer Q *will* accept peer P as a partner. For peer P, since peer Q has less content, it needs to make a decision whether to accept peer Q. This is because less content normally indicates lower upload bandwidth of a peer. We propose a simple scheme for peer P to make such a decision. Peer P agrees to form a partnership with probability p_P , which is simply given by:

$$p_P = \frac{S - s_P}{S}. \quad (1)$$

A peer with less substreams available is more likely to accept a partner, even though this partner has less content.

2) *Partnership maintenance*: After two partners establish the substream trading relationship, both of them monitor the service quality of each other. When a peer cannot fulfill the trading commitment due to lack of substream content or bandwidth, its partner will adapt this peer. We propose a *leaky bucket algorithm* for partners to monitor the trading procedure with each other.




Fig. 2. Leaky bucket algorithm for monitoring partner's service quality.

As shown in Figure 2, a peer maintains a leaky bucket for each of its partners. It counts the number of chunks received from partner n . When the peer receives a chunk from partner n , it will fill Δ bytes into the corresponding bucket. The decrement rate of the leaky bucket is set to r . In our design, the leaky bucket has B bytes initially. Whenever the leaky bucket is empty, the peer will break the trading relationship with its partner n . The leaky bucket algorithm can handle Internet jitter, short term content and/or bandwidth deficiency. The initial setup of the leaky bucket (with B bytes) shares some similarity with the “optimistic unchoking” strategy in BitTorrent. A newcomer can get served first and use the

received substream to trade other substreams before the leaky bucket empties.

C. Substream Selection




Fig. 3. (a)Substream maps; (b)Abstracted substream maps.

In our substream trading system, a peer should determine which substreams should be obtained from which partners. Before substream selection, the peer periodically requests substream maps from all its partners. Figure 3(a) shows the typical substream maps of peer P and its partners P1, P2, and P3. As an example, the substream map of P1 indicates that it has three substreams, and the sequence numbers of the last chunks from substreams 1, 2, 3 are 100, 101, 94, respectively. Assume that there is no chunk loss during the transmission (e.g., by using the TCP connections for chunk delivery, or by inserting sufficient FEC chunks), the sequence number of the last chunk is sufficient to indicate the chunk availability of a particular substream. For P1, although it has three substreams available, only substreams 1 and 2 can be pulled by peer P, since peer P already has more chunks from substream 3 than P1. Thus, in Figure 3(b), only substreams 1 and 2 are indicated as available in P1. Note that this automatically eliminates possible loops while delivering a substream in a mesh network. We record this processed data availability of partners in *abstracted substream maps*, as shown in Figure 3(b).

With the above definition, we can formulate the substream selection problem as an optimization problem. Assume a peer has N partners $1, \dots, N$ for requesting substreams. The set of available substreams in partner n is defined as \mathcal{S}_n . Let $x_{sn} = 1$ denote that substream s is received from partner n . Since a partner can send at most one substream to a peer (with the virtual partner definition), x_{sn} is subject to the following constraint:

$$\sum_{s \in \mathcal{S}_n} x_{sn} \leq 1, \quad n = 1, \dots, N.$$

Furthermore, since a substream only needs to be sent from one partner, we have the following constraint as well:

$$\sum_n x_{sn} \leq 1, \quad s = 1, \dots, S.$$

By substream selection, a peer tries to maximize the received video quality. With different video coding schemes, the importance of each substream could be different. For example, with single-layer coding and MDC, the substreams have equal importance, and the peer only needs to maximize the number of received substreams; while with layered coding, the substreams have unequal importance, and the peer needs to take into account the importance of different substreams and

maximize the received video quality. To reflect the importance of substreams, we assign weights to each substream, with a larger weight indicating a more important substream. With these weights, the optimal substream selection problem can be converted to maximizing the weighted sum of all substreams as follows:

$$\begin{aligned} \max \quad & \sum_{s=1}^S w_s x_{sn} \\ \text{s.t.} \quad & \sum_{s \in \mathcal{S}_n} x_{sn} \leq 1, \quad n = 1, \dots, N, \\ & \sum_n x_{sn} \leq 1, \quad s = 1, \dots, S, \end{aligned} \quad (2)$$

where w_s denotes the weight of substream s . This is the classical maximum weight matching problem in a bipartite graph as shown in Figure 4, which can be solved with a complexity of $O(S^3)$ [18]. We will give examples of an appropriate assignment of weights for single-layer video and layered video in Section IV.




Fig. 4. A bipartite graph representing the substream selection algorithm.

D. Altruistic Peers

A peer is altruistic at any given time if its aggregate upload rate is higher than its aggregate download rate. With the presence of altruistic peers, bandwidth-deficient peers can possibly receive video at rates higher than their contributions. We do not force a peer to donate bandwidth. We assume that a bandwidth-rich peer will only consider donating if it is receiving all substreams (i.e., the full video rate) and still has surplus upload bandwidth.

Assume a peer is willing to contribute upload bandwidth C where $C/r > S$. This peer can donate $C/r - S$ substreams, i.e., it can provide other peers substreams without reciprocation. In our design, it is the benefactor that determines who will be its beneficiaries. For simplicity, an altruistic peer can randomly select its beneficiaries. A biased donation can also be used. For example, the benefactor can first donate the substreams to its existing partners, and then other peers. As we will see in Section V, such a biased scheme helps to combat free-riders.

IV. VIDEO CODING

In this section, we show how substream trading can be applied to a variety of different video coding schemes, including single-layer video, layered video, MDC, and simulcast.

A. Single-Layer Video

With single-layer video, a video is time-divided into S substreams, each of rate r . If a peer receives all S substreams,

it can reconstruct the video perfectly; otherwise, the peer will reconstruct a corrupted video or experience frame freeze and discontinuous video playback. With substream trading, if the upload bandwidth of a peer is higher than the video rate, it can trade all substreams and obtain a continuous video playback; otherwise, it can only trade part of substreams and obtain a less continuous video playback. This provides the basic incentives for peers to contribute upload bandwidth.

Note that not all peers in a single-layer system are necessarily self-supported, with an upload bandwidth higher than the video rate. If no peer contributes at a rate higher than the video rate, the peers with low upload bandwidth cannot receive all S substreams. This may discourage such peers from using the application. But as in BitTorrent, where peers do not always immediately quit after receiving the entire file, we expect to see some altruistic behavior in P2P live streaming [19].

With single-layer video, substreams have equal importance. Thus, the weight for each substream for selection can be defined as $w_s = 1, s = 1, \dots, S$.

B. Layered Video

In recent years, significant advances have been made in layered coding. Now H.264/SVC (layered coding) achieves a rate-distortion performance comparable with H.264/AVC (single-layer coding), with the same visual reproduction quality typically achieved at $\pm 10\%$ bit rate [20]. It is reported that a real-time system with H.264/SVC encoder and decoder has been successfully implemented [21]. Thus, thanks to these recent advances, layered coding is a viable candidate for P2P live streaming systems. Furthermore, layered video is a particularly useful concept for P2P, even more so than for client-server. *In P2P, layered video responds to heterogeneous upload rates as well as heterogeneous download rates and congestion.*

With substream trading, a peer with a higher upload contribution will trade more layers and consequently obtain a better video quality. Furthermore, with layered video, even a small number of layers can lead to passable video quality without discontinuity. Peers with low upload bandwidth can therefore be self-sustaining and less dependent on altruistic peers.

To reflect the layer dependencies, the weights for substream for selection can be set to $w_s = 2^{S-s}, s = 1, \dots, S$. With these weights a lower layer is more important than the sum of all its upper layers, which is consistent with layered coding.

C. MDC

Like layered video, MDC generates multiple substreams. But unlike layered video, each of the substreams is of equal importance, so that video quality is only a function of the total number of substreams received and not of which substreams are received. Because all substreams have equal importance, designing a P2P live streaming system using MDC (rather than layered video) is appealing and more straightforward. A number of recent papers have investigated combining a large number of MDC substreams with P2P to create P2P video streaming systems [12], [13], [22], [23]. Like layered video,

the proposed substream trading *can* be applied with MDC. In this case, the w_s for each description should be equal.

The efficiency of MDC depends on the trade-off among the achievable qualities with different number of descriptions [24]. MDC is inherently inefficient when a large number of descriptions are created. Among the few proposed methods for generating a large number of descriptions (> 4), MD-FEC [25] together with scalable video coding, and temporal subsampling [26], both introduce significant ($> 60\%$) overhead, compared with single-layer video. *The inefficiency of MDC largely prevents its usage in practical P2P live streaming systems.* For this reason, we do not further consider MDC in this paper.

D. Simulcast

With simulcast, the video source encodes a video into multiple independent streams by using single-layer coding, with each stream having a different rate. Each stream then gets distributed within a separate torrent, with no interaction among the torrents. Compared to layered video, simulcast requires more source bandwidth. For example, a set of 5 video simulcast streams, from 200 kbps to 1 Mbps with a 200 kbps step size, would minimally require 3 Mbps at the source; the corresponding layered design would minimally require 1 Mbps. When the sources are residential broadband connections, this becomes a critical issue. When the sources have a sufficient upload capacity to support several torrents, substream trading can be applied within each torrent, as discussed in the single-layer video system. However, such a design would suffer from lack of sharing across the simulcast streams, as we briefly discuss in Section V-C.

V. TRACE-DRIVEN SIMULATION

We conduct extensive simulations to evaluate the performance of the single-layer system and the layered system with substream trading. We also investigate the impact of cheating behavior.

A. Simulation Setup

We developed a chunk-level discrete-event simulator in C++. In our simulations, we assume that the end-to-end bandwidth bottleneck is at the access links and not in the Internet core. We do not simulate the delay induced by routing in the Internet core; instead, we randomly assign the end-to-end propagation delay between each pair of peers to be between tens of milliseconds to hundreds of milliseconds. Such an abstraction speeds up the simulation and we believe it can still give us an accurate evaluation of the system.

Peer dynamics are simulated by traces collected from a real-world P2P live streaming system – PPLive [6], [7]. The traces record the arrival and departure times of the users for different channels. We select the trace of a popular Chinese TV channel, CCTV3, to drive our simulations. This one-day trace was collected from Nov 22nd 17:43, 2006 to Nov 23rd 17:43, 2006, and there were totally more than 100,000 video sessions during this period. Figure 5 shows the evolution of the number of users viewing this channel. This trace covers a variety of




Fig. 5. Evolution of number of users viewing a TV channel in the PPLive network.

typical scenarios in P2P live video networks, such as small systems (less than 200 concurrent users), large systems (more than 9000 concurrent users), short video sessions (shorter than one minute), long video sessions (longer than 16 hours), and flash crowds.

In our simulations, the upload bandwidth of the source is set to 2 Mbps. The upload capacities of peers are assigned randomly according to the distribution of Table I. To come up with an accurate bandwidth distribution of Internet users, we jointly consider the measurement studies in [27] and [28]. The overall distribution of residential peers and Ethernet peers is obtained from [27], while the detailed bandwidth distribution of residential peers is obtained from [28]. We exclude modem peers and ISDN peers due to their low upload and download capacities. Note that the upload capacities of Internet users are highly heterogeneous. Because peers may not be willing to contribute their entire upload bandwidth, in our simulations we assume that the peers only contribute portions of their upload bandwidth for trading, which are indicated in Table I. For example, the 256 kbps peers contribute 150 kbps for trading.

The playback lag between the peers and the source is set to ten seconds. Every one second, a peer exchanges substream maps with its partners. Accordingly, the peer re-selects the substreams based on the most recent substream maps. The substream maps and request notifications can be piggy-backed in the video chunks. The initial level of the leaky bucket is set to $10 \cdot r$, i.e., 10 seconds of video at rate r . In our simulations, a peer can at most contact eight peers to check if they are the suitable partners during one second. If we let a peer contact more peers simultaneously, it can locate the partners faster. But this will introduce higher overhead.

B. Single-Layer System

In this section, we investigate the substream trading system with single-layer video. We begin by assuming all peers in the system follow the proposed protocol, without tampering with the protocol to maximize their own benefits. We then consider free-riding and cheating behavior of peers.

1) *Differentiated services:* We evaluate the single-layer video system in two scenarios. In the first scenario, the system is underloaded and the supplied bandwidth (i.e., average upload bandwidth of the entire system) is higher than the demanded bandwidth (i.e., the video rate). In the second scenario, the system is overloaded and the supplied bandwidth is lower than the demanded bandwidth. To represent video continuity, we introduce the *received chunk ratio*, which is

TABLE I
PEER UPLOAD BANDWIDTH DISTRIBUTION (KBPS)

Total upload bandwidth (kbps)	256	320	384	448	512	640	768	1024	1500	> 3000
Distribution (%)	10.0	14.3	8.6	12.5	2.2	1.4	6.6	28.1	1.4	14.9
Contributed upload bandwidth	150	250	300	350	400	500	600	800	1000	1000




Fig. 6. Cumulative distribution of received chunk ratio. (a) Underloaded scenario; (b) Overloaded scenario.

defined as the ratio between the number of received video chunks and the number of encoded video chunks.

With the upload bandwidth distribution shown in Table I, the average contributed upload bandwidth in the system is about 540 kbps. To investigate both the underloaded and overloaded scenarios, we consider two video rates, 500 kbps and 700 kbps. The encoded videos are time divided into 10 and 14 substreams, respectively, with the rate of each substream being 50 kbps. Each chunk has a size corresponding to $\Delta = 250$ ms. We assume all peers become altruistic if they obtain the video at the full rate. In our simulations, the altruistic peers prefer to donate first to their trading partners.

In Figure 6, we show the CDF of the received chunk ratio for different upload bandwidths. From the ten types of peers (see Table I), five types of peers are shown in the figure, including the peer type with the lowest upload bandwidth (256 kbps), the peer type with the highest upload bandwidth (> 3000 kbps), the peer types with modest upload bandwidth and with many peers (448 kbps and 1024 kbps) and the peer type with the fewest peers (640 kbps). Figure 6(a) shows the results of the underloaded scenario. We observe that almost all peers have a high received chunk ratio that is close to 1.0, indicating that all peers can receive a continuous video quality. But we emphasize that the video qualities of the low bandwidth peers are highly dependent on the altruistic behavior of the high bandwidth peers.

Figure 6(b) shows the CDF of the received chunk ratio under the overloaded scenario. In this case, the upload bandwidth in the system cannot support the video rate for all peers. On average, each peer can at most receive 77% (540/700) of




Fig. 7. Cumulative distribution of the received chunk ratio for free-riders.

the video chunks. This means that some peers will have very discontinuous video quality. With our substream trading design, the peers that have an upload contribution higher than 700 kbps are self-supported and receive continuous video quality. This is verified in the figure, where the peers with 1024 kbps and higher upload bandwidth can receive almost all video chunks. For the peers whose upload contribution is lower than 700 kbps, the received chunk ratio increases with their upload contribution. This incentivizes peers to contribute more upload bandwidth.

2) *Free-riding and cheating*: We now consider free-riding and cheating behavior. As observed in BitTorrent, uncooperative users may tamper with the BitTorrent protocol to maximize their downloading speed [19]. Similarly, in an open P2P live streaming system, a free-rider may try to receive the same video quality as regular peers with minimum upload bandwidth. An open P2P live streaming system should be able to discourage free-riding, by providing minimum video quality for free-riders.

We assume the free-riders try to receive the video rate without contributing any upload bandwidth. In our simulations, we consider one type of cheating, where the free-riders untruthfully announce that they have a high upload bandwidth for trading but do not have any content currently. In other words, a free-rider is always pretending to be a newcomer to the system. It is possible for free-riders to take advantage of the initial state of the leaky bucket algorithm and get served for free during a period B/r , and then repeat this behavior to obtain more free service. We now examine whether a free-rider can receive good video quality under the overloaded scenario. In our simulations, we randomly select 10% of peers as free-riders and assume that a free-rider establishes partnerships with up to 14 peers.

Figure 7 shows the CDF of the received chunk ratio of the free-riders. We observe that even with cheating, the free-riders get a very low received chunk ratio. Since a free-rider has no content and bandwidth to trade, it will be dropped by its partners after B/r seconds. Furthermore, since altruistic peers prefer to donate spare bandwidth to their trading partners, it is less likely for a free-rider to get the donation. Additionally, compared with a regular peer, a free-rider has a much higher




Fig. 8. Behavior of typical peers.

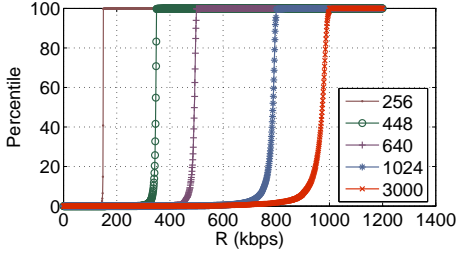


Fig. 9. Cumulative distribution of average useful download rate.

overhead of searching partners and maintaining partnerships.

We nevertheless acknowledge that it may be possible for a free-rider to obtain an acceptable chunk ratio if it continuously establishes partnerships with many peers ($\gg 14$). This will come at the cost of increased overhead. This situation is similar to BitTorrent [4], and it therefore appears impossible to fully defend against free-riders in a single-layer system. However, for a layered system, we will see even greater protection against free-riders.

3) *Summary of single-layer system:* The single-layer video substream trading system has the following properties:

- In an overloaded system, where it is not possible for all peers to get acceptable quality, the peers that upload at rates higher than the video rate do receive all substreams and have maximal quality.
- In an underloaded system where some peers have upload capacity lower than the video rate and others higher than the video rate, the system provides maximal quality to all peers, provided that the high-capacity peers are altruistic.
- Unless free-riders are extremely zealous about cheating, free-riders obtain poor-quality video.

C. Layered System

We now investigate substream trading with layered video. In our simulations, the video is encoded into 20 layers, with each layer being 50 kbps and the full video rate being 1 Mbps. For this system, since each peer has a maximum upload rate of 1 Mbps or less, none of the peers can be altruistic.

1) *Differentiated service:* Figure 8 shows the number of decodable layers of five randomly selected video sessions during their first 20 minutes. We observe that each peer receives a number of layers that is commensurate with its upload contribution. All peers reach a stable state within 100 seconds. Once a peer reaches its stable state, the video quality is generally smooth, without significant variation.




Fig. 10. Cumulative distribution of the smoothness index.

Figure 9 shows the CDF of the received video rates across all video sessions. We observe that almost all peers receive a video rate that is commensurate with their upload contributions. This demonstrates that substream trading provides differentiated services in a layered video system. Furthermore, the system has no bias among the different peer types. For example, the peers with 640 kbps upload bandwidth, which only make up 1.4% of the peer population, also obtain a video quality that is commensurate with their upload contributions.

2) *Video quality:* The received video rate can largely determine the received video quality. However, in addition to a high received video rate, it is desirable to receive continuous and smooth video quality. Unlike single-layer coding, which needs to receive the full video rate to decode and play a video, layered coding can provide basic video quality if only the base layer is received. For simplicity, we assume the base layer is fully encapsulated in layer 1, and a discontinuity occurs only if the video chunks of layer 1 cannot be received correctly. In our simulation, we observe that all video sessions receive more than 99.99% of video chunks from layer 1; thus, the peers rarely experience playback discontinuity.

With layered video, a receiver may receive a varying number of layers, which degrades the user experience. We introduce a smoothness index to evaluate video smoothness in our simulations. The smoothness index is defined as follows:

$$\Phi = \frac{1}{t-1} \sum_{i=2}^t \frac{|a(i) - a(i-1)|}{a(i-1)}, \quad (3)$$

where $a(i)$ is the number of received layers in time slot i , and t is the duration of a video session in terms of time slots. The smoothness index indicates how frequently and dramatically the number of received layers is changing. When Φ is larger, the video is less smooth. In our simulations, we set the length of a time slot equal to the chunk duration Δ .

Figure 10 shows the CDF of the smoothness index of the selected video sessions. We observe that the smoothness index is very low for all peer types. (A smoothness index of 0.005 is extremely low, and most peers have an index much lower than 0.005.) This verifies that peers receive a very smooth video quality, which can also be observed in Figure 8. Note that although the highest upload bandwidth peers have slightly larger smoothness index, this does not mean these users see more variation of video quality than the low bandwidth peers. When a large number of layers is received on average, the quality is already very good and having slightly more or less number of layers does not produce as much variation in video




Fig. 11. Cumulative distribution of start-up delay across all video sessions.




Fig. 12. Interaction across peer types.

quality, as in the case where on average a low number of layers are received.

3) *Start-up delay*: In P2P live streaming, start-up delay is the time from when a channel is selected until actual playback starts on the screen. This is a critical performance issue, particularly for users who do a lot of channel surfing. Before playback can begin, a peer needs to build an initial reservoir of video chunks to deal with Internet jitter and peer churn. With single-layer video, a peer needs to build the initial reservoir for *all* substreams; but with layered video, the peer only needs to build the initial reservoir for layer 1. In our simulations, if a peer finishes building the reservoir of video chunks for the next three seconds, it starts decoding and playing the video. Figure 11 shows the CDF of the start-up delay for both the single-layer video system (with 500 kbps video rate) and layered video system. We observe that the start-up delay of the layered video system is significantly shorter than that of the single-layer video system. This is because with layered video, only layer 1 is needed to build the initial reservoir and provide passable video quality.

4) *Interaction across peer types*: A natural question is whether the resulting layered system essentially creates a stratified system, where peers of the same type primarily share amongst each other and not with peers of other types. To explore this issue, for each peer, we record the download traffic from peers that have higher upload bandwidth (denoted as Higher), peers that have the same upload bandwidth (denoted as Equal), and peers that have lower upload bandwidth (denoted as Lower), and normalize by the peer's total download traffic. We average these values over all peers with the same upload bandwidth, and plot them in Figure 12. We observe that for all five types of peers, there exists a large amount of interaction across peer types. This is especially true for the peer types that have relatively few members, e.g., the peers with 640 kbps upload bandwidth.

Recall that with simulcast, peer types are separated into different simulcast torrents and there is no interaction across




Fig. 13. Video quality of free-riders. (a) Behavior of typical free-riders; (b) Cumulative distribution of smoothness index.

peer types. P2P live streaming with layer trading, however, as shown in Figure 12, provides major synergies across peer types. For example, a low-bandwidth peer may serve a high-bandwidth peer with a lower layer. More importantly, with layer trading, a peer type with a small number of members can easily find partners outside its type for trading. This can greatly improve the overall quality of service for the system.

5) *Free-riding and cheating*: We investigate free-riding and cheating in the layered video system. We consider the same type of cheating as discussed in the single-layer video system. In order to receive the full video rate (1 Mbps), a free-rider attempts to locate 20 partners simultaneously. Figure 13 shows the video quality of free-riders in the layered video system. Figure 13(a) plots the behavior of a typical free-rider. We observe that free-riders rarely receive video at an average rate higher than 200 kbps. Figure 13(b) shows the CDF of the smoothness index of the free-riders. Most free-riders have a smoothness index that is ten times higher than that of the regular peers, which verifies that the free-riders receive very variable video quality. This is because the free-riders are frequently dropped by their partners. On average, a free-rider is dropped by one of its partners every 0.6 seconds, leading to a high overhead for locating and managing partners. The low video quality and high overhead cost should largely discourage free-riders.

With our proposed partner selection mechanism, it is less likely for a free-rider to establish partnerships with high upload bandwidth peers. Even though a free-rider can locate a large number of partners, most of these partners will be low upload bandwidth peers. With layered video, since these partners only have lower layers, the free-rider can only obtain the lower layers and consequently low video quality. Therefore, the layered system is more robust to free-riders.

6) *Summary of layered video system*: The layered video substream trading system has been shown to have many desirable properties:

- It provides differentiated service – the video quality that a peer receives is commensurate with its upload rate. Thus peers have an incentive to upload as much as they can.
- For non-free-riding peers, there is little variation in video quality due to fluctuations in the received number of layers.
- The start-up delay is small, and significantly shorter than the start-up delay of single-layer system.
- Peers in different upload bandwidth categories synergistically share layers with each other.
- Aggressive free-riders receive the video at a low rate with relatively high quality fluctuations.

VI. RELATED WORK

Over the past few years, there has been a number of proposals for live P2P video in the research community [5], [11]–[14]. None of these papers, however, addresses built-in (tit-for-tat) incentives or the design of open P2P streaming systems. Within the context of cooperative peers, Sung *et al.* have recently proposed an MDC-based multiple-tree scheme that uses a novel taxation scheme to provide differentiated services [22]. However, this proposal does not include built-in incentives, assumes cooperative peers, and furthermore uses MDC encoding (which is inherently inefficient as discussed in Section 4.3).

There are three very recent proposals on using tit-for-tat incentives in the context of P2P live video streaming. In a workshop paper, we proposed a tit-for-tat scheme for layered video for *chunk-based systems* [29]. The scheme proposed in this paper has several advantages over that in [29]. First, in this paper we trade substreams rather than chunks, which significantly reduces playback lag and overhead. Second, the framework of this paper can be applied to a variety of coding schemes, including layered coding, MDC, and single-layer coding. Mol *et al.* propose an MDC-based multiple-tree scheme that employs tit-for-tat incentives [23]. Each description is distributed over a separate tree, and peers belonging to different trees exchange descriptions with each other. This approach is based on MDC (which is inherently inefficient), cannot be easily adapted to layered video or single-layer video, and restricts a peer to trade only the description corresponding to the tree to which it belongs. Finally, Pianese *et al.* propose a chunk-based mesh-pull scheme with single-layer video [30]. The scheme applies a combination of tit-for-tat and donation strategies to provide incentives. In particular, peers with higher upload contribution have more buffered data and are more robust to peer dynamics. However, this scheme is limited to single-layer video and has low throughput.

Unlike [29] [23] [30], the current proposal provides a framework for providing incentives in live P2P video streaming systems. This framework can accommodate a variety of coding schemes. Furthermore, the framework has been optimized for performance, providing differentiated service, high throughput, resiliency to churn, and short start-up delays. The scheme proposed here can serve as a blueprint for an open P2P live video streaming system.

VII. CONCLUSION

We have argued that built-in incentives are critical for the design of an open P2P live video streaming system. In this paper, we proposed a framework with live video streaming which has built-in incentives and can accommodate a variety of video coding schemes. In particular, we have shown that substream trading with layered video has many desirable properties, including differentiated service, short start-up delays, synergies across peer types, and protection against free-riders.

REFERENCES

- [1] http://en.wikipedia.org/wiki/BitTorrent_client.
- [2] <http://torrentfreak.com/5-most-popular-bittorrent-trackers-070924/>.

- [3] B. Cohen, "Incentives build robustness in BitTorrent," in *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, June 2003.
- [4] N. Liogkas, R. Nelson, E. Kohler, and L. Zhang, "Exploiting BitTorrent for fun (but not profit)," in *IPTPS*, Santa Barbara, February 2006.
- [5] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "DONet: A data-driven overlay network for efficient live media streaming," in *IEEE INFOCOM*, Miami, March 2005.
- [6] PPLive, <http://www.pplive.com/>.
- [7] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross, "A measurement study of a large-scale P2P IPTV system," *IEEE Trans. on Multimedia*, vol. 9, no. 8, December 2007.
- [8] PPStream, <http://www.ppstream.com/>.
- [9] UUsee, <http://www.uusee.com/>.
- [10] C. Wu, B. Li, and S. Zhao, "Characterizing peer-to-peer streaming flows," *IEEE JSAC*, vol. 25, no. 9, December 2007.
- [11] Y. Chu, S. G. Rao, and H. Zhang, "A case for end system multicast," in *ACM SIGMETRICS*, Santa Clara, June 2000.
- [12] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai, "Distributing streaming media content using cooperative networking," in *ACM NOSSDAV*, Miami, May 2003.
- [13] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "SplitStream: High-bandwidth content distribution in a cooperative environment," in *IPTPS*, Berkeley, February 2003.
- [14] V. Venkataraman, P. Francis, and J. Calandrino, "Chunkyspread: Heterogeneous unstructured tree-based peer-to-peer multicast," in *IEEE ICNP*, Santa Barbara, November 2006.
- [15] N. Magharei, R. Rejaie, and Y. Guo, "Mesh or multiple-tree: A comparative study of live P2P streaming approaches," in *IEEE INFOCOM*, Anchorage, May 2007.
- [16] M. Zhang, Q. Zhang, and S.-Q. Yang, "Understanding the power of pull-based streaming protocol: Can we do better?" *IEEE JSAC*, vol. 25, no. 9, December 2007.
- [17] B. Li, S. Xie, G. Y. Keung, J. Liu, I. Stoica, H. Zhang, and X. Zhang, "An empirical study of the Coolstreaming+ system," *IEEE JSAC*, vol. 25, no. 9, December 2007.
- [18] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to algorithms*. MIT Press, 2000.
- [19] M. Zghaibeh and K. G. Anagnostakis, "On the impact of P2P incentive mechanisms on user behavior," in *NetEcon+IBC*, San Diego, June 2007.
- [20] M. Wien, H. Schwarz, and T. Oelbaum, "Performance analysis of SVC," *IEEE TCSVT*, vol. 17, no. 9, September 2007.
- [21] M. Tchen, R. Cazoulat, A. Graffunder, A. Hutter, and P. Amon, "Real-time system for adaptive video streaming based on SVC," *IEEE TCSVT*, vol. 17, no. 9, September 2007.
- [22] Y.-W. Sung, M. Bishop, and S. Rao, "Enabling contribution awareness in an overlay broadcasting system," in *ACM SIGCOMM*, Pisa, September 2006.
- [23] J. D. Mol, D. H. P. Epema, and H. J. Sips, "The Orchard algorithm: Building multicast trees for P2P video multicasting without free-riding," *IEEE Trans. on Multimedia*, vol. 9, no. 8, December 2007.
- [24] Y. Wang, A. R. Reibman, and S. Lin, "Multiple description coding for video delivery," *Proceedings of the IEEE*, vol. 93, no. 1, January 2005.
- [25] R. Puri and K. Ramchandran, "Multiple description source coding using forward error correction codes," in *Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, October 1999.
- [26] F. H. Fitzek, B. Can, R. Prasad, and M. Katz, "Overhead and quality measurements for multiple description coding for video services," in *WPMC*, September 2004.
- [27] C. Huang, J. Li, and K. W. Ross, "Can Internet VoD be profitable?" in *ACM SIGCOMM*, Kyoto, August 2007.
- [28] M. Dischinger, A. Haerberlen, K. P. Gummadi, and S. Saroiu, "Characterizing residential broadband networks," in *ACM IMC*, San Diego, October 2007.
- [29] Z. Liu, Y. Shen, S. Panwar, K. W. Ross, and Y. Wang, "Using layered video to provide incentives in P2P streaming," in *ACM SIGCOMM P2P-TV*, Kyoto, August 2007.
- [30] F. Pianese, D. Perino, J. Keller, and E. W. Biersack, "PULSE: an adaptive, incentive-based, unstructured P2P live streaming system," *IEEE Trans. on Multimedia*, vol. 9, no. 8, December 2007.