

Investigating the Scheduling Sensitivity of P2P Video Streaming: An Experimental Study

Chao Liang[†], Yang Guo[‡], and Yong Liu^{†*}

[†]Polytechnic Institute of NYU, Brooklyn, NY, USA 11201

[‡]Corporate Research, Thomson, Princeton, NJ 08540 USA

Email: cliang@photon.poly.edu, Yang.Guo@thomson.net, yongliu@poly.edu

Abstract—Peer-to-Peer (P2P) technology has recently been employed to deliver large scale video multicast services on the Internet. Considerable efforts have been made by both academia and industry on P2P streaming design. While academia mostly focus on exploring design space to approach the theoretical performance bounds, our recent measurement study on several commercial P2P streaming systems indicates that they are able to deliver good user Quality of Experience with seemingly simple designs. One intriguing question remains: *how elaborate should a good P2P video streaming design be?* Towards answering this question, we developed and implemented several representative P2P streaming designs, ranging from theoretically proved optimal designs to straightforward “naive” designs. Through an extensive comparison study on PlanetLab, we unveil several key factors contributing to the successes of simple P2P streaming designs, including system resource index, server capacity and chunk scheduling rule, peer download buffering and peering degree. We also identify regions where naive designs are inadequate and more elaborate designs can improve things considerably. Our study not only brings us better understandings and more insights into the operation of existing systems, it also sheds lights on the design of future systems that can achieve a good balance between the performance and the complexity.

I. INTRODUCTION

With the fast penetration of broadband residential accesses, Video-over-IP applications are quickly becoming the new generation “killer” applications on the Internet. Traditionally, videos are streamed from a server to clients through unicast connections. Due to the unpredictable packet losses, delays and delay jitters inside the “best-effort” Internet, it is challenging to maintain a continuous video playback on a client. Recently, several ISPs provide IPTV services through the combination of multicast and unicast inside their own private IP networks, such as FiOS of Verizon and U-Verse of AT&T. Content Delivery Networks (CDNs) have also been employed to push video content to the edges of networks, e.g. Youtube [2]. More recently, Peer-to-Peer (P2P) technology has enabled large scale video multicast on the Internet with the minimum server and network infrastructure support [3], [4], [5]. Compared with the traditional client-server based solutions, P2P streaming faces additional challenges, such as random peer arrivals and departures, peer uploading bandwidth bottlenecks [6] and content bottlenecks [7]. Considerable efforts have been made by academia and industry on P2P

streaming design. While academia mostly focus on exploring the design space to approach theoretical performance bounds, e.g. [8], [9], commercial P2P streaming systems are able to deliver good user Quality of Experience(QoE) with seemingly simple designs [10], [11].

One intriguing question remains: *how elaborate should a good P2P video streaming design be?* There is a whole spectrum of answers representing different tradeoffs between the performance and the complexity. Optimal streaming designs from academia reside at one end of the spectrum where the optimal performance is obtained through elaborate designs. Commercial system designs reside at the other end of the spectrum where the protocol simplicity is a high-order design rule, which is often followed at the price of a certain degree of performance loss. In this paper, we address this question through an extensive experimental comparison study of several proposed P2P video streaming designs, ranging from theoretically proved optimal designs to straightforward “naive” designs. Among all the proposed P2P streaming mechanisms, our study unveils several key mechanisms contributing most to the successes of P2P streaming on the Internet, including system resource index, server capacity and chunk scheduling rule, peer download buffering and peering degree. We also identify regions where naive designs are inadequate and more elaborate designs can improve things considerably. Our study not only brings us better understandings and more insights into the operation of existing systems, but also sheds lights on the design of future systems that can achieve a good balance between the performance and the complexity.

A comparison study like this can be done using different approaches. The first approach is modeling and analysis. However, tractable models are often based on unverified assumptions and many system details have to be abstracted out. The second approach is simulation. Due to the difficulty of simulating packet level events in large networks, P2P systems are mostly simulated at the application layer only. Consequently, important details of the underneath native networks, such as topology, congestion, and packet delays, cannot be faithfully represented in the evaluation. The third approach is experiments on the Internet. The ideal places to conduct such experiments are deployed commercial P2P streaming systems. However, many commercial systems are proprietary, and they cannot afford user performance degradations introduced by large scale experiments on their running systems. In this study, we resort to controlled experiments on the PlanetLab [12], a

A preliminary version of this work is to appear in Proc. IEEE ICDCS 2008 [1].

testbed overlaid on the Internet. To the best of our knowledge, it is the first work to conduct specialized comparison study via experiments on real Internet network environment. Our experimental results provide us with valuable quantitative answers regarding the sensitivity and the insensitivity of P2P streaming performance to different design mechanisms. Specifically, our contribution is three-fold:

1) We developed and implemented several representative P2P streaming designs, including our novel adaptive queue-based streaming design [16] (AQCS) and an improved design based on a theoretically proved optimal streaming algorithm (RUPF) from [8]. Those designs were systematically evaluated through extensive experiments on the PlanetLab. Our optimal designs consistently achieve good performance in a wide range of tested scenarios. Compared with naive random pull scheduling, the system with optimal design generally can have shorter chunk delay, higher bandwidth utilization and less chunk loss.

2) Our experiments demonstrated that the performance of P2P streaming is *insensitive* to scheduling when the streaming rate is low and long playback delays are tolerable. Even simple random scheduling can achieve a close-to-optimal performance. We further quantitatively showed that P2P streaming performance becomes *highly sensitive* to scheduling when the streaming rate approaches the maximum supportable rate of the system and the tolerable delays become small.

3) Our rich experimental results along different dimensions deepened our understandings on several fundamental tradeoffs in P2P streaming designs. Firstly, we showed that high streaming rates and short playback delays are achievable through elaborate designs with fine tuned peer management and data scheduling. Several simple and efficient mechanisms allow a design to achieve a good balance between system performance and protocol complexity. Secondly, our study quantitatively demonstrated the importance of server infrastructure and super peers. We showed that increasing server capacity and exploring peer heterogeneity can significantly “bootstrap” the streaming performance on all peers. While the major advantage of P2P streaming is cost-savings on the infrastructure, our results suggest that it is possible to dimension a server infrastructure in a P2P streaming system to achieve the optimal cost-performance tradeoff. Thirdly, buffering is a key mechanism in P2P streaming to ensure smooth playbacks on peers. We demonstrated that elaborate peer buffer management allows one to simultaneously achieve good peer delay performance and playback continuity.

The remaining of this paper is organized as follows. We briefly describe the related work in Section II. Different P2P streaming designs and detailed implementations are introduced in Section III. We then describe our evaluation methodology in Section IV. Major experimental results and key findings are presented in Section V. The paper is concluded with a discussion of the implications on the future P2P streaming designs in Section VI.

II. RELATED WORK

P2P streaming systems can be broadly classified into two categories, namely tree-based and mesh-based. The tree-based

systems, such as ESM [3], have well-organized overlay structures and typically distribute video by actively pushing data from a peer to its children peers. One major drawback of tree-based streaming systems is their vulnerability to peer churn. A peer departure will interrupt video delivery to all peers in the subtree rooted at the departed peer. In a mesh-based P2P streaming system, peers are not confined to a static topology. Instead, the peering relationships are established/terminated dynamically based on the content availability and bandwidth availability on peers. At any given time, peer connects to a subset of peers in the system. Peers periodically exchange information about their data availability. A video chunk is pulled/pushed between a peer and one of its neighbors who has already obtained that chunk. Many mesh-based P2P streaming systems have been proposed, such as Chainsaw [13], DONet/CoolStreaming [4], PRIME [14], etc. Most of the real P2P streaming systems deployed over Internet are based on this type of design ([10], [5]), also including those [15] with network coding mechanism. More recently, there have been ongoing efforts to improve resource utilization and achieve the maximum streaming rate allowed by the system. Random useful packet forwarding [8] is a randomized distributed algorithm that can converge to the maximum streaming rate. Adaptive queue based chunk scheduling [16] is another distributed optimal P2P streaming algorithm. Mutualcast [17] also provides a theoretical optimal solution aided with the peer queue length information. Other related optimal content delivery work can be found in [18][19][20] to discuss the scheduling with network coding and multi-party conferencing application.

Given the large body of literature on mesh based P2P streaming, little work has been done that offers comprehensive comparison study of various schemes. Since it is infeasible to compare a large number of existing schemes in one shot, we select several schemes that are most representative in our view. The comparison study of these schemes over the Internet deepens our understandings of various P2P streaming mechanisms. Our study exposes the fundamental design tradeoffs. It also sheds new lights on the design of future P2P streaming systems.

Previous works may happen to get partial similar observation result through their simulations. The authors in [21] conducted an experimental study on pull-based peer-to-peer streaming system. It was shown that, with proper configuration, simple pull-based scheduling comes close to the optimum in terms of bandwidth utilization and throughput. Our study confirms some of their results. In addition, we examine the system performance from multiple angles, e.g., system resource index, server scheduling rule, buffering/playback delay, peering degree, signaling overhead, etc. Hence our work offers a more comprehensive treatment of the performance comparison study.

III. P2P STREAMING: DESIGN AND IMPLEMENTATION

In this section we describe several P2P scheduling algorithms that will be used later on in performance comparison. We use the term of streaming design and scheduling algorithm

interchangeably in the following. We select these algorithms either due to their provable optimality, or due to their design simplicity and their adoptions in existing P2P streaming solutions. For each scheduling algorithm, we fine tune its design and carefully select the configuration parameters. This ensures individual scheduling algorithm to operate at its peak performance. Our goal is to compare the performance of these representative scheduling algorithms and identify factors that contribute most to the performance of a P2P streaming system. The performance comparison also allows us to answer fundamental questions such as whether scheduling is indeed important in P2P streaming.

We start with introducing a layered model of P2P streaming. So far, P2P streaming designs have taken a holistic approach. To facilitate the comparison of different scheduling algorithms and pinpoint the root cause of performance difference, we propose a layered model that dissects a P2P streaming design into four basic components forming a layered stack.

A. Layered Model of P2P Streaming

P2P streaming provides TV-like live streaming service to users. In client-server service model, a live video stream is transmitted from a server to clients in the same order as it is created. Hence a small buffer is required at client end to avoid buffer underflow due to bandwidth variations inside networks. In P2P live streaming, a live stream is divided into equal size chunks. The server does not upload all chunks to all clients/peers. Peers also need to retrieve missing chunks from each other. Since the chunks obtained from the server and multiple peers may arrive out of order, a download window is required for a peer to assemble out-of-order chunks. A large download window increases the opportunities of chunk swapping among peers. However it imposes long startup delays. Different P2P scheduling algorithms employ different strategies in dimensioning download window and diffusing chunks among peers.

A typical P2P streaming design can be dissected into four basic components: mesh construction, peer selection, data exchange, and download window control. These four components form a layered architecture where each layer interfaces with its upper/lower layers, and communicates with corresponding layers at neighboring peers (Fig. 1). In the following we use the term of layer and component interchangeably.

Upon arrival, a peer connects to multiple peers already in the system via the mesh construction component running at different peers. A P2P mesh is constructed and maintained by the mesh construction components (probably with the help of directory services provided by peer trackers). The neighboring peer list is passed to the upper layer, peer selection component. Meanwhile, peers periodically exchange available chunk information with neighboring peers using buffer maps. Each bit in a buffer map indicates the availability of one chunk, a single bit of '1' indicates the chunk is available and '0' indicates a chunk is missing.

The peer selection layer selects a subset of neighbors to whom the local peer serves data. The other neighbors are snubbed. The selection process is executed periodically, or is

invoked by pre-defined events. The neighbor selection criteria, the epoch of period, and the events that invoke the peer selection process are design parameters that vary from systems to systems. If required, the peer selection component also collects the list of neighbors that have signaled the willingness to serve the data, and passes the list to the upper layer, the data exchange layer.

The data exchange layer works in either push mode or pull mode. In push mode, once the underlying peer selection component selects the neighboring peers to serve, the data exchange component determines which specific data chunks to push to the selected peers. In pull mode, the peer selection component informs the data exchange component the peers who are willing to serve. The data exchange component then computes the set of data chunks for each peer, and requests to pull in the chunks from the corresponding peers.

In order to support TV-like live streaming, a moving download window is required. At any time, a peer only downloads data chunks falling in the download window. The download window continuously moves forward as the content in the buffer is played out. Intuitively, a large window helps to diversify the content among peers and better utilize peers' uplink bandwidth. However, large window size leads to long playback delays, which is not desirable to users.

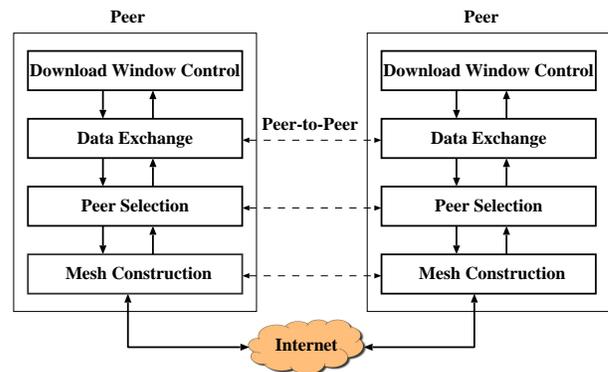


Fig. 1. Layered model of P2P streaming

To comparatively study the fundamental design tradeoffs, we select several representative P2P streaming schemes from a pool of existing solutions. Next we describe three basic scheduling algorithms: Random Pull (RP), Adaptive Queue Based Chunk scheduling (AQCS), and Random Useful Packet Forwarding (RUPF). AQCS and RUPF have been theoretically proved to be optimum [16], [8]. RP scheduling is a simple scheduling and bears similarity to the scheduling used in real P2P streaming solution [10]. Two variations of RP, Random Pull with Fresh chunk first (RPF) and Random Pull with Most-Deprived peer selection (RPMD), are also introduced for comparison.

B. Baseline Scheduling : Random Pull Scheduling

As a baseline scheduling, Random-Pull (RP) works randomly in each layer. The peers form a random mesh at mesh construction layer. A peer in RP scheduling randomly select M peers to serve at peer selection layer. The value of M is

chosen to be proportional to the peer's upload capacity. RP employs pull mode at the data exchange layer. Pull mode is widely employed in P2P streaming since it avoids duplicate data chunks. A peer can request up to K data chunks from one serving peer. The value of K is set so that K chunks can be downloaded within *one window sliding period*, the epoch at which download window moves forward. At the data exchange layer, RP randomly selects missing chunks to pull. Figure 2 illustrates a simple example of the above process.

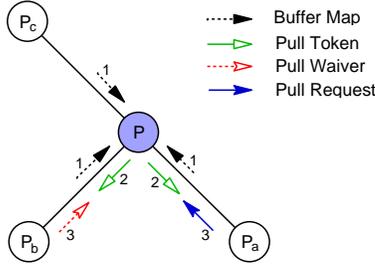


Fig. 2. Random pull scheduling

Peer P randomly selects peer P_a and P_b to serve by issuing pull tokens to them. Peer P_a then sends a pull request back to P for some missing chunks which P owns. P_b sends a pull waiver message instead since all of its missing chunks have been scheduled to be pulled from other peers. After receiving the waiver message, P can send the token to another unselected peer. The value K is selected by peer P according to its uplink capacity, and is piggybacked on the pull tokens.

C. Optimal Scheduling I: Adaptive Queue-based Chunk Scheduling

Adaptive queue-based chunk scheduling (AQCS) [16] is a deterministic distributed scheduling algorithm. AQCS can support the optimal streaming rate in practical networking environment. The mesh construction layer in AQCS maintains a fully connected mesh among participating peers. Data chunks are pulled/pushed from the server to all peers, cached at peers' forwarding queues, and relayed from peers to their neighbors. The server's peer selection layer determines the peers that issue the pull signals. The received chunks at peers are relayed to all other peers hence peers' peer selection layer in AQCS is not required. Server's data exchange layer always pushes out the fresh chunks that have not been served before, while peers' data exchange layer conducts chunk relaying.

The availability of spare upload capacity on a peer is inferred from the occupancy of its forwarding queue. Signals are passed between the server and peers to adaptively allocate chunk forwarding workloads to peers based on their current bandwidth availabilities. Figure 3 illustrates a simple example with one source server and three peers. Each peer maintains several queues including a forwarding queue. Using peer a as an example, the signal and data flow is described next. A pull signal are sent from peers a to the server whenever its forwarding queue becomes empty (or falls below a threshold) (step 1 in Figure 1). The server responds to a pull signal by sending three data chunks back to peer a (step 2). These

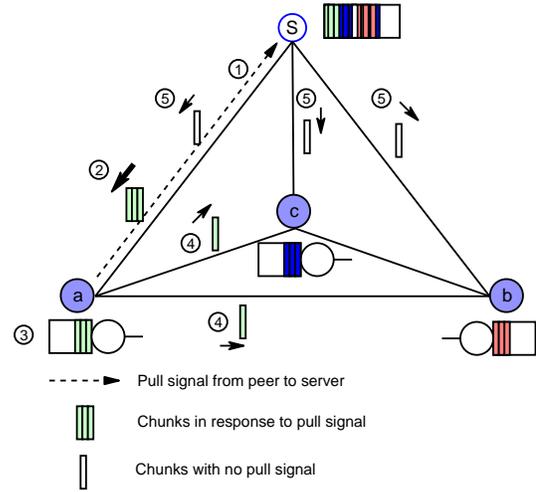


Fig. 3. Queue-based P2P system with four nodes. 1. peer a sends pull signal to the server; 2. the server replies with three chunks in response to the pull signal; 3. three chunks are cached in the peer a 's forwarding queue; 4. cached chunks are forwarded to neighbor peers sequentially; 5. duplicate chunks are sent to all peers when the server has finished all 'pull' signals processing.

chunks will be stored in the forwarding queue (step 3) and be relayed to peer b and peer c (step 4). When the server has responded to all 'pull' signals in its 'pull' signal queue, it uploads one data chunk to all peers (step 5). This data chunk will not be stored in peers' forwarding queue and will not be relayed further.

The pull signal issued by a peer is used by the server as a realtime feedback on the peer's available upload bandwidth. A peer with higher upload bandwidth can forward chunks more quickly and drain out its forward queue faster. Consequently, it will send more pull signals to the server. Using the feedback of peer bandwidth, the server can either reply existing pull signal with chunks for forwarding, or broadcast chunks for all peers to fully utilize its bandwidth when it receives no pull signal from any peer. It is this adaptiveness that enables AQCS to achieve high utilization of peer and server bandwidth. The basic AQCS design requires that peers are fully connected. To extend the system scalability, a hierarchically clustered AQCS framework is proposed in [22] by groups peers into clusters and implement AQCS scheduling inside each cluster to achieve the global optimal performance. In this paper, for the comparison purpose, we use the basic AQCS design to obtain the insights of fundamental design tradeoffs. Various conclusions can be easily extended to hierarchical AQCS.

D. Optimal Scheduling II: Random Useful Packet Forwarding

Random Useful Packet Forwarding (RUPF) [8] is a randomized distributed algorithm that can converge to the optimal streaming rate. The optimality of the method is proved under the assumption that the mesh is fully connected. However the scheduling can be extended to arbitrary mesh networks. The key in RUPF is for peers to identify the neighboring peer that needs data the most, namely the *most-deprived peer*. Peers serve data chunks to their respective most-deprived peers. The served data chunks are randomly selected from the missing chunk set of the receiving peer. The source server in RUPF

TABLE I
LAYERED MODEL COMPARISON OF P2P SCHEDULING ALGORITHMS

	RP(Peer)	AQCS(Peer)	RUPF(Peer)	RP(Server)	AQCS(Server)	RUPF(Server)
Mesh-construction	Any	Full mesh	Any	N/A	N/A	N/A
Peer-selection	Random	All peers	Most deprived	Random	Pull signal	Most deprived
Data-exchange	Random useful	Queued chunks	Random useful	Random useful	Fresh chunk first	Fresh chunk first

estimates the average upload and download rate every ten seconds. The information is recorded at local log files and is transmitted to one public sink after the experiment is over. The experiment results are analyzed off-line.

B. Performance Metrics

The following performance metrics are investigated to evaluate the performance of peer-to-peer streaming systems.

- 1) *chunk miss ratio*: During the video playback, chunks have to be received before their playback deadline. The chunks missing the deadlines are skipped or the video playback freezes. Both downgrade users' viewing quality. *Chunk miss ratio* is defined as the number of data chunks which miss the playback deadline over the total number of data chunks that a peer should receive. Since chunks are of the same size, this term directly reflects the playback quality at user end. *Average chunk miss ratio* is the average of chunk miss ratios across all peers. It takes time for the system to reach its steady state. In our off-line analysis, the first 60 seconds data is skipped to minimize the impact of transient system behavior.
- 2) *peer bandwidth utilization*: The bandwidth utilization reflects how effective scheduling algorithms utilize the bandwidth resources. A poor scheduling scheme may under-utilize a peer's upload capacity while other peers are "starved" of content, the so-called *content bottleneck* phenomenon. The bandwidth utilization is computed by estimating the average upload rate and dividing it by the peers' pre-set upload capacity.
- 3) *playback delay*: Playback delay is the elapsed time from a chunk is generated at the source to the moment it is played at the peer/user.
- 4) *signaling overhead*: The mesh-based P2P streaming requires peers/source server to exchange information, such as buffer map and peer management info, periodically. A well designed scheduling should incur small signaling overhead that does not affect the overall performance.

In the following experiments, we study the impact of different design parameters by examining these performance metrics.

V. EXPERIMENTAL RESULTS

In this section, we present experiment results obtained in our comparison study of five different P2P streaming designs, Random Pull (RP), Random Pull with Fresh first policy (RPF), Random Pull with Most-Deprived peer selection (RPM), Random Useful Packet Forwarding (RUPF) and Adaptive Queue-based Chunk Scheduling (AQCS).

A. Impact of System Resource Index

As shown in [6], the maximum supportable streaming rate in a P2P streaming system with n peers is

$$r_{max} = \min \left\{ u_s, \frac{u_s + \sum_{i=1}^n u_i}{n} \right\} \quad (1)$$

where u_s is the uploading capacity of the server and u_i is the uploading capacity of peer i . Denote by \bar{u} the average upload capacity per peer, i.e., $\bar{u} = \frac{u_s + \sum_{i=1}^n u_i}{n}$. If $u_s \leq \bar{u}$, $r_{max} = u_s$. The video rate is bounded by the server uploading rate u_s . We call such scenario *server resource poor scenario*. In contrast, if $u_s > \bar{u}$, $r_{max} = \bar{u}$. The server is no longer the performance bottleneck. We call it *server resource rich scenario*. The bandwidth resource index ρ is defined as

$$\rho \triangleq \frac{\bar{u}}{r} = \frac{u_s + \sum_{i=1}^n u_i}{nr}, \quad (2)$$

which is the ratio of the total uploading resources in the system and the total resource demand at streaming rate of r .

Below we compare the performance of scheduling algorithms at different resource indexes by varying the streaming rate. In AQCS, the chunks size is set to be 1KByte, and the server replies each pull signal with only one chunk. We experiment with other parameters and the current setting gives us the best performance. The server increases the streaming rate by increasing the number of chunks generated per second. In RUPF and RP, we fix the buffer map size at 120bits (15 bytes). The download window is set to be 30 seconds and moves forward every 10 seconds. The server produces four chunks per second and increases the streaming rate by increasing the chunk size (this way the buffer map size remains the same). We select 40 stable PlanetLab nodes for this experiment. Using Trickle, we set the uplink bandwidth according to the distribution as stated in Table II. A full mesh is formed among nodes. The content source server's uploading capacity is set at 1 Mbps. We gradually increase the streaming rate from 480 kbps to 960 kbps. The resource index of the system decreases correspondingly from 2.2 to 1.1. For each streaming rate, we conduct one set of experiments for each scheduling algorithm. Each experiment lasts for 300 seconds. Based on our off-line data analysis, the experiment duration of 300 seconds is appropriate since the systems goes into the steady state within tens of seconds. Figure 5(a) shows the average chunk miss ratios of the various scheduling algorithms. When the streaming rate is 480kbps ($\rho = 2.2$) or smaller, all scheduling methods have zero chunk miss ratio. This indicates that when the system has high resource index, the chunk miss ratio is insensitive to scheduling algorithms. Even random scheduling can have very good performance. Indeed, most current commercial P2P streaming systems on the Internet operate at streaming rates of 400kbps or lower.

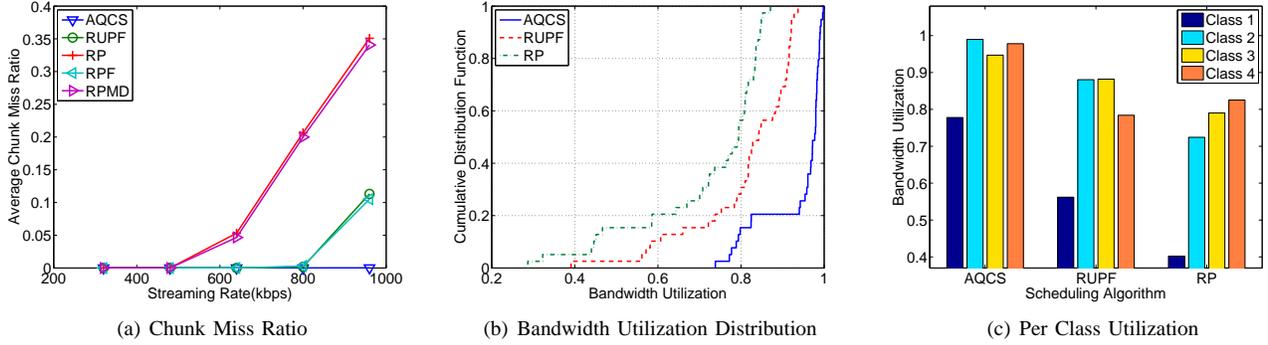


Fig. 5. Streaming performance of different scheduling algorithms at different resource indexes when server capacity is 1Mbps. (a) Chunk miss ratio at different streaming rates. The performance is insensitive to scheduling at high resource indexes, and becomes sensitive to scheduling at low resource indexes; (b) Peer bandwidth utilization distribution at low resource index $\rho = 1.1$; (c) Bandwidth utilization for peers with different upload capacity. At low resource indexes, scheduling with higher peer bandwidth utilization achieves better performance.

Our experimental results explain the successes of simple P2P streaming designs.

As we increase the streaming rate above 480kbps, the performance of RP and RPMD drops sharply. Meanwhile, with more elaborate designs, RPF and RUPF still perform satisfactorily until the streaming rate approaches 960kbps. At this point, AQCS outperforms all others and sustains zero chunk miss ratio. This suggests that intelligent scheduling does make difference when the resource index is low.

To understand the root cause of the performance difference among different scheduling algorithms, we look into peers' uploading bandwidth utilization when the streaming rate is 960kbps ($\rho = 1.1$). At this streaming rate, RP and RUPF has around 35% and 11% chunk miss ratio, respectively, while AQCS has zero loss. Figure 5(b) shows the cumulative distribution of peer average bandwidth utilization during the experiment. Peers in AQCS have average bandwidth utilization around 94%, while that in RUPF and RP is around 80% and 72%, respectively. Moreover, the peers of different classes (see Table II) have different uploading bandwidth. Figure 5(c) depicts the average bandwidth utilization of each peer class. All peer classes in AQCS achieve the highest bandwidth utilization. The bandwidth utilization of peer classes in RUPF is better than RP, except for the class four peers with the lowest bandwidth. When the system resource index is close to one, it becomes critical to utilize all peers' uploading bandwidth. In this experiment, RUPF and RP encounter content bottleneck [7] with small resource index. Content bottleneck wastes bandwidth since the peers cannot obtain new chunks in time and have nothing to upload to other peers.

B. Impact of Server Scheduling Rule and Capacity

It is interesting to notice the performance difference between random pull (RP) and random pull with fresh chunk first (RPF) as shown in Figure 5(a). RPF scheduling is a variation of RP and implements fresh-chunk-first scheduling at the server. The server gives strict priority to push out fresh chunks so that new content can be quickly distributed among peers. RPF out-performs RP significantly and its performance is very close to the performance of RUPF. In contrast, random pull with most deprived peer selection (RPM) implements most

deprived peer selection as in random useful packet forwarding (RUPF). The chunk miss ratio in RPMD is almost as bad as in RP. The experiment results suggest that (i) the fresh-chunk-first scheduling at source server plays an important role in improving the system performance; and (ii) most deprived scheduling, although has been theoretically shown to be optimal [8], does not seem to bring performance improvement in our experiments.

Next we investigate the impact of server bandwidth. In the previous experiment, the server uploading bandwidth is 1Mbps. Since the average peer uploading capacity in the system is slightly above 1Mbps, it corresponds to the server resource poor scenario. In the following experiments, we increase the server bandwidth to 3.2Mbps, leading to the server resource rich scenario. We repeat the previous experiment with a wider range of streaming rates and plot the results for AQCS, RUPF and RP in Figure 6(a). (We skip RPMD and RPF because their performance lies in between that of RP and RUPF). The curves show that all miss ratios are nearly zero, even for RP, when the streaming rate is less than 960kbps. The chunk miss ratio goes up rapidly for RUPF and RP when streaming rate goes beyond 960kbps while AQCS still maintains zero miss ratio up to 1,100kbps. The miss ratio for AQCS increases linearly at the end, since the system resource index drops below one when the streaming rate is larger than 1,100kbps. This again verifies that *the streaming performance becomes insensitive to scheduling when the system resource index is high, and sensitive to scheduling when resource index is low*.

More interestingly, increasing server capacity from 1Mbps to 3.2Mbps extends the scheduling insensitive region dramatically from 480kbps ($\rho = 2.2$) in Figure 5(a) to 960kbps ($\rho = 1.155$) in Figure 6(a). Obviously, increasing server capacity can increase the resource index of the system. And system performance will improve as resource index increases. However, according to equation (2), the increase in server bandwidth will be factored down by the number of peers in the system. In fact, in this experiment, the resource index at 960kbps is 1.155 which is only slightly higher than the resource index of 1.1 at 960kbps in the previous experiment. The chunk miss ratios for RP and RUPF improve from 35%

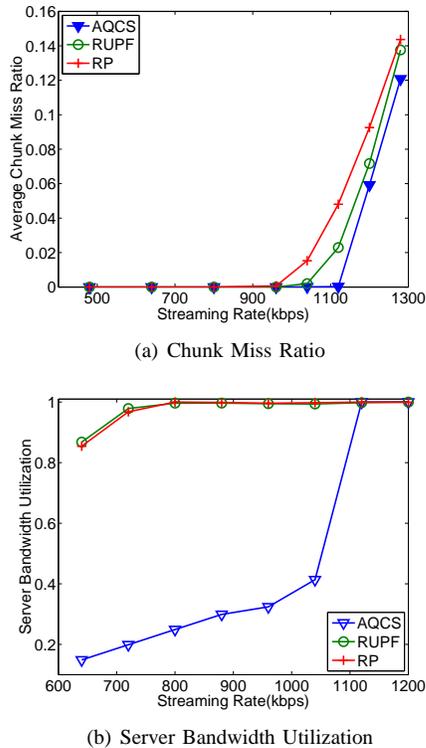


Fig. 6. Streaming performance of different scheduling algorithms at different resource indexes when server capacity is 3.2Mbps. (a) Increasing server bandwidth dramatically extends the scheduling insensitive region; (b) More efficient P2P scheduling impose lighter load on server.

and 11% to zero respectively. This suggests that increasing server uploading capacity brings in more dramatic performance improvement than barely increasing system resource index. Figure 6(b) illustrates the average server bandwidth utilization under different streaming rates. The server utilizations in RUPF and RP are always high. This suggests that the server bandwidth plays an important role in reducing the chunk miss ratio for these two scheduling algorithms. On the contrary, the server utilization in AQCS is low until the streaming rate exceeds 1Mbps and the resource index falls below 1. This suggests that an efficient P2P streaming design can effectively bring down the server load and improve the system scalability.

To further examine the impact of server bandwidth, we conducted additional experiments by continuously varying server upload capacity. We fix the streaming rate at 640kbps and increase the server bandwidth from 600kbps to 1.2Mbps. All other experiment setting is the same as the previous experiments. Figure 7 shows the chunk miss ratio for five different scheduling algorithms at different server bandwidth. When the server bandwidth is low, there are performance gaps among them. Again, algorithms with fresh-chunk-first rule have better performance. As the server bandwidth approaches 1.2Mbps (twice the streaming rate), the performance becomes insensitive to scheduling and all of them have nearly zero chunk miss ratios. This again shows the unique impact of the server bandwidth on the whole system. One explanation is that a server with high bandwidth can simultaneously upload a chunk to many peers. This effectively increases the fan-

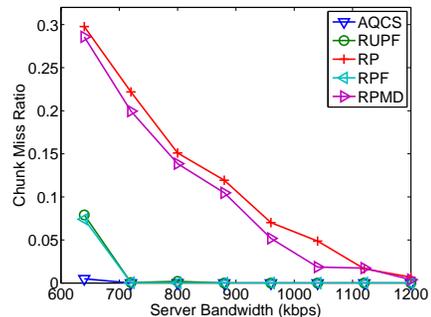


Fig. 7. Chunk miss ratio under different server bandwidth. Increasing the server bandwidth can quickly improve the system performance.

out degrees of chunk delivery trees at their roots thus chunks can be delivered to all peers quicker. The same amount of bandwidth increase on a regular peer does not have such significant impact. Our results here suggest that *investing on server bandwidth can dramatically bootstrap the performance of P2P streaming systems*, which is consistent with findings from other analysis [9] and experimental studies [21].

C. Signaling Overhead

Signaling overhead has been a concern for mesh-based P2P streaming systems due to the requirement of frequent buffer map exchange among peers. Figure 8(a) shows the system overhead of these systems under different streaming rates, calculated as the ratio of the signaling traffic volume and the video data traffic volume. AQCS still outperforms other scheduling algorithms due to its simple but effective signaling strategy. At different streaming rates, the overhead ratio is nearly constant because the server always replies with a fixed number of chunks to each pull signal. RUPF and RF have larger system overhead than AQCS. However, the overhead is less than 4%. As the streaming rate increases, the chunk size becomes larger, which leads to slight decrease of signaling overhead.

Although the absolute traffic volume of signaling in RUPF and RF is not large compared to the traffic volume of video chunk, the number of signaling packets is large due to its small size. Figure 8(b) shows the ratio of the number of signaling packets over data chunks. The number of signaling packets in RUPF and RP is comparable to the number of data packets. In practical system, typically the network bandwidth is the bottleneck and the user machines have sufficient processing power and memory, results show that the signaling overhead is not a big issue for a well designed P2P streaming system.

D. Impact of Buffering Delay

In client-server based video streaming, client side video buffering is necessary for continuous playback in face of network bandwidth variations. In P2P streaming, each peer maintains a moving window that specifies the range of video chunks to be downloaded. The window normally advances at the video playback rate. The window size determines the length of playback delay. The larger window size gives peers

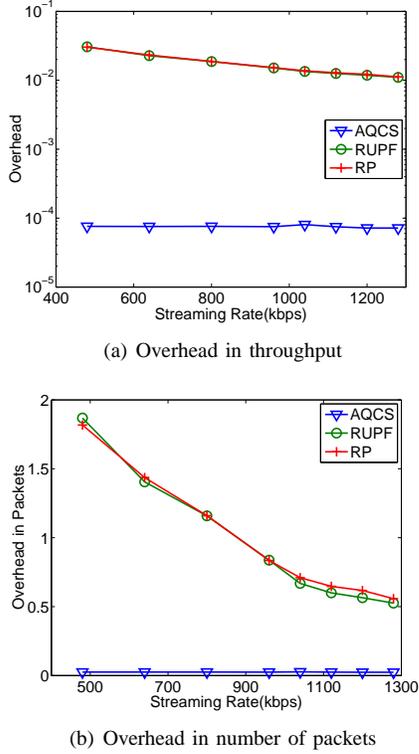


Fig. 8. Signaling overhead. Efficient scheduling leads to low system overhead.

more time to download chunks. However, the larger window size imposes longer playback delay. The dimensioning of buffering delay is indeed a trade-off between the streaming delay performance and playback continuity. In this section, we investigate the impact of buffering delay on the performance of different P2P streaming designs.

We first study the video chunk delays under different streaming rates. We measure the chunk delays in the server resource rich experiments presented in Section V-B. For each chunk, the server records its generation time, and peers record the chunk's arrival time. In RUPF, the download window is 30 seconds and moves forward every 10 seconds. In AQCS, the download window is 15 seconds and moves forward every 1 second. Figure 9(a) shows the arrival times of chunks on one peer and their generation times at server side when the streaming rate is 480kbps in RUPF. With this low streaming rate, the system has high resource index. We can observe from Figure 9(a) that chunks arrive at the monitored peer in batches. The moving window advances once every 10 seconds. After each window advance, the peer can quickly finish the downloading of all missing chunks entering into the moving window. The peer stays idle to wait for the next window advance. The fast downloading at low streaming rate enables us to reduce the download window size to achieve shorter playback delays on all peers.

We then compare the delay performance of RUPF and AQCS at the streaming rate of 1120kbps. From Figure 6(a), the chunk miss ratio in RUPF is around 3%. For chunks downloaded before their playback deadlines, we plot their delay distribution in Figure 9(b). The chunk delays in RUPF

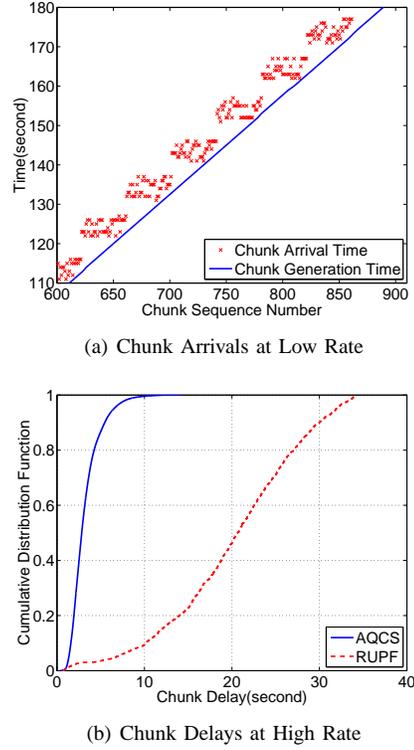
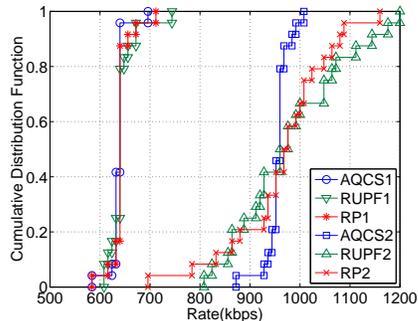


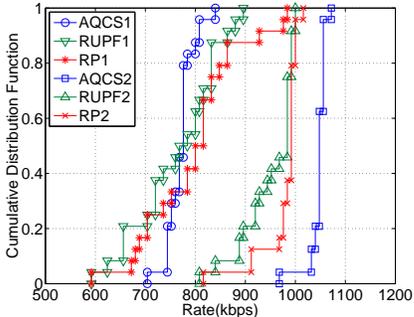
Fig. 9. Delay performance under different streaming rates. (a) Chunk arrival patterns in RUPF at rate 480kbps, high resource index leads to fast chunk downloading. (b) Chunk delay performance comparison between AQCS and RUPF at rate 1120kbps, low resource index leads to high delay variability in RUPF.

have a wide variability with the average delay of 20.2 seconds. In comparison, AQCS not only achieves zero chunk miss ratio, but also has superb delay performance. All chunks can be received within 10 seconds after they are generated by the server. The average chunk delay is 3.2 seconds. This shows us that *efficient scheduling can help reduce the video delays and require less buffering at peers.*

Next we examine the root cause of the delay difference in different scheduling algorithms. In P2P streaming, a peer downloads chunks either from the server or from other peers. Therefore, the variability in the uploading rates and downloading rates on all peers collectively determine peer delay performance. To verify this, we conduct two sets of experiments at different streaming rates. For each set of experiments, we keep track of peers' average uploading and downloading rates every 10 seconds. Fig 10 presents the receiving and sending rate distribution of one class 2 peer at two streaming rates. The digit '1' and '2' in the legends represent results for the streaming rate 640kbps and 960kbps, respectively. From Figure 10(a), we see that when the streaming rate is low, for all three scheduling algorithms, the peer's download rates are stable and tightly centered around the streaming rate. This indicates that the monitored peer can have a smooth video playback. However, when the streaming rate is high, while AQCS is still pretty stable, RUPF and RP incur lots of oscillations and a long buffering is necessary to ensure smooth playback. Figure 10(b) shows the distribution of uploading



(a) Download Rate Distribution



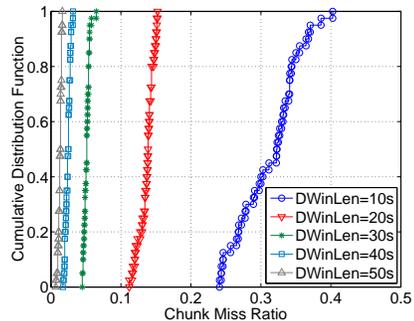
(b) Upload Rate Distribution

Fig. 10. Downloading and uploading rates of one peer under different scheduling algorithms. AQCS has stable peer downloading and uploading. Long buffering is necessary to smooth out the uploading and downloading fluctuations in RUPF and RP.

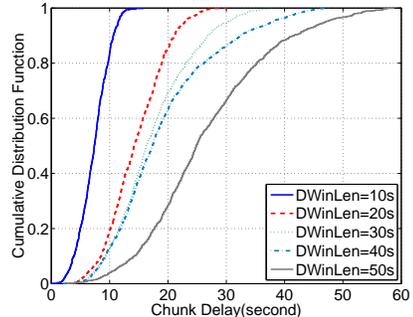
rates on the same peer. Again, AQCS has stable uploading at low and high streaming rates. Uploading rates in RUPF and RP oscillate significantly.

To reduce chunk miss ratio under uploading and downloading rate oscillations, one can introduce a large buffering delay on peers. We examine the performance of RP with various buffer lengths. The streaming rate and the server bandwidth are fixed to be 640kbps and 1Mbps. We then vary the download window size of all peers from 10 seconds to 50 seconds.

Figure 11(a) shows the chunk miss ratio distribution of all nodes with different buffer lengths. As the buffer length increases, the chunk miss ratio decreases. With 10 seconds buffer length, all peers encounter larger than 23% chunk loss. The average chunk miss loss drops to 5% when the buffer length increases to 30 seconds. As the buffer length continues to increase, the improvement becomes less because fewer and fewer chunks miss the playback deadline. To examine the impact of buffer length on the delay performance, we plot the delay distribution of chunks received by one specific peer with different buffer lengths in Figure 11(b). As the buffer length increases, the average and worst chunk delay also become larger. The results confirm that the chunk miss ratio can be reduced by increasing peer buffering, although that may introduce longer delays. From this experiment of RP, we see that *while random scheduling can achieve zero chunk miss ratio with long buffering delays, a careful scheduling design is necessary to simultaneously achieve low chunk miss ratio and low chunk delay.*



(a) Chunk Miss Ratio



(b) Chunk Delay

Fig. 11. Chunk miss ratio and delay trade-offs under different buffer lengths. (a) chunk miss ratio decreases as buffer length increases. (b) chunk delay increases as buffer length increases.

E. Impact of Peering Topology and Degree

The peering degree is another important parameter in P2P streaming designs. In a random mesh based streaming system, a peer only connects to a subset of randomly selected neighbors and exchanges data with them. Due to the connection and signaling overhead, the number of neighbors of each peer cannot be too large. In this section, we compare the performance of two random mesh based designs, RUPF and RP, with different peering degrees. We use 100 Planetlab nodes for the evaluation. The streaming rate is 640 kbps and the server bandwidth is 1 Mbps.

We start with homogeneous random peering topology where each peer has the same number of randomly selected neighbors. First we compare the streaming performance of RUPF scheduling algorithm at different peering degrees. Figure 12(a) shows the cumulative distribution of all peer chunk miss ratio under different peering degrees. As the peering degree increases, the chunk miss ratio drops. When each peer only connects to 6 neighbors, the average chunk miss ratio is around 20%. Chunk miss ratio drops to 5% when peering degree increases to 10. At peering degree 14, the chunk miss ratio drops to 3%. This is because that with more neighbors a peer has a better chance to locate and download missing chunks in its download window. We can also observe that as the peer degree increases, the improvement slows down. The miss ratio improvement from peering degree 14 to 18 is less than the improvement from 6 to 10. Figure 12(b) shows the distribution of the chunk miss ratios with the homogenous setting in RP algorithm. Similar trends can be observed as that in RUPF.

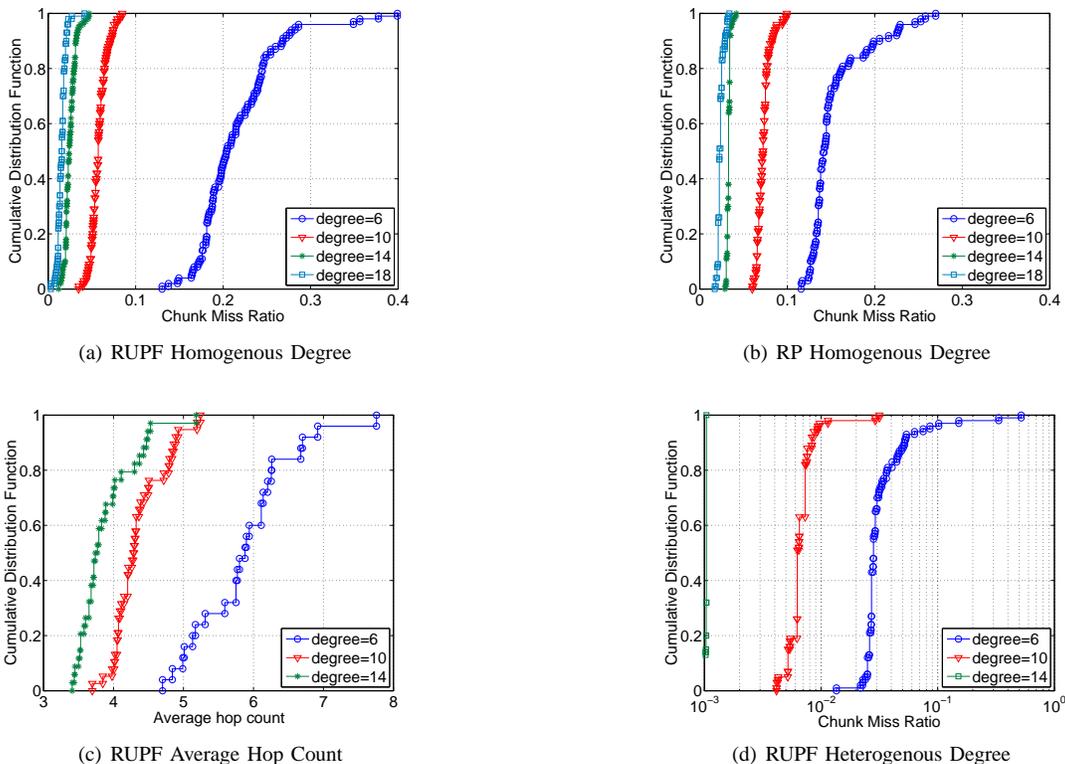


Fig. 12. System performance with different peering topologies and degrees. a), b) chunk missing ratio decreases as peering degree increases for homogenous RUPF and RP; c) Increasing peering degree leads to shorter chunk delays; d) Heterogeneous peering topology has better performance.

We can conclude that *increasing peering degree can reduce chunk miss ratio on peers*.

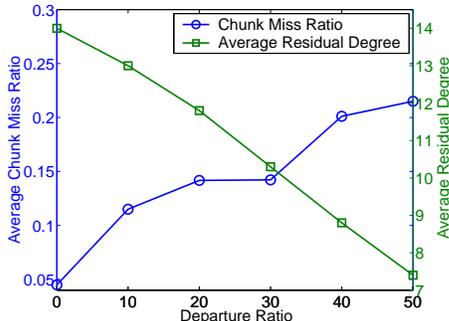
Even in mesh-based P2P streaming systems, the delivery paths of a single chunk still form a tree. As peer degree increases, a peer can upload one chunk to more neighboring peers, which in turn increases the fan-out of the chunk delivery tree. Therefore, the average number of hops that a chunk needs to traverse to reach all peers decreases. Consequently, peers can download chunks faster and the chunk miss ratio becomes smaller. To verify the path length of each chunk, we append a hop counter to each chunk, which records how many hops a chunk has traversed. Upon receiving a chunk, each peer increments the chunk's hop counter and forward it to its descendants in the chunk delivery tree. Figure 12(c) depicts the distribution of the average hop count of randomly sampled chunks received by all peers. When all peers connect to 6 neighbors, each chunk needs to traverse in average 6 hops from server. As peer degree increases to 14, the average path length drops down to below 4.

Next we examine the impact of peering degree in heterogeneous random topology. Swaplink [24] can build heterogeneous random graph with different node degrees based on their upload capacity. We implemented the Swaplink topology construction algorithm and incorporated it into our peering management protocol in the experiments. Since peers with high bandwidth capacity are capable of serving more neighbors, we set the number of neighbors of a peer proportional to its uploading bandwidth. The degrees of peers in class 1 are the same as those in class 2, and the degrees of peers in class 3

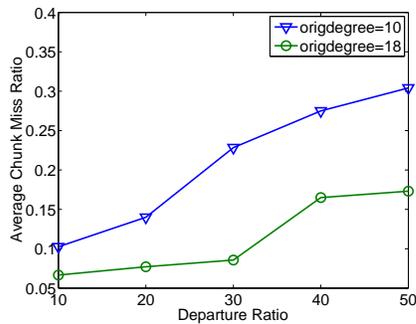
are the same as those in class 4. The degrees of peers in class 1 and class 2 peers are twice the degrees of class 3 and class 4. Figure 12(d) illustrates the chunk miss ratio under different average peering degrees. Compared with homogeneous results in Figure 12(a), the performance in the heterogeneous random topology is much better than that of the homogenous topology with the same average degree. With the average degree of 6, the average chunk miss ratio is less than 5%, while the average miss ratio in the homogeneous case is 20%. This suggests that peers with high bandwidth should be assigned with more neighbors in order to improve the whole system performance. Since the existence of super peers can dramatically improve performance, more considerations should be given to them during the system design.

Other than achieving good streaming performance, such as small chunk missing ratio and low chunk delay, how to make P2P streaming systems robust against peer churn is another major design consideration. In the rest of this section, we investigate the impact of peering degree on the resilience of P2P streaming systems. In this experiment, all 100 peers join the system at the beginning and each peer has homogenous number of neighbors. After the system enters its steady state, we create peer churn events by removing a certain percentage of peers from the system in a batch. To evaluate the resilience of random mesh topologies at different degrees, we sample the performance of the system after each batch peer removal without implementing any churn recovery mechanisms. In practice, a peer losing neighbors can obtain more neighbors through directory service such as tracker. The performance

shown here is hence the worst case scenario and represents the performance peers may experience during the transition period. After each batch peer removal, a random set of peers in the system are sampled to evaluate their playback performance. Figure 13(a) shows the chunk miss ratio and the average



(a) Miss ratio & Residual degree



(b) Different initial degree

Fig. 13. Robustness against peer churn at different peering degrees. a) chunk miss ratio and peer connectivity after batch peer departures; b) higher peering degree makes the system more resilient to peer churn.

residual peering degree of the sampled peers after different percentages of peers leave the system. At the beginning of the experiment, every peer connects to 14 neighbors to construct the mesh network. When 10% peers leave, the average chunk miss ratio on the sampled peers is around 12% and the average number of connections of each peer drops to around 13. While 40% peers leave, the chunk miss ratio increases to 20% and the average residual peer number drops below 9. As larger percentage of peers leave, the remaining peers have less neighbors and the performance becomes worse.

Large peering degree helps in improving system robustness in the face of peer churn. We compare the resilience of the system when the initial peering degree is set to 10 and 18, respectively. Figure 13(b) illustrates the chunk miss ratios. The system with initial peer degree 18 has better resilience than others. Even when 50% peers leave, it can still has around 15% chunk loss, while the ratio of that with initial peer degree 10 is nearly twice as much.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we conducted an extensive comparison study for five representative designs. Rich numerical results obtained from experiments on the PlanetLab allow us to quantitatively investigate the sensitivity and insensitivity of P2P streaming

performance to different design mechanisms. Our study unveils several key factors contributing most to the successes of simple P2P streaming designs on the current Internet. More importantly, our study sheds lights on the design of the future generation of P2P streaming systems that achieve a good balance between the performance and the complexity.

Future work will be conducted in the following directions. Firstly, due to the size of the current PlanetLab and the heavy usage from the research community, the scale of our experiments is confined to around one hundred nodes. Nevertheless, many findings from our current experiments are consistent with results obtained from measurement studies of various commercial systems with thousands of real users. P2P streaming systems are designed to be scalable. We are interested in exploring different ways to extrapolate our results to draw more general conclusions on P2P systems with different sizes. To address the difficulty of large scale experiments on the Internet, we plan to test the scalability of our findings by incorporating simulations and analysis.

Secondly, through experiments, we identified several fundamental P2P streaming design tradeoffs deserving further investigations. We showed that a server infrastructure can bootstrap the performance on all peers. We will systematically study the server infrastructure provisioning problem to achieve the optimal cost-performance tradeoff in a P2P streaming design. Buffering is a key mechanism in P2P streaming to ensure smooth playbacks on peers. Buffer dimensioning is a tradeoff between peers' delay performance and playback continuity. We will develop buffer dimensioning guidelines that take into account peer behaviors and network characteristics. In this work, we mostly focused on the streaming performance. The impact of peer churn on system performance is an important research problem. It deserves more attentions/efforts beyond the results presented in this paper. It is an ongoing effort to investigate this deeper with the assistance of model analysis and trace-driven simulations. Another important consideration in practical systems is robustness and diagnosability. The tradeoff between the protocol complexity and the system's robustness and diagnosability deserves further study.

Lastly, one major advantage of Video-over-IP is its flexibility and efficiency in providing Video-on-Demand (VoD) services. Several P2P commercial systems, such as PPLive and PPstream, have started to offer VoD services with reasonably good user quality of experience. Compared with P2P live streaming, P2P VoD has its own design advantages, e.g. prefetching and longer tolerable delays, and challenges, e.g. less concurrent peers. We will extend our current prototype system to support P2P VoD service. Experimental study will be conducted to test the applicability of live streaming results to VoD, and obtain new results unique to P2P VoD designs.

REFERENCES

- [1] C. Liang, Y. Guo, and Y. Liu, "Is Random Scheduling Sufficient in P2P Video Streaming?" in *Proceedings of International Conference on Distributed Computing Systems (ICDCS)*, 2008.
- [2] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User Generated Content Video System," in *Proceedings of Internet Measurement Conference*, 2007.

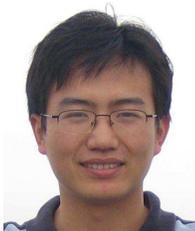
- [3] Y.-H. Chu, S. G.Rao, and H. Zhang, "A case for end system multicast," in *Proceedings of ACM SIGMETRICS*, 2000.
- [4] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "DONet/CoolStreaming: A data-driven overlay network for live media streaming," in *Proceedings of IEEE INFOCOM*, 2005.
- [5] C. Wu, B. Li, and S. Zhao, "Magellan: Charting Large-Scale Peer-to-Peer Live Streaming Topologies," in *Proceedings of International Conference on Distributed Computing Systems (ICDCS)*, 2007.
- [6] R. Kumar, Y. Liu, and K. Ross, "Stochastic fluid theory for p2p streaming systems," in *Proceedings of IEEE INFOCOM*, 2007.
- [7] N. Magharei, R. Rejaie, and Y. Guo, "Mesh or Multiple-Tree: A Comparative Study of Live P2P Streaming Approaches," in *Proceedings of IEEE INFOCOM*, 2007.
- [8] L. Massoulié, A. Twigg, C. Gkantsidis, and P. Rodriguez, "Randomized decentralized broadcasting algorithms," in *Proceedings of IEEE INFOCOM*, 2007.
- [9] Y. Liu, "On the Minimum Delay Peer-to-Peer Video Streaming: how realtime can it be?" in *Proceedings of ACM Multimedia*, 2007.
- [10] X. Hei, C. Liang, J. Liang, Y. Liu, and K. Ross, "A Measurement Study of a Large-Scale P2P IPTV System," *IEEE Transactions on Multimedia*, November 2007.
- [11] X. Hei, Y. Liu, and K. Ross, "Inferring Network-Wide Quality in P2P Live Streaming Systems," *IEEE Journal on Selected Areas in Communications, the special issue on advances in P2P streaming*, 2008.
- [12] PlanetLab, "PlanetLab Homepage," <http://www.planet-lab.org>.
- [13] Y. Pai, K. Kumar, K. Tamilmani, V. Sambamurthy, and A. E. Mohr, "Chainsaw: Eliminating trees from overlay multicast," in *IPTPS*, 2005.
- [14] N. Magharei and R. Rejaie, "PRIME: Peer-to-Peer Receiver-driven Mesh-based Streaming," in *Proceedings of IEEE INFOCOM*, 2007.
- [15] M. Wang and B. Li, " R^2 : Random Push with Random Network Coding in Live Peer-to-Peer Streaming," in *IEEE Journal on Selected Areas in Communications*, 2007.
- [16] Y. Guo, C. Liang, and Y. Liu, "AQCS: Adaptive Queue-based Chunk Scheduling for P2P Live Streaming," in *Proceedings of IFIP Networking*, 2008.
- [17] J. Li, P. A. Chou, and C. Zhang, "Mutualcast: An Efficient Mechanism for Content Distribution in a P2P Network," in *Proceedings of Acm Sigcomm Asia Workshop*, 2005.
- [18] D. M. Chiu, R. W. Yeung, J. Huang, and B. Fan, "Can network coding help in p2p networks?" in *Proceedings of IEEE NetCod*, 2006.
- [19] M. Chen, M. Ponc, S. Sengupta, J. Li, and P. A. Chou, "Utility Maximization in Peer-to-Peer Systems," in *Proceedings of ACM SIGMETRICS*, 2008.
- [20] S. Sengupta, M. Chen, P. A. Chou, and J. Li, "On Optimality of Routing for Multi-source Multicast Communication Scenarios with Node Uplink Constraints," in *Proceeding of 2008 IEEE International Symposium on Information Theory (ISIT 2008)*, 2008.
- [21] M. Zhang, Q. Zhang, L. Sun, and S. Yang, "Understanding the Power of Pull-based Streaming Protocol: Can We Do Better?" *IEEE Journal on Selected Areas in Communications*, 2007.
- [22] C. Liang, Y. Guo, and Y. Liu, "Hierarchically clustered p2p streaming system," in *Proceedings of GLOBECOM*, 2007.
- [23] Trickle, "Trickle Homepage," <http://monkey.org/~marius/pages/?page=trickle>.
- [24] V. Vishnumurthy and P. Francis, "On heterogeneous overlay construction and random node selection in unstructured p2p networks," in *Proceedings of IEEE INFOCOM*, 2006.



Yang Guo is a senior member of technical staff at Thomson Corporate Research, Princeton, NJ. His research interests include peer-to-peer networking and content distribution, IPTV service, media streaming, real-time system, and network modeling and performance evaluation. He received his B.S. and M.S. degrees from Shanghai Jiao Tong University, and his Ph.D. from University of Massachusetts at Amherst.



Yong Liu has been an assistant professor at Electrical and Computer Engineering department of Polytechnic Institute of NYU since March, 2005. He received his Ph.D. degree from Electrical and Computer Engineering department at the University of Massachusetts, Amherst, in May 2002. He received his master and bachelor degrees in the field of automatic control from the University of Science and Technology of China, in July 1997 and 1994 respectively. His general research interests lie in modeling, design and analysis of communication networks. His current research directions include robust network routing, Peer-to-Peer IPTV systems, overlay networks and network measurement. He is a member of IEEE and ACM.



Chao Liang received his B.Engr. and M.Engr. degrees from Department of Electronic and Information Engineering, Huazhong University of Science & Technology (HUST), China, in 2000 and 2002, respectively. He is currently a Ph.D. candidate at the Department of Electrical and Computer Engineering, Polytechnic University, Brooklyn, New York. His research interests include network optimization in overlay and wireless networks, corresponding algorithm and protocol design.