

Inferring Network-Wide Quality in P2P Live Streaming Systems

Xiaojun Hei, Yong Liu *Member, IEEE* and Keith W. Ross *Fellow, IEEE*

Abstract—This paper explores how to remotely monitor network-wide quality in mesh-pull P2P live streaming systems. Peers in such systems advertise to each other *buffer maps* which summarize the chunks of the video stream that they currently have cached and make available for sharing. We demonstrate how buffer maps can be exploited to monitor network-wide quality. We show that the information provided in a peer’s advertised buffer map correlates with that peer’s viewing-continuity and startup latency. Given this correlation, we remotely harvest buffer maps from many peers and then process these buffer maps to estimate the video playback quality. We apply this methodology to a popular P2P live streaming system, namely, PPLive. To harvest buffer maps, we build a buffer-map crawler and also deploy passive sniffing nodes. We process the harvested buffer maps and present results for network-wide playback continuity, startup latency, playback lags among peers, and chunk propagation patterns. The results show that this methodology can provide reasonably accurate estimates of ongoing video playback quality throughout the network.

Index Terms—Measurement, Peer-to-Peer streaming, IPTV, Video Quality Monitoring

I. INTRODUCTION

P2P live streaming systems, using mesh-pull (also known as data-driven or BitTorrent-like) architectures, have enjoyed many successful deployments to date. These deployments include CoolStreaming[1], PPLive[2], PPStream[3], UUSee[4], SopCast[5], TVAnt[6], and many more. These systems have mostly originated from China; however, they have a significant user base in Europe and North America [7], [8], [9], [10]. We will likely observe continued growth in P2P live streaming over the Internet. In the upcoming years, we can expect these systems to distribute more international television channels, reaching larger and more diverse audiences. Moreover, thanks to the low distribution cost of P2P streaming, we also expect user-generated live video content (emanating from web cams and wireless devices) to be distributed using P2P mesh-pull streaming architectures. A video channel in a P2P live streaming system may be simultaneously watched by thousands or even millions of users from all over the world.

It is important to monitor the network-wide quality in a mesh-pull P2P live streaming system. Service quality metrics

include video playback continuity, initial startup latency, and playback time lags among users. Numerous parties have interests in monitoring video quality:

- The users’ viewing experience is crucial for a successful service deployment. If users experience frequent freezes in video playback, significant delays at startup or when switching channels, or large time lags among users for the same video, then the users may abandon the service. Service providers would like to detect when service quality degrades, so that they can add, retroactively and dynamically, additional uploading capacity, either using servers that they operate or that are provided by Content Distribution Networks (CDN). When the service quality is satisfactory, the infrastructure capacity can be released for other purposes.
- In order to compare the service quality of different channels across different P2P streaming providers, third-party companies may be interested in monitoring the channels independently. This information can then be provided to users for reference in selecting P2P video providers. It can also be provided to the advertisers who wish to advertise using the streaming systems.
- Given the knowledge of the actual performance of the streaming systems, ISPs and CDNs can dimension their networks and expand capacity more efficiently, attracting more residential and corporate customers.
- Finally, from a research perspective, monitoring P2P live streaming systems allows researchers to gain an in-depth understanding of various system design choices, including algorithms for creating peering partners, scheduling algorithms (for both uploading and downloading), video encoding algorithms, and so on. Such insights will guide the design of future P2P live streaming systems.

The most direct approach to monitor network-wide quality is to include quality monitors in each peer (for example, within the P2P software). Each monitor would record when media objects arrive at the peer and report this information to a central server for processing. Based on the sequence of arrival times, the central processing server can estimate quality metrics such as freezing and start up latencies at each node, and finally combine this information across nodes to obtain network-wide quality measures. However, this approach requires the installation (or integration) of monitors in the peers, which is only possible with the direct control of the P2P live streaming system. Without this control, it is a challenging problem to monitor the quality of thousands of peers scattered around the world.

Manuscript received March 15, 2007; revised July 31, 2007. This work was supported by the National Science Council under Grant ITR-0325726 and Grant CNS-0519998. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Gunnar Karlsson.

Xiaojun Hei and Keith W. Ross are with Department of Computer and Information Science, Polytechnic University, Brooklyn, NY, USA 11201. (Email: heixj@poly.edu and ross@poly.edu)

Yong Liu is with Department of Electrical and Computer Engineering, Polytechnic University, Brooklyn, NY, USA 11201. (Email: yongliu@poly.edu)

In this paper, we examine the problem of remotely monitoring network-wide quality in P2P mesh-pull live streaming systems. Similar to BitTorrent, peers in P2P mesh-pull streaming systems advertise to each other the chunks availability of the video stream that they currently have cached. These advertisements are called *buffer maps*. We illustrate how these buffer maps can be exploited to monitor network-wide quality. In particular, we show that the information provided in a peer's advertised buffer map correlates with that peer's viewing-continuity quality and startup latency. Given this correlation, we can remotely harvest buffer maps from many peers and then process these buffer maps to estimate ongoing quality metrics.

We apply the monitoring methodology to a popular P2P live streaming system, namely, PPLive. Harvesting buffer maps from PPLive is a challenging problem, since PPLive is a proprietary protocol. To harvest buffer maps, we develop a buffer-map crawler and also deploy passive sniffing nodes. We process the harvested buffer maps and present results for network-wide playback continuity, startup latency, playback lags among peers, and chunk propagation. The results show that this methodology can provide reasonably accurate estimates of ongoing quality throughout the network. Although our experiments are with PPLive, the methodology can in principle be used for any mesh-pull P2P live streaming system.

This paper is organized as follows. We end this section with a summary of the related work. In Section II, we provide an overview of mesh-pull P2P streaming systems. In Section III, we present our methodologies for harvesting buffer maps from a mesh-pull P2P streaming system. In Section IV, we demonstrate a strong correlation between the video availability reported in buffer maps and the streaming service quality. In Section V, we use the methodology to estimate network-wide video quality in PPLive, including playback continuity, startup latency, and playback lags. In Section VI, we present measurement results on network-wide video chunk propagation patterns. Finally we conclude the paper in Section VII.

A. Related Work

Video quality is conventionally characterized by video distortion; the quality measurement normally involves the video rate measurement from which video distortion is estimated utilizing empirical rate-distortion curves or various proposed rate-distortion models [11]. However, this approach varies among the video sequences, and is not suitable for video quality measurement over the Internet because these measurements typically require extensive video traffic measurements [12], [13], [14], [15], [16]. In contrast, our approach does not process video directly, but instead employs buffer maps, which summarize the availability of video chunks at various peers. This approach generates significantly less measurement traffic in the network.

There have been a few measurement studies in quantifying large scale P2P streaming systems. If one has the control of the P2P system, the system can conduct various measurements and provide system logs for analysis. In [1], performance

evaluation was conducted for CoolStreaming over PlanetLab and the system logs show that mesh-pull live streaming systems achieve significant more continuous media playback than tree-push streaming systems. Utilizing the system logs provided by UUSee, Wu et al. characterized the topology properties of the UUSee overlay in [17] and the UUSee streaming traffic patterns in [18]. However, these system logs are usually proprietary and not accessible for third parties. The general measurement techniques for P2P systems fall into two categories: passive sniffing and active crawling. For passive sniffing, our previous work [7] is the first measurement study of a large-scale P2P streaming system. We investigated traffic patterns and peer dynamics of the PPLive IPTV system. Our work was followed by two other passive measurement studies [9] and [10]. Ali et al. [9] studied the traffic characteristics of controlled PPLive peers on PPLive and SopCast. Passive sniffing was also utilized to study the traffic patterns of PPLive, PPStream, TVAnts and SopCast in [10].

Passive sniffing techniques are often constrained to measure a small set of controlled peers. To understand network-wide peer churn, we developed active crawling apparatus to measure the global view of the PPLive network [8]. Subsequently, another crawler-based measurement study was conducted in [19]. Vu et al. [19] examined a small number of PPLive channels and studied various system properties including peer behaviors, overlay structure and channel sizes.

There are a number of measurement studies of other types of P2P systems, including file sharing, content-distribution, and VoIP (see [8] and the references therein). Recently, Isdal et al. [20] proposed to measure the block/piece availability via the BitField message to infer the download rate of end-hosts in BitTorrent. The utilization of this BitField message shares some similarity with our measurement of buffer maps; nevertheless, the major difference is that we infer the video playback performance of streaming peers.

II. OVERVIEW OF MESH-PULL P2P LIVE STREAMING SYSTEMS

Mesh-pull streaming systems (including CoolStreaming [1], PPLive [2], PPStream [3] and UUSee [18]) have a common generic architecture. As shown in Fig. 1, mesh-pull P2P architectures have the following characteristics:

- 1) A live video stream is divided into media chunks (e.g., each chunk consisting of one second of media data) and is made available from an origin server.
- 2) A host, interested in viewing the video, requests from the tracker server a list of hosts which are currently watching the video (Step 1 in Fig. 1). The host then establishes partner relationships (e.g., TCP connections) with a subset of hosts on the list (Step 2 in Fig. 1). The host may also establish a partner relationship with the origin server.
- 3) Each host viewing the video caches and shares chunks with other hosts viewing the same video. In particular, each host receives *buffer maps* from its current partners. A buffer map indicates the chunks that the partner has cached for sharing. Using a scheduling algorithm, each

host requests from its partners the chunks that it will need in the near future.

- 4) Similar to BitTorrent, each host continually seeks new partners from which it can download chunks.

An important characteristic of mesh-pull P2P systems is the lack of an (application-level) multicast tree - a characteristic particularly desirable for the highly dynamic, high-churn P2P environment [1]. Although these mesh-pull systems have similarities with BitTorrent [21], [22], BitTorrent by itself is not a suitable streaming delivery architecture, since it does not account for the real-time needs of streaming applications.

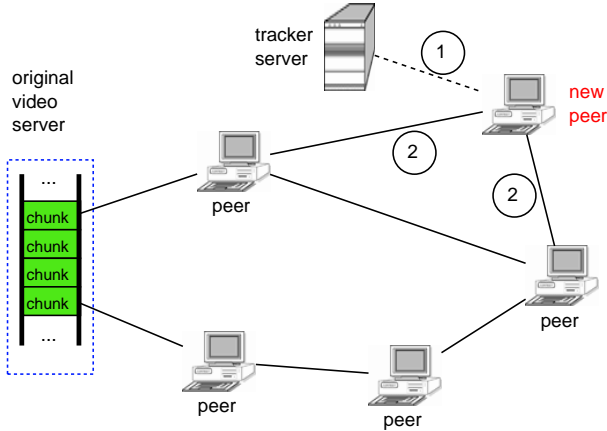


Fig. 1. Mesh-pull P2P live streaming architecture

Fig. 2 shows the software architecture of a peer in a mesh-pull system. The peer includes a P2P streaming engine and a media player. The streaming engine has the job of (i) retrieving chunks from partner peers and (possibly) from the origin server; (ii) storing the retrieved chunks in a cache; (iii) sharing media chunks stored in its cache with its partners; (iv) sending a copy (of the data) of each chunk it receives to the media player. As shown at the bottom of Fig. 2, the peer sends a buffer map to each of its partner peers. The partner peer, having learned from the buffer map what chunks the peer has, issues requests for specific chunks. The peer then sends the requested chunks to the partner.




Fig. 2. A peer includes P2P streaming engine and media player. The streaming engine trades chunks with partner peers.

A. Buffer Map and P2P Streaming Engine

The buffer maps (BMs) play a central role in our approach to infer network-wide quality. At any given time instant, the P2P streaming engine caches up to a few minutes video chunks within a sliding window. Some of these chunks may be chunks that have been recently played; the remaining chunks are chunks scheduled to be played in the next few minutes. Peers download chunks from each other. To this end, peers send to each other buffer map messages; a buffer map message indicates which chunks a peer currently has buffered and can share. Each chunk is assigned with a unique increasing ID. Specifically, the buffer map message includes the offset (the ID of the first chunk), the width of the buffer map, and a string of zeroes and ones indicating which chunks are available (starting with the chunk designated by the offset), as illustrated in Fig. 3.




Fig. 3. A peer’s buffer map, which indicates the chunks it currently has cached.

A peer can request a buffer map from any of its current partner peers. After peer A receives a buffer map from peer B, peer A can request one or more chunks that peer B has advertised in the buffer map. A peer may download chunks from tens of other peers simultaneously. The streaming engine continually searches for new partners from which it can download chunks. Different mesh-pull systems may differ significantly with their peer selection and chunk scheduling algorithms. Peers may also download chunks from the origin server. The chunks are typically sent over TCP connections, although in some mesh-pull systems, video chunks are also delivered using UDP.

B. Media Player

When the client application is started, the media player is launched and the URL of the video stream is provided to the media player. From the client’s perspective, the server of the media content is the P2P streaming engine (which is in the same host as the media player). Once the media player is initialized, it sends (typically) an HTTP request to the P2P streaming engine. After having received the request, the P2P streaming engine assembles the header information and these video chunks in order into a media file and delivers of the file to the media player. Because new chunks continually arrive at the streaming engine, the streaming engine continually adds data to the file. Because some chunks may not arrive before the playback deadline, there may exist “gaps” in the file. When the media player begins to receive the video from the streaming engine, it buffers the video before playback. When it has buffered a sufficient amount of continuous video content, it begins to render the video.

A video stream is commonly encoded in one of four video formats: Window Media Video (WMV), Real Video (RMVB), Quicktime video or Macromedia flash video. For the sake of concreteness, let us consider mesh-pull P2P streaming systems that use WMV (as is often the case today). WMV videos use the Advanced Systems Format (ASF) [23]. An ASF file starts with a header object and is followed by a series of media data objects. The file header object includes various meta information for the file (i.e., the file identifier, buffering time and overall file bit-rate), number of available audio and video streams carried in the file, and meta information for each audio (i.e., sampling rate) and video stream (i.e., spatial resolution, compression algorithm). Each data object starts with an object header. This object header includes the sequence number of the media object, the number of streams carried in the object, playback duration of the object, and so on. In the payload of data objects, multiple media objects from different streams (for example, a video stream and an audio stream) may be carried and interleaved.

Because of the real-time nature of live streaming, each media chunk has a playback deadline (which can be different from one peer to another by several minutes). When a chunk does not arrive before its playback deadline, the peer has two options: it can **freeze** the playback with the most recently displayed video chunks and wait for the missing chunks to arrive; or it can **skip** the playback of the missing video chunks and advance the deadlines for the subsequent chunks accordingly. The freeze is a natural consequence of the media playback when there are no video chunks in the player buffer. If there are still chunks available in the player buffer, the player continues playback even though those chunks might not be continuous; in this case, video playback skipping occurs. We have also observed that in many P2P live streaming systems, when the P2P streaming engine has starved for video chunks over an extended time period, the engine terminates the connection with the player and re-initializes the entire streaming process; we refer to this impairment as **rebooting**.

III. HARVESTING BUFFER MAPS

As described in Section II, peers in mesh-pull streaming systems exchange buffer maps between themselves. These buffer maps can be used to infer many network-wide characteristics of P2P live streaming systems. These characteristics include playback quality, time lags among users, and the video chunk propagation patterns throughout the network. Our methodology has two major steps: (i) continually harvest buffer maps from peers scattered throughout the network; (ii) process (either in real-time or off-line) the buffer maps to infer characteristics about network-wide system behaviors. In this section we describe two approaches for harvesting buffer maps: *active buffer-map crawling* and *buffer-map sniffing*. We also describe how we specifically harvested buffer maps from the PPLive streaming system.

A. Buffer-Map Crawling

Buffer-map crawling is the process of remotely contacting a large number of peers, and continually requesting and

receiving buffer maps from these peers. To perform active buffer-map crawling, we accomplish the following tasks:

- Track the IP addresses (and port numbers) for the peers in the system, accounting for peer churn. Tracking the peers is typically achieved by requesting lists of peers from the tracker and/or from the peers themselves [8].
- Using the knowledge of the specific protocol of the P2P live streaming system, send buffer-map request messages to the tracked peers.
- Receive buffer-map messages from the tracked peers, and again using the knowledge about the specific protocol, extract the actual buffer maps from the buffer-map messages.

It is challenging in building a buffer-map crawler. Most of the P2P mesh-pull live streaming systems are proprietary. In order to establish partnerships with peers, generate buffer map requests and interpret buffer map responses, it becomes necessary to conduct an analysis of the proprietary protocol. If one is privy to the signaling protocol, then these tasks are relatively straightforward. Otherwise, this is typically accomplished by collecting packet traces and analyzing them using one's basic knowledge of how a mesh-pull P2P live streaming system operates. Second, a peer usually retrieves buffer maps by establishing a TCP connection from a remote peer; however, operating systems in general have constraints on the maximum rate to accept new TCP connections. When the remote peer has received a large number of TCP connection requests, there may not be a response to a buffer map request by the operating system. In addition, the peer application may not respond to buffer map requests when the application has already been overloaded with TCP connections. Utilizing TCP to probe the remote peers for buffer maps also incurs unavoidable overhead due to the establishment and termination of TCP connections. Therefore, a single crawler may not be able to harvest quickly an accurate snapshot of the buffer map distribution of peers in a popular channel.

B. Buffer-Map Sniffing

Recall that each peer participating in the system sends and receives buffer maps to and from other peers. We can therefore harvest buffer maps as follows:

- Connect several peer clients to the P2P live streaming system.
- At each connected client, run a sniffer which captures all traffic sent and received by the client.
- Using the knowledge about the specific knowledge about the protocol, extract the buffer maps, along with the origin (IP addresses) of the buffer maps.

The sniffing nodes should be distributed globally throughout the network, in order to capture a diverse set of buffer maps. There is a tradeoff with respect to the number of sniffing nodes deployed. On one hand, we would like to have a large number of sniffers, which increases the number of buffer maps collected. On the other hand, each sniffer node perturbs the system, as it downloads and redistributes chunks.

C. Harvesting Buffer-Maps in PPLive

PPLive is a typical mesh-pull P2P streaming system. PPLive provides 320+ channels as reported in Feb. 2007 [8]. The bit rates of video programs mainly range from 300 kbps to 800 kbps. PPLive does not own video content; the video content is mostly feeds from TV channels, TV series and movies in Mandarin. The channels are encoded in two video formats: Window Media Video (WMV) or Real Video (RMVB). The encoded video content is divided into chunks and distributed to users through the PPLive P2P network.

It is challenging to conduct measurement studies over PPLive because the PPLive protocol is proprietary. In particular, in order to build the buffer map harvesting tools that are used in this paper, we had to analyze the PPLive protocol using packet traces from passive sniffing. With these efforts, we were able to understand critical portions of PPLive’s signaling protocols.

Our PPLive buffer-map crawler utilizes both UDP and TCP transport for collecting peer lists, and TCP for collecting buffer maps from many remote peers during the same time period. Note that the PPLive protocol is still evolving. Recently, we observe that buffer maps start to be carried using UDP in PPLive. This brings forth the possibility to develop a light-weight buffer map crawler using UDP. Fig. 4 shows the architecture of our PPLive buffer-map crawler. Recall from Section II that each peer watching a particular channel maintains a peer list, which lists other peers currently watching the same channel. In PPLive, any peer A can send to any other peer B, within a UDP datagram or a TCP segment, a request for peer B’s peer list. The crawler sweeps across the peers watching the channel and obtains the peer list from each visited peer. Given the IP addresses and the TCP/UDP service ports of the remote peers on the peer list, we augment the crawler with a TCP component that retrieves the buffer map from these remote peers. To this end, as we crawl each peer, the crawler sends a request for the peer’s buffer map. We then parse the buffer maps off-line, to glean information about buffer resources and timing issues at the remote peers throughout the network.

In general, operating systems have a maximum TCP thread constraint that one host is able to maintain. To track a large set of the peers, we deploy a master-slave crawling architecture to harvest buffer maps through the network. As shown in Fig. 4, in step 1, the peer list crawler tracks the peers of one channel in the PPLive network. This peer list crawler serves as a buffer map crawling master. At the same time, we initiate dozens of buffer map crawling slaves over PlanetLab [24]. The buffer map master assigns 30 peers to each slave. Each slave then probes these 30 peers to harvest their buffer maps.

While running the crawler, we utilize WinDump [25] to passively sniff from three PPLive nodes under control in different geographical locations: one in New York with the residential cable-modem access; one in New York with the university Ethernet access; and one in Hong Kong with the university Ethernet access. The buffer maps are embedded in the application data in these captured traffic traces. Given the knowledge of the PPLive signaling protocol, we are able to

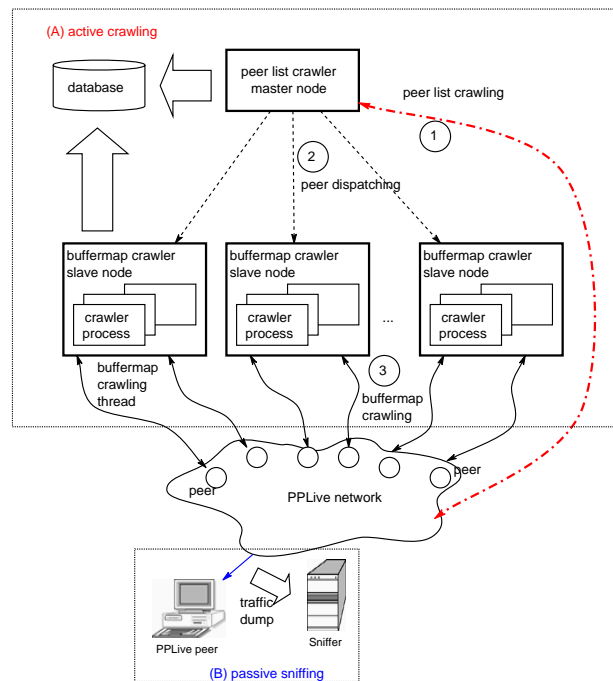


Fig. 4. Harvest buffer maps using crawling and sniffing

extract the buffer maps from these traces for our analysis. A summary of the five sets of buffer maps collected (from the crawler and sniffers combined) is given in Table I. The “x-star” popularity in the table is marked for each PPLive channel. This popularity parameter is computed based on the peer statistics and is provided by the PPLive web site. We use it for a reference of the channel popularity.

IV. INFERRING PLAYBACK QUALITY FROM BUFFER MAPS

In the previous section, we showed how buffer maps can be harvested from a P2P live streaming system. In this section, we illustrate how the buffer maps are used to infer important performance metrics of the P2P live streaming system, including playback continuity and start-up latency.

A. Automatic Playback Continuity Evaluation

In order to study the relationship between buffer-map values and playback continuity, it is desirable to have a mechanism that can automatically evaluate the playback continuity of a video stream received at a media player. To this end, we developed a stream analyzer to capture the arrival process of the media objects to the media player from the streaming engine, and estimate the user perceived quality accordingly. Our analyzer tracks the sequence numbers of the arriving media objects and the presentation time of the latest media object, from which we estimate the playback impairments including freezing and rebooting. We now describe how we estimate when playback freezing occurs.

Typically, a media file consists of two streams, video and audio stream. We use the video stream for illustration. A video stream consists of a sequence of video objects. Each video object is identified by an increasing sequence number i ($i =$

TABLE I
SUMMARY OF HARVESTED BUFFER MAPS

channel name	popularity	bit rate (bps)	chunk size (byte)	buffer maps	total peers	duration (minutes)
CCTVNews	5-star	381,960	4,774	1,952,837	8,120	635
CCTV1	4-star	381,960	4,774	1,440,117	6,668	636
CCTV3	4-star	381,960	4,774	1,789,331	12,392	650
CCTV8	3-star	381,960	4,774	1,359,236	6,985	751
CCTV9	1-star	516,802	6,350	1,263,097	2,125	1,031

1, 2, ...). Let t_i be the relative presentation time of the i -th object, which is obtained from the video object header. Let a_i be the arrival time, which is recorded when our analyzer receives this object. Let τ be the initial playback pre-buffering delay, which is specified by default in the header of video stream or it can be configured manually by users in the player settings. Let b_i and f_i be the actual playback starting time and finishing time, respectively.

The presentation time stamped on the video objects can be used to estimate when video freezing occurs. Fig. 5 depicts how we estimate video freeze and resume when all video objects arrive in order. After the video object p_2 finishes playback at time f_2 , there are no video objects in the buffer ($f_2 < a_3$); hence, our analyzer declares that playback freezes at time f_2 . During the playback freezing period, the analyzer continues receiving video objects from the streaming engine. The video object p_3 arrives at time a_3 . The analyzer starts to buffer video objects without playback. When it receives sufficient video objects, the time duration of the buffered objects ($t_4 - t_3$) is larger than the pre-buffering threshold τ and the playback restarts. Hence, the analyzer estimates the video freezing period as $a_4 - f_2$.




Fig. 5. Playback time diagram with video freezing

We have just described how our analyzer estimates when freezing occurs and the freezing duration for the case when all video objects arrive at the analyzer in order. Our analyzer also estimates when rebooting occurs. We have validated the analyzer, by physically watching PPLive channels and observing the analyzer, to ensure that those impairments actually occur.

B. Correlation between Buffer Maps and Video Continuity

In this section, we demonstrate that there exists a correlation between the buffer map values and the playback impairment events as determined by our analyzer. Because of this correlation, we are able to quantify network-wide playback quality from the harvested buffer maps in Section V and VI.

We performed a number of experiments, and in each of these experiments such a correlation exhibits. As an illustrative example, we connected a campus peer to a popular, four-star PPLive channel (CCTV1) for a three-hour period. During this three-hour period, the number of peers viewing the channel increased from about 1000 to about 2800, as shown in Fig. 6. These values were obtained from our peer-list crawler. During this time period, the output of the PPLive streaming engine was fed into our streaming analyzer. In addition, we captured the buffer maps that were sent from this monitored campus peer to its partner peers.





Fig. 6. Peer evolution in trace CCTV1


In Fig. 7(a), we plot the Buffer Map (BM) width and the BM playable video, which are obtained from the captured buffer maps. As shown in Fig. 3, the **BM width** is the difference between the newest and oldest chunk ID advertised in a buffer map message. The **BM playable video** is the number of contiguous chunks in the buffer map, beginning from the offset. We also plot the **analyzer playable video** (in seconds) as determined by our stream analyzer in Fig. 7(b). It is important to note that the BM playable video, as determined from the buffer maps, and the analyzer playable video, as determined from the stream analyzer, are not equal because the actual playback point is somewhere in the middle of the BM playable video.

We observe that the shape of the BM playable video curve, obtained from the buffer maps, is very similar to the analyzer playable video, obtained from the analyzer. Note in particular that both curves show a brief period of no buffered video in the vicinity of $t = 481$ minutes.

In Fig. 8 we zoom in on the buffer map evolution and also plot the time instants of the playback impairment events reported by the analyzer, including two streaming reboots and four playback freezes. We observe that playback impairments indeed occur only when the BM playable video, determined from the buffer maps, becomes very small. Thereafter, we introduce the following heuristic to detect freezing events:



(a) Buffer width and playable video determined from buffer maps



(b) Buffered playable video determined from analyzer

Fig. 7. Video buffering for a campus peer watching CCTV1

given a sequence of buffer maps for one peer, record the time instant t_1 of the first buffer map when the BM playable video is less than 5 chunks. We then continue to track the BM playable video of consecutive subsequent buffer maps. If the BM playable video remains less than 5 chunks for 4 seconds, we say a freezing event occurs at t_1 . A freeze event ends when the BM playable video is larger than the pre-buffering time of the stream.




Fig. 8. Playback impairments of a campus peer in trace CCTV1

In addition to the playable video size, the offset value in the buffer map can be used to infer the playback continuity. Under normal circumstances, the buffer map offset should increase at the video playback rate: if this rate varies significantly, a playback reboot occurs. Fig. 9 shows the buffer map offset increasing rate of this campus peer. For this video trace CCTV1, the media playback bit rate is 381 kbps; the buffer map offset normally increases at $r = 10$ chunks/sec = $4774 \cdot 8 \cdot 10$ kbps = 381 kbps. We observe that when there is a sudden change in the offset rate, a reboot occurs, which also explains the first reboot occurrence in Fig. 8. We use the following heuristic to detect reboot events: compute the average rate at which the

offset increases over 40 seconds. If this average rate is 10% higher than the nominal playback rate or it is less than 90% of the nominal playback rate, the potential reboot alarm is issued. If in 3 consecutive offset rate checking rounds, which lasts over 120 seconds, all three reboot alarms are issued, we say a playback reboot occurs. Otherwise, the reboot alarm is cleared.




Fig. 9. The rate of the buffer map offset of a campus peer in trace CCTV1

C. Correlation between Buffer Maps and the Start-up Latency

Buffer maps can also be used to infer the start-up latency when the streaming engine starts up the video player for playback. As shown in Fig. 10, in PPLive, when a peer joins the network and has no chunks in its engine buffer, it sends buffer maps with the playback offset of zero and a buffer size of zero. After a peer obtains chunks from other peers, the offset is set to a non-zero value. The streaming engine then starts the player when the size of the cached playable video size exceeds a specified pre-buffering video threshold. This threshold is a system parameter of the streaming engine, which an end user may configure in the option settings of the engine. We use the default pre-buffering threshold value specified in the video file header when collecting statistics in our experiments. The consecutive buffer maps provide snapshots of the video chunks in the buffer. Therefore, we can infer the start-up time of the playback of a peer given the sequence of the buffer maps of this peer with the first buffer map having a zero-valued offset. In this example, the pre-buffering video size is 5 seconds in time (roughly 50 chunks); hence, the startup time is roughly 17 seconds.




Fig. 10. Buffer map evolution at start-up of a campus peer in trace CCTV1

V. NETWORK-WIDE SERVICE QUALITY

Using the methodologies described in the previous section, we can infer video playback quality on remote peers that we

TABLE II
PLAYBACK IMPAIRMENTS ON PEERS OF FIVE CHANNELS WITHIN TEN HOURS

time(hour)	[0,1]	[1,2]	[2,3]	[3,4]	[4,5]	[5,6]	[6,7]	[7,8]	[8,9]	[9,10]
channel	CCTVNews									
# of sampled peers	1206	950	868	973	1230	1414	1253	1054	974	827
reboot peer ratio	18.8%	38.5%	41.9%	14.1%	3.4%	2.2%	2.6%	3%	4.4%	4.1%
freeze peer ratio	1.8%	3.1%	2.9%	2.0%	1%	0.6%	0.7%	0.9%	2.7%	3.1%
channel	CCTV1									
# of sampled peers	301	382	605	1107	1104	926	714	1038	1438	1141
reboot peer ratio	20%	15%	4.8%	2.1%	1.4%	3.6%	1.7%	2.2%	2.4%	1.3%
freeze peer ratio	9.0%	9.0%	2.5%	0.6%	0.8%	2.1%	0.6%	0.8%	1.3%	1%
channel	CCTV3									
# of sampled peers	899	920	996	956	1413	2141	2835	2307	1848	1215
reboot peer ratio	7.6%	6.6%	7.3%	5%	2.8%	7.9%	10.7%	9.5%	8.6%	14.4%
freeze peer ratio	1.1%	1.7%	2.2%	1.4%	0.9%	2%	4.3%	4.2%	2.3%	3.1%
channel	CCTV8									
# of sampled peers	435	393	564	667	860	1188	1313	1426	1085	974
reboot peer ratio	3%	5.1%	3.7%	4.2%	2.7%	2%	0.7%	0.5%	1.2%	1%
freeze peer ratio	1.4%	1%	1.2%	1.4%	0.9%	1.2%	0.7%	0.4%	0.2%	0.3%
channel	CCTV9									
# of sampled peers	207	173	160	184	199	247	340	397	353	311
reboot peer ratio	26.6%	44.5%	38.1%	35.3%	32.2%	40.1%	29.7%	26.2%	15.9%	16.7%
freeze peer ratio	4.4%	20.2%	23.1%	20.7%	12.1%	14.6%	12.7%	6.8%	9.6%	5.5%

cannot monitor directly. In this section, we report network-wide service quality measurement results for several important quality metrics of PPLive on playback continuity, including freezes and reboots, start-up latency and time lags among peers.

A. Playback Continuity

The most important measure of the service quality of a streaming system is the continuity of video playback at end-hosts. We processed the buffer maps in Table I using the heuristics described in Section IV. Table II reports playback continuity, freezes and reboots, experienced by the sampled peers during 10 one-hour periods. For each one-hour period, we select peers with at least 50 captured buffer maps. Among the sampled peers, we count the number of the reboot/freeze peers that have at least 1 playback reboot/freeze event in this one-hour period. The reboot/freeze peer ratio is defined as the ratio between the reboot/freeze peer and the total number of the sampled peers in this one-hour period.

In Table II, we observe that the peers in the monitored channels have smooth video playback most of the time. For channel CCTV3 and CCTV8, more than 90% peers never experienced any freezes or reboots within one hour time periods. For channel CCTVNews and CCTV1, the freeze peer ratio and the reboot peer ratio vary over time. The general trend is that with a larger number of the peers in the system, the freeze and the reboot ratio are lower. Compared with the other four channels, channel CCTV9 has relatively high freeze and reboot ratio over the ten hours period. We find the number of peers in this channel is much lower than in the other four channels.

To reveal more details on playback impairments, we zoom in one time period, from 2.5 hours to 3.5 hours, when a high ratio of peers in channel CCTVNews experience freezes and reboots. Table III lists the freeze and reboot ratios among all

the sampled peers and the number of the playback impairment events on two campus peers under our control every 10 minutes. We observe high reboot ratios ($> 30\%$) before $t = 182$ minutes among the sampled peers and those two campus peers experienced a number of reboots during these 10-minute intervals.

To understand the causes of the frequently occurring playback impairments, we plot in Fig. 11 the BM playable video on the three peers which are located in New York and Hong Kong. Before $t = 182$ minutes, all three peers have a small and highly variable amount of the BM playable video in their buffers. After $t = 182$ minutes, the playable videos in all three peers stay at a high level.

We further look into the rate at which the buffer map offsets increase for the three peers during the investigated intervals. Since the playback time for a chunk is 0.1 second, to maintain steady playback, the buffer map offset should increase at a rate of 10 chunks per second. In Fig. 12, we observe that before $t = 182$ minutes, the three peers cannot advance their buffer map offset at the playback rate of 10 chunks/sec. They experience severe playback impairments. After $t = 182$ minutes, the increasing rate of the buffer map offset reaches the nominal playback rate for all three peers.

Based on the strong correlation of playback impairments on three peers, we infer that this is a network-wide quality degradation. The network-wide service quality is largely determined by the total usable upload bandwidth of the peers in the network. Two factors impact the total usable upload bandwidth: the number of the peers, particularly those super peers with large upload bandwidth, and peer churn. In trace CCTVNews, we suspect the frequent playback impairments before $t = 182$ is because the number of peers in that time range is not large enough. To verify this conjecture, we plot in Fig. 13 the evolution of peer population from time 0 to 210 minutes. The peer number before $t = 180$ minutes oscillates around 700. Starting from $t = 180$ minutes, more peers join

TABLE III
PLAYBACK IMPAIRMENTS ON PEERS OF ONE CHANNEL WITHIN ONE HOUR

time(minute)	[150,160]	[160,170]	[170,180]	[180,190]	[190,200]	[200,210]
playback performance over sample peers						
# of sampled peers	320	305	332	310	315	339
reboot peer ratio	41.8%	41.3%	39.1%	36.8%	5.7%	3.5%
freeze peer ratio	1.9%	2.0%	0.3%	1.3%	3.8%	0.6%
HK campus peer						
# of reboot events	9	11	14	4	0	4
# of freeze events	0	0	0	0	0	0
NY campus peer						
# of reboot events	11	10	17	7	0	0
# of freeze events	0	0	0	0	0	0




Fig. 11. Correlation on BM playable video of three peers in trace CCTVNews




Fig. 12. Correlation on buffer map offset increase rates of three peers in trace CCTVNews

the channel and the peer number keeps increasing until the end of the observation period. This may indicate that the peer population largely affects the network-wide service quality. An in-depth characterization of the total usable upload bandwidth of the P2P system requires to track the upload capacity of each peer, in particular those super-peers, and peer churn. This may lead to an interesting direction for further work.

B. Start-up Latency

The very first quality measure that a user perceives after joining a channel is the start-up latency. We use the heuristic described in Section IV-C to estimate the start-up delay of the peers in trace CCTV9. The total number of the tracked start-up processes is 220. We estimated the upper bound and lower bound of the start-up delay of these tracked peers, as plotted in Fig. 14.




Fig. 13. Evolution of the number of peers in trace CCTVNews

Table IV tabulates the start-up delays for peers in five different traces. For each channel, we report the median values of the lower bound and upper bounds of peers in that trace. The variance of the startup delay is not significantly large.




Fig. 14. Distribution of start-up latencies of peers in trace CCTV9

For example, for CCTV9, the standard deviation (STD) of the lower bound of the startup delay is 41.7 seconds and the STD of the upper bound is 51.3 seconds.

TABLE IV
MEDIAN OF STARTUP LATENCIES

Trace	Captured Peers	Lower bound (sec)	Upper bound (sec)
CCTVNews	38	11.8	17.5
CCTV1	74	13.3	19.8
CCTV3	1405	50.3	94.9
CCTV8	92	5.98	12.8
CCTV9	220	14.2	19.7

C. Playback Lags among Peers

In a mesh-pull P2P streaming system, the progress of video playback on a peer is determined by how quickly it collects video chunks from the system. The variations in chunk retrieval times from peer to peer introduce time lags among the peers.

1) *Maximum Time Lags*: One measure of the peer playback synchronization is the time lags between the “fastest” and the “slowest” peer. To obtain maximum time lags, we merge buffer map traces collected by three monitors and group buffer maps into time bins of 5 seconds with around 150 probed peers in these bins. For each bin, we calculate the difference between the maximum and the minimum offset of the buffer maps in that bin. The difference indicates the maximum playback lag among the peers monitored in that time bin. Fig. 15 shows that the maximum time lags among peers fluctuate around 150 seconds within a one hour time period.




Fig. 15. Maximum time lags among peers in trace CCTVNews

2) *Tiering Effect*: Other than the maximum time lags, we also intend to explore the structure embedded in the delay performance of peers in a mesh-pull P2P streaming system. In a tree-push video streaming system, a user’s playback delay is determined by its distance to the server. When grouped into tiers according to their distances, users in a tier close to the source consistently have smaller playback time delays than users in a tier farther away. In a mesh-pull P2P streaming system, to deal with peer churn, there is no fixed tree structure. Peering links are established/removed dynamically. Consequently, unlike in a tree-push system, the regularity in the relative delay performance among peers cannot be assumed. On the other hand, in a mesh-pull based system, the peering links are *not* purely random. They are driven by peer visibility and chunk availability, which are largely determined by the peers’ network connections and their locations in the streaming system. Hence, one can expect some structure in the peering topology, and the peer delay performance as a result. By examining buffer map traces collected from peers with different geographical locations and different network access, we demonstrate a tiering effect in the playback time lags. In particular, one controlled campus peer is located in Hong Kong and two controlled peers with the Ethernet access and the cable-modem access, respectively, are in New York,

To expose the ordering of the playback delays on different peers, we infer the playback progress of a remote peer by checking the offsets of buffer maps collected from that peer. We define the relative playback difference between two peers as the offset difference of these two peers. More specifically, suppose that we obtain one buffer map of peer A at t_A with the buffer map offset Δ_A , and one buffer map of peer B at t_B with the offset Δ_B . The playback difference is computed as

$$\delta = \frac{(\Delta_A - \Delta_B) * L}{r} - (t_A - t_B),$$

where L denotes the chunk size and r denotes the playback bit rate.

We use three different traces captured during the same time period to explore the tiering effect. For each trace, using the local playback time on the monitor as the reference point, we calculate the relative playback differences for the 6 peers from which we collected the most buffer maps. We plot in Fig. 16 the evolution of the relative playback differences over a one hour time period. We observe that there exists a clear ordering among peers’ playback lags in all three traces. In addition, the relative playback differences are quite stable over the one-hour time period. From the playback lags between two peers, one can also infer the direction of video data flow between them. It is unlikely for a peer to download video from a peer whose playback time is tens of seconds behind.

In the trace collected from the Hong Kong campus peer, Fig. 16(a) shows that only one peer out of the six peers is ahead of the local peer. This confirms that the Hong Kong campus peer is close to the video source server. This peer downloads video chunks quickly from the server and redistributes them to other peers in Hong Kong or places farther away. In Fig. 16(b), in the trace collected from the campus peer in New York, about half of its peers are ahead




Fig. 16. Tiering effect in trace CCTVNews: peers have a stable relative order on their playback lags

of it and the other half are behind it. It appears that this peer downloads video chunks from peers in Asia and serves as a proxy to upload video chunks to peers in the US. Although in the same city as the New York campus peer, the cable peer in New York has very different delay performance relative to its peers. Fig. 16(c) shows that most of its peers are ahead and this residential peer uploads little traffic to its peers. We also calculate the relative playback delay among the three monitors. The Hong Kong campus peer's local playback time is around 80 seconds ahead of the New York campus peer; the New York residential peer's local playback time is 20 seconds behind the New York campus peer. This suggests that both the location and network access of a peer may significantly impact its playback delay.

VI. VIDEO OBJECT PROPAGATION

Buffer maps provide valuable information for inferring how video chunks propagate among peers inside the network. Chunk propagation in a P2P streaming system is governed by peering strategies and chunk scheduling schemes. It is not possible to sniff traffic on all peers and trace the detailed propagation of a chunk accordingly. However, we can utilize the harvested buffer maps to infer when a chunk is downloaded by a monitored peer. Thereafter, we obtain a sampled view of the availability of chunks among all peers in the system. After a chunk is generated by the server, it is downloaded by peers. After a peer successfully downloads the chunk, the peer shares it for other peers to download until that chunk falls outside its buffer map window. How long a chunk is available in the system and how many peers can upload that chunk at a given time can be inferred from the buffer maps.

A. Chunk Life Time

We first investigate the chunk life time. The life time of a chunk is defined as the duration from the first appearance of that chunk in the collected buffer maps to the last time of its appearance. The chunk life time characterizes how long a chunk is available for download in the network. Since we only collect buffer maps from a subset of peers, the reported chunk life time is an under-estimate of the real chunk life time. In Fig. 17(a), a chunk stays in one peer's buffer for about 210 seconds; however, this chunk may still be available in some other peers' buffer. We observe that the chunk life time in the network is about 250 seconds.

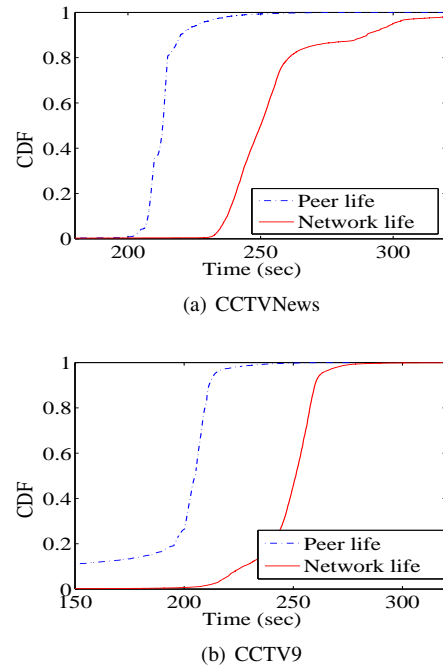


Fig. 17. Distribution of the chunk lifetime by tracing the campus peer in New York

B. Evolution of Chunk Availability

We define the availability of a chunk at a given time as the fraction of peers that make that chunk available for other peers to download. With a higher the chunk availability, the upload supply for that chunk is higher and the download demand is lower. Fig. 18 shows the availability snapshots for a range of chunks among randomly sampled peers in two channels. In Fig. 18(a), we tracked 4500 consecutive chunks among the sampled 122 peers. The tracking results are reported every 90 seconds. We observe that the chunk availability percentage can reach as high as 80%. For a given snapshot, a wide window of chunks for 150 seconds have the availability above 70%. The position of the window advances at the video playback rate from one snapshot to another. Shown in Fig. 18(b), in the 1-star channel CCTV9, the chunk availability percentage is much smaller than that in the popular channel CCTVNews. For a given snapshot, chunks falling in the window only have availability around 40%. This indicates that the peers

in CCTV9 have difficulty in locating chunks and the playback is thus susceptible for various playback impairments.

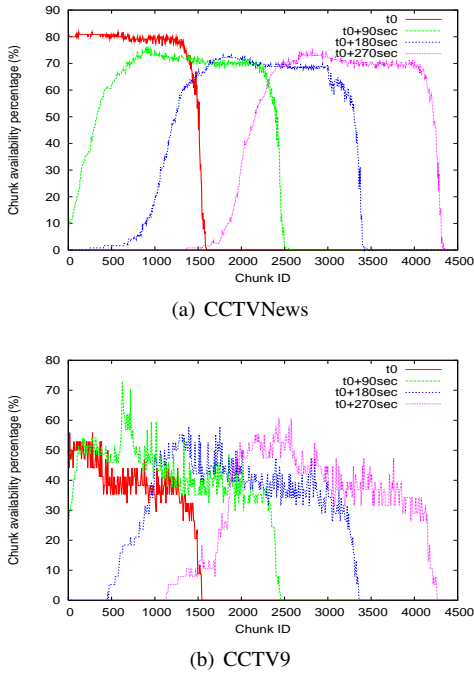


Fig. 18. Chunk availabilities among peers at four time snapshots

C. Evolution of Chunk Retrieval Ratio

The playback quality on a peer is largely determined by whether the peer can retrieve video chunks before their playback deadlines. Using buffer maps, for a specific chunk, we can calculate the fraction of peers that have retrieved the chunk at any time instant. The chunk retrieval ratio among peers should grow over time until it reaches close to 1. The growth rate is a good measure of the P2P video distribution efficiency. Fig. 19 shows the retrieval ratios of five chunks from 122 randomly sampled peers in two channels: CCTVNews with a 5-star popularity and CCTV9 with a 1-star popularity. The five chunks are chosen such that two adjacent chunks have presentation time difference of 10 seconds. All chunk retrieval ratios evolve in a similar pattern. As shown in Fig. 19(a), they increase quickly to about 65% within 20 seconds. The increase trend slows down after the knee point until it converges to around 95%. However, as illustrated in Fig. 19(b), in the less popular channel CCTV9, the knee points come earlier, at around 30%. The chunk retrieval ratios converge in a wider range from 50% to 75%. This suggests that the chunk propagation in a small P2P streaming system is slower and more variable than that in a large system. One direct consequence is that less popular channels have worse quality.

To reveal the impact of the chunk retrieval ratio on individual peer’s playback quality, we plot in Fig. 20(a) the evolution of buffer sizes on the Hong Kong campus sniffer. Fig. 20(b) shows the chunk retrieval ratio among its peers during the same time period. Whenever the sniffer has a low video buffer

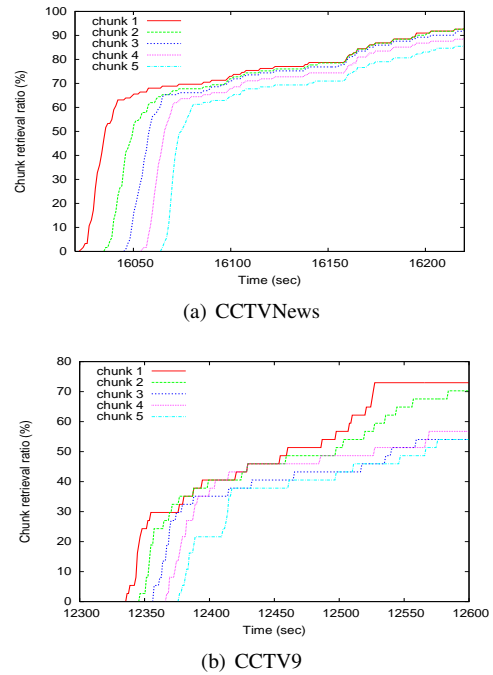


Fig. 19. Evolutions of chunk retrieval ratios for five sampled chunks

level, the chunk retrieval ratio among its peers is below 50%. Streaming service providers may use this chunk retrieval ratio as the performance indicator so that they can add additional uploading capacity to improve the service quality if necessary.

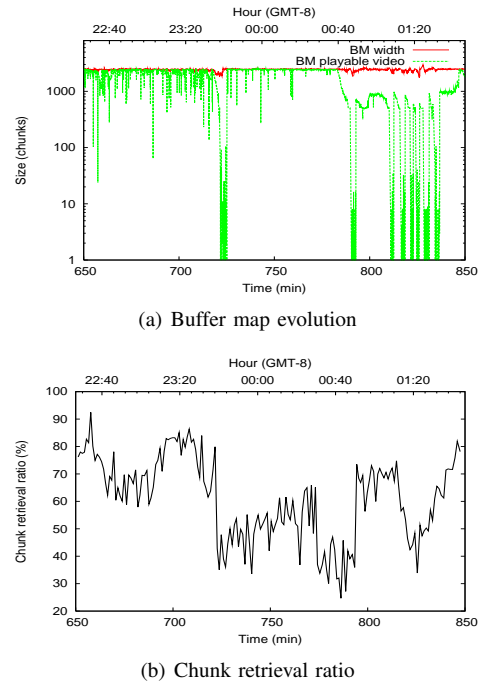


Fig. 20. Correlations between the playback quality of a monitored peer and the chunk retrieval ratio in the whole network

VII. CONCLUSION

Based on remotely harvesting buffer maps, we have developed a methodology to estimate network-wide quality in mesh-pull P2P live streaming systems. A peer's buffer map provides a light-weight snapshot of the chunk availability at that peer. We have shown that the video playback quality can be inferred from buffer maps. We illustrated this methodology by applying it to a specific P2P live streaming system, namely, PPLive. The methodology enabled us to estimate several important network-wide quality measures, including network-wide viewing continuity, startup latency, and time lags. It also enabled us to track the propagation on chunks throughout the P2P network. As discussed in the Introduction, this methodology can be used by service providers, by third parties who monitor performance on behalf of customers and advertisers, and by researchers who seek a deeper understanding of the behaviors of mesh-pull P2P live streaming systems.

ACKNOWLEDGMENT

We thank Jian Liang and Chao Liang for their help in the early stages of this research. The authors also thank anonymous reviewers' constructive comments that help to improve the quality of this paper. We thank the editor Prof. Gunnar Karlsson for his efforts in handling the paper.

REFERENCES

- [1] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "DONet/CoolStreaming: A Data-driven Overlay Network for Peer-to-Peer Live Media Streaming," in *IEEE INFOCOM*, vol. 3, Mar. 2005, pp. 2102 – 2111.
- [2] "PPLive," <http://www.pplive.com>.
- [3] "PPStream," <http://www.ppstream.com>.
- [4] "UUSee," <http://www.uusee.com>.
- [5] "SopCast," <http://www.sopcast.org/>.
- [6] "TVAnts," <http://www.tvants.com>.
- [7] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross, "Insights into PPLive: A measurement study of a large-scale P2P IPTV system," in *IPTV workshop in conjunction with WWW*, May 2006.
- [8] —, "A measurement study of a large-scale P2P IPTV system," *IEEE Transactions on Multimedia*, vol. 9, no. 8, Dec. 2007.
- [9] S. Ali, A. Mathur, and H. Zhang, "Measurement of commercial peer-to-peer live video streaming," in *First Workshop on Recent Advances in Peer-to-Peer Streaming*, Aug. 2006.
- [10] T. Silverston and O. Fourmaux, "Measuring P2P IPTV systems," in *Proc. NOSSDAV*, June 2007.
- [11] M. Pinson and S. Wolf, "A new standardized method for objectively measuring video quality," *IEEE Transactions on Broadcasting*, vol. 50, no. 3, pp. 312 – 322, Sep. 2004.
- [12] D. Loguinov and H. Radha, "Large-scale experimental study of Internet performance using video traffic," *SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 1, pp. 7–19, 2002.
- [13] S. Winkler and R. Campos, "Video quality evaluation for Internet streaming applications," in *Proc. SPIE*, vol. 5007, June 2003, pp. 104–115.
- [14] A. Reibman, S. Sen, and J. V. der Merwe, "Video quality estimation for Internet streaming," in *ACM WWW*, 2005, pp. 1168–1169.
- [15] S. Tao, J. Apostolopoulos, and R. Guerin, "Real-time monitoring of video quality in IP networks," in *Proc. NOSSDAV*, 2005.
- [16] L. Guo, E. Tan, S. Chen, Z. Xiao, O. Spatscheck, and X. Zhang, "Delving into Internet streaming media delivery: a quality and resource utilization perspective," in *ACM IMC*, Oct. 2006, pp. 217–230.
- [17] C. Wu, B. Li, and S. Zhao, "Magellan: Charting large-scale peer-to-peer live streaming topologies," in *IEEE ICDCS*, June 2007.
- [18] —, "Characterizing peer-to-peer streaming flows," *IEEE JSAC*, 2007.
- [19] L. Vu, I. Gupta, J. Liang, and K. Nahrstedt, "Measurement and modeling a large-scale overlay for multimedia streaming," in *Proc. QShine*, Aug. 2007.
- [20] T. Isdal, M. Piatek, A. Krishnamurthy, and T. Anderson, "Leveraging BitTorrent for end host measurements," in *Proc. PAM*, Apr. 2007.
- [21] B. Cohen, "Incentives Build Robustness in BitTorrent," in *Proc. P2P-Econ*, June 2003.
- [22] "BitTorrent," <http://bittorrent.com/>.
- [23] "Advanced systems format (ASF) specification," 2004, <http://www.microsoft.com/windows/windowsmedia/forpros/format/asfspec.aspx>.
- [24] "Planetlab," <http://www.planet-lab.org>.
- [25] "WinDump," <http://www.winpcap.org/windump>.