

IPTV over P2P Streaming Networks: The Mesh-Pull Approach

Xiaojun Hei, Yong Liu, and Keith W. Ross, Polytechnic University

ABSTRACT

IPTV, an emerging Internet application, would revolutionize the entertainment and media industries; however, IPTV also has the potential to overwhelm the Internet backbone and access networks with traffic. To date, IPTV over P2P streaming networks has advanced significantly using two different approaches: tree-push versus mesh-pull. In particular, the mesh-pull streaming approach has achieved a number of successful commercial deployments. In this article, we examine the current progress in the research and development of mesh-pull P2P streaming systems. We provide an overview of the general mesh-pull streaming architecture and review various challenges, design issues, and interesting research problems in this approach. We discuss the construction costs for providing an IPTV service with service guarantees. We outline a measurement technique for monitoring the video playback quality of mesh-pull streaming systems. We emphasize that the future P2P IPTV systems should be designed to meet the expectations of users for quality-of-experience. We also identify a few other important issues for IPTV over P2P streaming networks, including the traffic pressure on ISPs, various security concerns, and the necessity to re-examine the most appropriate P2P architecture. Insights obtained in this study will be valuable for the development and deployment of future P2P IPTV systems.

INTRODUCTION

With the widespread adoption of broadband residential access and the advances in video compression technologies, IPTV may be the next disruptive Internet application. Considering a scenario where hundreds of millions of users watch video programs online at a bit rate of 500 kb/s or more, IPTV would revolutionize the entertainment and media industries; however, IPTV also has the potential to overwhelm the Internet backbone and access networks with traffic. Given this possible tidal wave of new Internet traffic, an in-depth understanding of the IPTV delivery mechanisms is essential, particularly for the delivery architectures that hold the greatest promise for broad deployment in the near future.

Peer-to-peer (P2P) streaming networks do not rely on a dedicated delivery infrastructure and hence offer the possibility of rapid deployment at low cost. The upload capacity of peers can be utilized for video transmission so as to reduce the server load dramatically. Therefore, P2P streaming appears to be the most promising mechanism for the IPTV deployment. Content owners often are not equipped with a strong content delivery infrastructure; the low cost incurred in the P2P streaming paradigm is particularly appealing to these content owners. Nevertheless, the large volume of video traffic generated by P2P applications has been raising the traffic load on the network infrastructures of telecommunication operators.

There exist two major design issues for constructing a P2P streaming network:

- How to form an overlay topology between peers
 - How to deliver video content efficiently
- The current approaches can be classified into two categories:
- Peers form a tree-shaped overlay, and video content is pushed from the origin server to peers, namely the *tree-push* approach
 - Peers form a mesh-shaped overlay, and they pull video from each other for content delivery, namely the *mesh-pull* approach.

Over the years, many tree-push systems have been proposed and were evaluated in academia, and they achieved some success. However, they never took off commercially. Nevertheless, mesh-pull IPTV systems have enjoyed a number of successful deployments to date, such as Cool-Streaming [1], PPLive (www.pplive.com), and others [2–4]. The major advantages of mesh-pull systems are the simple design principle and inherent robustness, particularly desirable for the highly dynamic, high-churn P2P environment. A well-written tutorial on P2P video streaming over the Internet can be found in [5]. Our study in this article dissects the general architecture of the current, most popular mesh-pull P2P streaming systems.

The aim of this article is to explore the general design space and challenges of mesh-pull systems. First, we provide an overview of the general mesh-pull streaming architecture. The operations of the mesh-pull system rely on the bandwidth contribution from each peer. Peers

may have different upload capacity due to different types of access links. Next, we discuss the implication of the heterogeneity in peer uploading capacities on system scalability. User experiences play an important role in a successful IPTV application. We then review the service quality issues and introduce a quality monitoring methodology for the mesh-pull system. Mesh-pull streaming systems employ both TCP and UDP as the underlying transport control protocol. The congestion control problem is revisited. Security is another vulnerable aspect of most current mesh-pull systems. Also, we outline the potential peer poison and video pollution attacks in these systems. Different P2P streaming architectures have different advantages and drawbacks in video playback quality perceived by end users. Pure mesh-pull architectures may find it difficult to meet the stringent delay requirement of future IPTV services. Next, we briefly discuss a hybrid tree-mesh architecture for the next generation P2P IPTV systems. Finally, we conclude the article by highlighting the design challenges for future mesh-pull streaming systems.

MESH-PULL P2P STREAMING ARCHITECTURE

A number of mesh-pull streaming systems have been deployed to date. However, most mesh-pull streaming systems provide little information about their proprietary technologies. Through our measurement studies and protocol analysis on a few mesh-pull streaming systems [2, 6], we gained significant insights into the protocols and streaming mechanisms of mesh-pull streaming systems. In spite of different features and implementation details, they share a common generic architecture that we describe in this section.

SYSTEM OVERVIEW

As shown in Fig. 1, mesh-pull P2P architectures have the following characteristics. A video is divided into media chunks and is made available from an origin server for broadcast. All the video information is accessible for users at the channel server. A host, interested in viewing the video, requests the available video streams from the channel server (Step 1 in Fig. 1). The tracker server maintains a list of hosts who are interested in watching the same video. After a host selects the video in which it is interested, it retrieves a list of hosts currently watching the same video (Step 2 in Fig. 1). The host then establishes partner relationships (TCP/UDP connections) with a subset of hosts on the list (Step 3 in Fig. 1). These peers help each other and deliver video traffic cooperatively. The host also may establish a partner relationship with the origin server. 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 from a remote partner indicates the chunks that are available on that partner. Using a chunk scheduling algorithm, each host requests the chunks that it will require in the near future from its part-

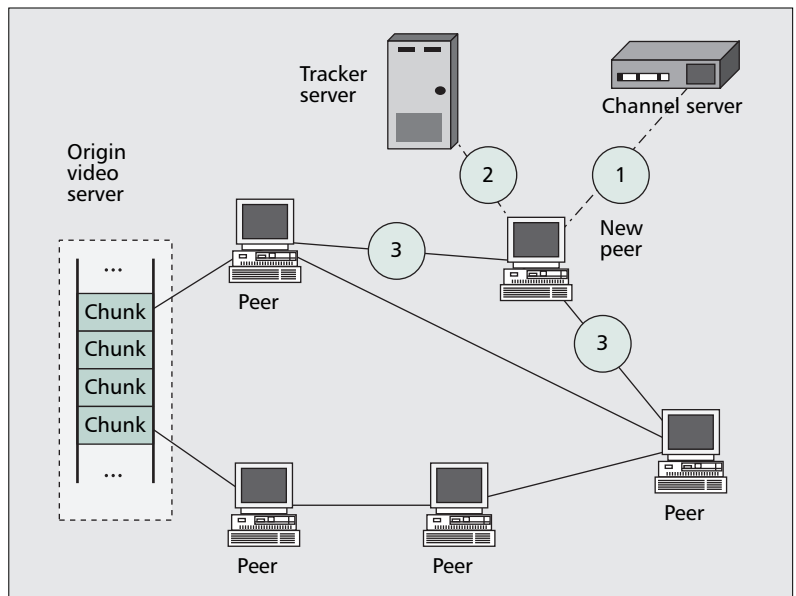


Figure 1. Mesh-pull P2P live streaming architecture.

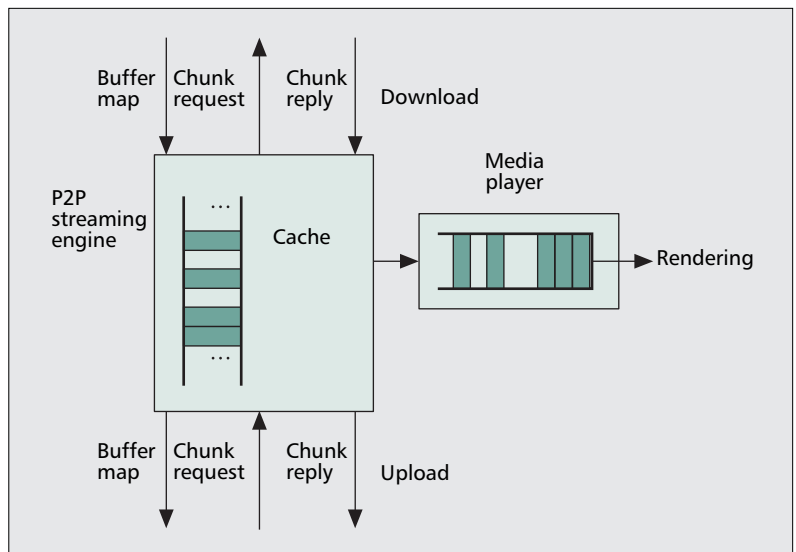


Figure 2. A peer includes a P2P streaming engine and a media player. The streaming engine trades chunks with partner peers.

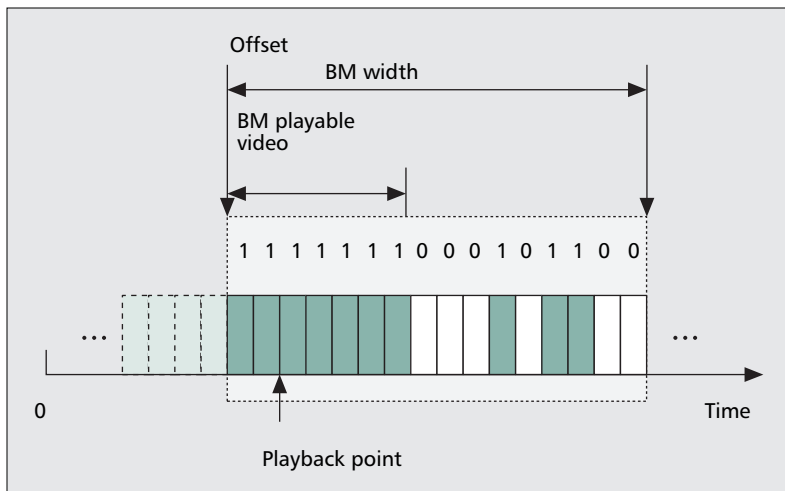
ners. Each host continually seeks new partners from which it can download chunks.

SOFTWARE COMPONENTS

Figure 2 shows the software components of a peer in a mesh-pull system. The peer software includes a P2P streaming engine and a media player. The streaming engine has the job of:

- Retrieving chunks from partner peers and (possibly) from the origin server.
- Storing the retrieved chunks in a cache.
- Sharing media chunks stored in its cache with its partners.
- Sending a copy (of the data) of each chunk that it receives to the media player.

As shown at the bottom of Fig. 2, the local peer sends a buffer map to each of its partner peers. A partner peer, having learned which chunks that the local peer has from the buffer map,



■ **Figure 3.** A peer's buffer map, which indicates the chunks it currently has cached.

issues requests for specific chunks. The local peer then sends the requested chunks to that partner peer.

P2P Streaming Engine — At any given instant, the P2P streaming engine caches up to a few minutes worth of chunks within a sliding window. Some of these chunks may be chunks that recently were 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 the chunks that a peer currently has buffered and can share. A 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). A unique channel ID also is carried in each buffer map message. Figure 3 illustrates the structure of a buffer map.

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 its buffer maps. 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. The selection and scheduling algorithms used by CoolStreaming are documented in [1]. Peers also can download chunks from the origin server. The chunks typically are sent over TCP connections; however, recently, video chunks also are transferred using UDP in more mesh-pull systems.

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). After the media player is initialized, it

(typically) sends an HTTP request to the P2P streaming engine. After receiving the request, the P2P streaming engine assembles its chunks and header information into a media file and delivers the file to the media player. Because new chunks continually arrive to the streaming engine, the streaming engine continually adds data to the file. Because some chunks may not arrive before the playback deadline, there can be "gaps" in the received media 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.

MESH-PULL STREAMING: BITTORRENT REVISITED?

Bearing strong similarities to BitTorrent, mesh-pull streaming deviates significantly from BitTorrent in various aspects. BitTorrent by itself is not a feasible video streaming architecture, because it does not account for the real-time requirements of IPTV. In mesh-pull streaming, each video chunk has a corresponding playback deadline. Hence, video chunk scheduling is an indispensable component for assisting a timely video delivery.

Due to the stringent playback deadline of video chunks, fair resource sharing has not been addressed carefully in the current mesh-pull systems, in that there are no reciprocity mechanisms deployed in the current mesh-pull systems to encourage sharing between peers, such as Tit-for-Tat equipped in BitTorrent.

BitTorrent is targeted at a group communication of medium size (<1000); hence, peers retrieve peer neighbor information directly from the tracker server. However, a large-scale live streaming broadcast can easily attract tens of thousands of users. Hence, gossip-like peer search algorithms have been implemented in various mesh-pull systems to support large-scale group communication. However, the deployment of gossip algorithms produces various results, that is, delay may occur in searching peers; tracker servers may handle only part of the peers in the system and hence, lose the global view and control of the network; and so on.

SYSTEM SCALABILITY: CONSTRUCTION COSTS

In mesh-pull streaming systems, participating peers are very heterogeneous, particularly in terms of the amount of upload bandwidth they contribute [2]. In addition, peers may randomly join the system, watch the video for a random period of time, and then leave the system. These two factors, peer heterogeneity and churn, generate the major challenges in provisioning the P2P IPTV services to ensure that all participating peers can continuously playback the video (without freezing or skipping) with a small playback delay.

The current peers participating in mesh-pull streaming systems are usually broadband resi-

dential peers, with DSL and cable access, and institutional peers, with high-bandwidth Ethernet access. The residential peers, or ordinary peers, typically have upload capacity of 500 kb/s or less; the institutional peers, or super peers, often have upload capacity in excess of a few Mb/s. The service level of P2P streaming systems is largely determined by the resources contributed by all the peers in the system, that is, upload bandwidth. These two types of peers, ordinary and super peers, have different upload capacities. If the upload capacity of ordinary peers cannot sustain the video playback rate, super peers must contribute additional upload capacity to help the ordinary peers. As demonstrated in [7], there exists a critical value of the ratio between ordinary peers and super peers for sustaining a satisfied service level of P2P streaming systems. Roughly speaking, when this ratio exceeds this critical value, the system performs well; otherwise, the system performs poorly.

Most mesh-pull streaming systems to date deploy only limited delivery infrastructure, that is, channel list servers, tracker servers, and so on. The service quality of peers varies significantly at different time and locations. IPTV service providers may invest in additional streaming infrastructure, either using servers that they operate or that are provided by content distribution networks (CDN). To be cost-effective, it is crucial that service providers quickly detect when service quality degrades, so that they can add, retroactively and dynamically, additional uploading capacity. When service quality is satisfactory, the infrastructure capacity can be released for other purposes.

VIDEO PLAYBACK: QoS TOWARD QoE

The IPTV viewing experience of users is crucial for a successful service deployment. If users experience frequent freezes in the video playback, large delays for start up after switching channels, or significant time lags among users for the same video frame, then they may abandon the service. Important IPTV quality metrics include initial start-up delay, video switching delay, playback time among users, video playback continuity, and so on.

QUALITY METRICS

As theoretically demonstrated in [7], appropriate buffering can significantly improve video streaming quality. However, too much buffering in mesh-pull systems can lead to an unacceptable delay performance for an IPTV service. *Start-up delay* is the time interval from when one video is selected by a user until actual playback starts on his/her screen. For IPTV applications on the best effort Internet, startup buffering always has been a useful mechanism to deal with the rate variations of IPTV sessions. IPTV applications additionally must deal with peer churn, increasing the requirement for start-up buffering and delay. Although short start-up delay is desirable, a certain amount of start-up delay is required for continuous play-

back. End users also may switch to watch another video from the current one. Before users can watch the new video, the buffering in mesh-pull systems often incurs *video switching delay*. These delays are, of course, significantly longer than those experienced in traditional television. Hence, the state-of-the-art mesh-pull P2P streaming technology does not provide users with the same channel-surfing experience as traditional television.

Another unfortunate characteristic of a mesh-pull P2P streaming system is the possibility of *playback time lags* among peers due to the deployment of the buffering mechanisms. Specifically, some peers watch frames in a video minutes behind other peers. Thus, for the example of a live broadcast of a soccer game, some peers might see a goal a few minutes after other peers. Additionally, peers with large playback lags will not upload useful chunks to peers with smaller lags, decreasing the aggregate uploading capacity of the system.

In addition, because of the real-time nature of IPTV, each media chunk has a playback deadline (which can be different from one peer to another by a few 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 frame and wait for the missing chunk to arrive; or it can *skip* the playback of the frames in the chunk 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 buffer of the player. If there are still chunks available in the buffer of the player, the player continues playback although those chunks might not be continuous; in this case, video playback skipping occurs. In many P2P live streaming systems, when the playback freezes for an extended period of time, the engine terminates the connection with the player and restarts the entire streaming process; we refer to this impairment as *rebooting*.

QUALITY INFERENCE USING BUFFER MAPS

Various parties are interested in monitoring the service quality of IPTV applications. Service providers would like to detect when service quality degrades, so that they can add additional uploading capacity for maintaining a satisfied service level. This information could be provided to users as an aid in selecting P2P video providers. It also could be provided to advertisers who wish to advertise in IPTV systems.

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. However, this approach usually varies among the video sequences and is not suitable for video quality measurement over the Internet, because this approach typically requires extensive direct traffic measurements and hence, incurs high overhead and cost.

As we described earlier, peers in P2P mesh-pull streaming systems advertise buffer maps to each other for exchanging the availability of the

In many P2P live streaming systems, when the playback freezes for an extended period of time, the engine terminates the connection with the player and restarts the entire streaming process; we refer to this impairment as rebooting.

In a measurement study, the number of simultaneous users watching a live broadcast of the annual Spring Festival Gala for the Chinese New Year on January 28, 2006 reached over 200,000 users at a bit rate in the 400–800 kb/s range, corresponding to an aggregate bit rate in the vicinity of 100 Gb/s!

video chunks that they have cached. The information provided in these buffer maps correlates with the playback quality and start-up latency of that peer. Given this correlation, buffer maps can be exploited to measure network-wide quality [6]. In this approach, instead of monitoring video traffic between peers, we capture the exchange of peer buffer maps that summarize the availability of video data at the various peers. These buffer maps are of very small sizes. Therefore, this approach generates significantly less measurement traffic in the network.

As shown in Fig. 3, the offset value in a buffer map indicates the playback position of a peer. The string of zeroes and ones embedded in buffer maps shows the chunks that a peer currently has buffered and can share. When a peer joins the network and has no chunks in its engine cache, it sends buffer maps with a particular playback offset, that is, equal to zero. After a peer obtains chunks from other peers, the offset is set to a value of the playback starting position. The streaming engine will start the player when the size of the cached playable video size exceeds a specified pre-buffering video threshold. The consecutive buffer maps provide snapshots of the video chunks in the cache. Therefore, we can infer the start-up delay of the playback of a peer given the sequence of the buffer maps of this peer. In addition, by tracing the switching of the channel IDs carried in buffer maps, we are able to infer the channel switching delay of one peer. In tracking user playback continuity, the playback freezing events are inferred when the buffer map playable video, determined from the buffer maps, remains at a low level of size over a period of time. Under normal circumstances, the buffer map offset should increase at the video playback rate; hence, by tracking the variability of the offset increasing rate, playback reboot events can be inferred.

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 approaches one. The speed of the growth is a good measure of the P2P video distribution efficiency. The chunk retrieval ratio is a good metric for service providers for making capacity offload decisions. When this ratio is high, service providers may remove their capacity for other purposes; otherwise, service providers should add additional server capacity to sustain a high-level chunk retrieval ratio.

OPTIMIZATION FOR QUALITY-OF-EXPERIENCE

The conventional quality of service (QoS) metrics, such as delay, loss, and so on that are commonly used in network-layer service provisioning are important but not sufficient in evaluating a successful IPTV application. We find that most current mesh-pull IPTV systems incur start-up delays ranging from a few seconds to a couple of minutes, and the playback time lags up to several minutes among peers. Although research efforts were conducted to minimize delays and increase the video bit rate for high-quality video services, IPTV service providers can optimize

the IPTV application to achieve satisfied user experience, namely, quality of experience (QoE).

In addressing noticeable start-up delay, some advertisement videos can be pre-cached on the user side; when a user starts the IPTV application, these advertisement videos immediately are played back. During the playback of these commercials, videos are downloaded in the background. Therefore, users experience zero start-up delay, and the IPTV service provider has an avenue to obtain revenues. Another possibility is to provide multiple videos with different bit rates for the same channel. Dedicated servers are deployed to provide the low-quality video, and the P2P delivery architecture is utilized to transfer the high-quality video. When a user joins the service, the P2P engine downloads the low-quality video from the server to achieve a quick service initiation; then, the engine downloads the high-quality video from other peers. Technology can provide the foundation of the IPTV application; nevertheless, a deep psychological understanding of user behavior and expectations is helpful in optimizing the application to maximize the user viewing experience.

TRAFFIC PRESSURE: TCP VERSUS UDP

The possible tidal wave of IPTV traffic will create a great deal of pressure on the Internet. In a measurement study [2], the number of simultaneous users watching a live broadcast of the annual Spring Festival Gala for the Chinese New Year on January 28, 2006 reached over 200,000 users at a bit rate in the 400–800 kb/s range, corresponding to an aggregate bit rate in the vicinity of 100 Gb/s! In the future, we envision increasing IPTV traffic volume with thousands of streaming channels available on the Internet, each with a bit rate of 500 kb/s or more and each supporting tens to hundreds to thousands of users.

The huge traffic volume generated by IPTV applications is a growing concern among Internet service providers (ISPs) that must support the distribution cost without sufficient incentives. Video chunks are the basic data units in mesh-pull systems. ISPs might consider deploying caching infrastructures to cache these video chunks for assisting the peers within their domains and reducing the inter-ISP IPTV traffic. On the other hand, peers in mesh-pull streaming systems also should turn to more *locality-aware* [8] chunk delivery when selecting peers and video chunks for download.

The transmission of video chunks in mesh-pull streaming systems is carried using TCP or UDP. TCP provides a reliable and congestion-aware transport. However, the selection of the TCP transport introduces various implications on end-hosts, access links, and network cores. Operating systems, in general, have constraints on the maximum rate to accept new TCP connections. When a peer receives a large number of TCP connection requests, there may not be a response for other TCP requests. Other networking applications running on the same host, such as Web browsing, may suffer significantly. Residential users may have additional difficulties because access routers at

home are often of limited capacity and not capable of handling a large number of TCP connections. When one user is watching IPTV, other users using different computers within the same LAN may have poor network connectivity. Recently, we found that mesh-pull streaming systems tend to carry IPTV traffic using UDP instead of TCP. UDP incurs much less connection overhead than TCP; however, this increased usage of UDP raises a concern that this UDP traffic should be congestion-aware to avoid network collapse. Because UDP datagrams may be dropped in the networks, IPTV applications must address how to react to packet loss. The simplest strategy is to ignore the dropped UDP datagrams with the expectation that the viewing quality does not suffer too much in spite of a small amount of packet dropping; however, too much UDP segment loss will severely degrade the video playback performance. It may be beneficial to retransmit some of the lost UDP datagrams with the additional complexity in P2P applications. As a more promising alternative, this streaming traffic can be delivered over the Datagram Congestion Control Protocol (DCCP) [9] to achieve a good trade-off between timeliness and reliability.

SECURITY CONCERNS: ATTACKS AND DEFENSES

The distributed P2P architecture of mesh-pull streaming systems makes them prone to various security threats. A malicious peer in the system may mix video stream with bogus chunks that can significantly degrade the quality of the rendered media at the receivers. This peer can also advertise a large number of non-existing peers who are interested in the same channel; therefore, a legitimate peer may find it difficult to identify other legitimate peers to download video chunks. Due to the low distribution cost of P2P streaming, we expect user-generated live video content (emanating from Web cams and wireless devices) to be distributed using P2P mesh-pull architectures. The origin video server from a user is usually a computer with limited CPU power and network capacity. If malicious peers connect to this server and occupy its bandwidth, without sharing the video chunks with other peers, other peers are not able to enjoy the video at all. A mesh-pull streaming system potentially consists of hundreds of thousands of peers. If malicious peers advertise that one victim host has abundant video chunks, other peers may send chunk requests to this victim host, consuming the CPU power and network bandwidth of this host. As a result, this victim host may undergo denial-of-service (DoS) attacks.

Similar attacks have been studied heavily in P2P file-sharing applications; nevertheless, few attacks have been reported for mesh-pull IPTV systems. Due to the real-time communication in IPTV, the potential attacks on mesh-pull systems can be devastating. In [10], a chunk pollution attack is demonstrated to severely degrade the performance of an IPTV application. In the experiment, before launching the attack, a particular channel had about 3300-plus viewers before the attack; during the attack the number of viewers

dropped to about 500 within 30 minutes. The video quality became unacceptable for a large majority of peers, and they eventually left the system. In defending this attack, chunk signing is an effective mechanism. In chunk signing techniques, the so-called authentication information, or signature, must be transmitted to the receivers along with the chunks. This authentication information can either be provided by the source (in which case, the load on the source might be high) or could be distributed through the P2P system itself, in the form of a separate stream, or be piggybacked with video chunks. A peer receives each chunk and its corresponding signature one by one, verifies its integrity, and plays back (and forward) only if the chunk is valid; otherwise it rejects the chunk as being polluted. In facing the security challenges in mesh-pull streaming systems, researchers and developers should meet the real-time requirements for defense. They may benefit from the lessons in combating the pollution, poisoning, and DoS attacks in P2P file sharing.

ARCHITECTURE RETHINK: TREE-PUSH, MESH-PULL, OR PUSH-PULL?

Mesh-pull systems have enjoyed a number of successful commercial deployments; however, tree-push systems largely have been running at the research stage. Nevertheless, those mesh-pull systems often suffer from long start-up delays, significant video switching delays, and large peer playback time lags. Unlike mesh-pull systems, tree-push systems have smaller delay performance when the tree structure does not break down due to peer churn, and peers at the higher level of the tree have sufficient upload capacity to support the streaming of their children peers. The debate on whether mesh-pull or tree-push is more suitable for IPTV is on-going [11, 12]. The architecture design trade-offs will affect various system characteristics significantly [13], that is, system resilience and signaling overhead, and so on.

In mesh-pull systems, the fundamental task of a peer is to download video chunks quickly. To this end, a peer should first locate its peer neighbors as potential chunk providers; then, identify the chunks to download. Either its neighbors notify this peer of their buffer maps, or this peer inquires about the buffer maps of its neighbors, so that the chunk-availability information propagates among peers quickly. Finally, this peer pulls chunks from its neighbors according to the buffer map information. The amount of the time used in this process often contributes a large portion of the long service delay suffered in mesh-pull systems. Video also can be divided into chunks and be delivered via tree-push systems. Before the chunk delivery starts, the single/multiple tree structure is established, and the chunks are just forwarded from the root of the tree at the video source to the different levels of the tree hierarchy. There is no further communication overhead for the subsequent chunk delivery unless the trees must be repaired due to the departure of intermediate peers. Therefore, the delay performance in tree-push systems is in general good when the trees are stable.

Recently, a new type of push-pull architecture appears to promise to offer a good trade-off

The distributed P2P architecture of mesh-pull streaming systems makes them prone to various security threats. A malicious peer in the system may mix video stream with bogus chunks that can significantly degrade the quality of the rendered media at the receivers.

The increasing user base of P2P IPTV systems will surely attract more malicious attacks. It is urgent to address various vulnerabilities in current designs to improve the security of future P2P IPTV systems.

between the peer/chunk selection and scheduling overhead and the delay performance [12, 14]. The basic idea is to construct the chunk distribution trees between those peers that are stable and of high upload capacity — or super peers — so that video chunks can be pushed from the video source to these super peers quickly. Those ordinary peers, however, pull chunks from the super peers and other ordinary peers.

In this push-pull scenario, super peers contribute more upload capacity than ordinary peers; however, all the peers enjoy the same video playback quality. It is questionable what the incentives are for one peer to become a super peer. One way to encourage the contribution of peers is to provide differentiated service in IPTV applications. In [15], layered video can be used to provide such incentives among streaming peers. Ordinary peers download only a basic-layer video to receive the basic video playback quality; however, super peers download multiple-layer video to enjoy an enhanced video quality.

CONCLUSION

The current practice of mesh-pull P2P streaming systems demonstrates the feasibility of large-scale application layer multicast on top of the best-effort Internet. In this article, we provided an overview of existing mesh-pull systems. We briefly discussed several important design issues regarding performance, reliability, and security of mesh-pull systems. Despite its early successes, P2P IPTV is still at an early stage. Many open and interesting research problems remain to be addressed to design and deploy the next generation P2P IPTV services with superior user QoE. P2P IPTV systems should be enhanced with appropriate server infrastructure support to achieve a high level of QoS. On the other hand, the popularity of P2P IPTV applications on the Internet has posed a great challenge for network infrastructures to support tremendous P2P IPTV traffic. Therefore, it is critical to design new P2P IPTV systems in an ISP-friendly fashion to minimize the traffic stress that they impose on ISP networks. In addition, the increasing user base of P2P IPTV systems will surely attract more malicious attacks. It is urgent to address various vulnerabilities in current designs to improve the security of future P2P IPTV systems.

ACKNOWLEDGMENT

This work was supported by the National Science Council under Grant ITR-0325726 and Grant CNS-0519998.

REFERENCES

- [1] X. Zhang *et al.*, "DONet/CoolStreaming: A Data-Driven Overlay Network for Peer-to-Peer Live Media Streaming," *IEEE INFOCOM*, vol. 3, Mar. 2005, pp. 2102–11.
- [2] X. Hei *et al.*, "A Measurement Study of a Large-Scale P2P IPTV System," *IEEE Trans. Multimedia*, vol. 9, no. 8, Dec. 2007.
- [3] Y. Tang *et al.*, "Deploying P2P Networks for Large-Scale Live Video-Streaming Service," *IEEE Commun. Mag.*, June 2007.
- [4] A. Sentinelli *et al.*, "Will IPTV Ride the Peer-to-Peer Stream?" *IEEE Commun. Mag.*, June 2007.
- [5] J. Liu *et al.*, "Opportunities and Challenges of Peer-to-Peer Internet Video Broadcast," *Proc. IEEE*, 2007.

- [6] X. Hei, Y. Liu, and K. W. Ross, "Inferring Network-Wide Quality in P2P Live Streaming Systems," *IEEE JSAC*, vol. 25, no. 10, Dec. 2007.
- [7] R. Kumar, Y. Liu, and K. W. Ross, "Stochastic Fluid Theory for P2P Streaming Systems," *IEEE INFOCOM*, 2007.
- [8] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?" *ACM IMC*, 2005.
- [9] E. Kohler, M. Handley, and S. Floyd, "Datagram Congestion Control Protocol (DCCP)," RFC 4340, Mar. 2006.
- [10] P. Dhungel *et al.*, "The Pollution Attack in P2P Live Video Streaming Systems: Measurement Results and Defenses," *Sigcomm P2P-TV Wksp.*, 2007.
- [11] N. Magharei, R. Rejaie, and Y. Guo, "Mesh or Multiple-Tree: A Comparative Study of Live P2P Streaming Approaches," *IEEE INFOCOM*, 2007.
- [12] M. Zhang *et al.*, "Understanding the Power of Pull-Based Streaming Protocol: Can We Do Better?" *IEEE JSAC*, 2007.
- [13] V. Fodor and G. Dan, "Resilience in Live Peer-to-Peer Streaming," *IEEE Commun. Mag.*, June 2007.
- [14] F. Wang, Y. Xiong, and J. Liu, "mTreebone: A Hybrid Tree/Mesh Overlay for Application-Layer Live Video Multicast," *IEEE ICDCS*, 2007.
- [15] Z. Liu *et al.*, "Using Layered Video to Provide Incentives in P2P Streaming," *Sigcomm P2P-TV Wksp.*, 2007.

BIOGRAPHIES

XIAOJUN HEI [S'02] (heixj@poly.edu) received his B.Eng. degree in information engineering from Huazhong University of Science and Technology, China, in 1998. He obtained his M.Phil. degree in electrical and electronic engineering from Hong Kong University of Science and Technology in 2000. He is currently working toward his Ph.D. degree in the Department of Electronic and Computer Engineering at Hong Kong University of Science and Technology. From September 2005 to September 2007 he conducted a research visit on P2P networking at the Department of Computer and Information Science, Polytechnic University. His research interests include P2P networking, network modeling, and measurement.

YONG LIU [M'02] (yongliu@poly.edu) received his Ph.D. degree from the 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. He has been an assistant professor in the Electrical and Computer Engineering Department of Polytechnic University since March 2005. His general research interests are 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 ACM.

KEITH W. ROSS [S'82, M'86, SM'90, F'06] (ross@poly.edu) received a B.S.E.E from Tufts University, an M.S.E.E. from Columbia University, and a Ph.D. in computer and control engineering from the University of Michigan. He joined Polytechnic University as the Leonard J. Shustek Chair Professor in Computer Science in January 2003. Before joining Polytechnic University, he was a professor for five years in the Multimedia Communications Department at Eurecom Institute, Sophia Antipolis, France. From 1985 through 1997 he was a professor in the Department of Systems Engineering at the University of Pennsylvania. He has worked in peer-to-peer networking, Internet measurement, video streaming, Web caching, multiservice loss networks, content distribution networks, network security, voice over IP, optimization, queuing theory, and Markov decision processes. He is currently an Associate Editor of *IEEE/ACM Transactions on Networking*. He has served as an advisor to the Federal Trade Commission on P2P file sharing. He is coauthor (with James F. Kurose) of the popular textbook, *Computer Networking: A Top-Down Approach Featuring the Internet* (Addison-Wesley, 4th ed. 2007). The text is used by over 200 U.S. universities each academic year, is widely used internationally, and has been translated into 12 languages. He is also the author of the research monograph *Multiservice Loss Models for Broadband Communication Networks* (Springer, 1995). From 1999 to 2001 he took a leave of absence to found and lead Wimba, an Internet technology startup. Wimba develops and markets Java-based asynchronous and synchronous VoIP technologies, primarily for the online education and language learning markets.