Performance and Quality-of-Service Analysis of a Live P2P Video Multicast Session on the Internet

Sachin Agarwal Jatinder Pal Singh
Deutsche Telekom A.G., Laboratories

Ernst-Reuter Platz 7,
10587 Berlin, Germany
Email: {sachin.agarwal, jatinder.singh}@telekom.de

Aditya Mavlankar Pierpaolo Baccichet Bernd Girod
Department of Electrical Engineering
Stanford University
350 Serra Mall
Stanford, CA 94305, USA
Email: {maditya, bacci, bgirod}@stanford.edu

***Abstract*—We evaluate the performance of a large-scale live P2P video multicast session comprising more than $120,000$ peers on the Internet. Our analysis highlights P2P video multicast characteristics such as high bandwidth requirements, high peer churn, low peer persistence in the P2P multicast system, significant variance in the media stream quality delivered to peers, relatively large channel start times, and flash crowd effects of popular video content. Our analysis also indicates that peers are widely spread across the IP address space, spanning dozens of countries and hundreds of ISPs and Internet ASes. The analysis yields several QoS measures such as fraction of stream blocks correctly received, number of consecutive stream blocks lost, and channel startup time across peers. We correlate the observed quality with the network and peer behavior, suggesting several avenues for optimization and research in P2P video multicast systems.**

I. Introduction

Peer-to-peer (P2P) live and on-demand video multicast has become a notable Internet application offering diverse video content to millions of viewers. Many P2P multicast system providers develop P2P software clients for users who want to connect to these P2P video multicast systems. Popular commercial services available on the Internet include End System Multicast (ESM), Joost, Gridmedia, PPStream, PPLive, Zattoo, Vudo, Octoshape, Sopcast, Tvkoo, Roxbeam, Tribler, etc. Many service providers have licensing agreements with content providers for providing engaging content, with the result that some P2P video multicast-supported content delivery systems have become popular with a large number of users.

This popularity of P2P video multicast calls for a detailed study of the delivered quality-of-service of large-scale, representative P2P video multicast systems in order to quantify performance and to compare them with other Internet video delivery platforms like IP-multicast and traditional Akamai-like [1] content delivery networks.

Almost all P2P video multicast systems operate over the best-effort public Internet. The streaming bandwidth is usually provided through using the uplink bandwidth of participating peers to serve video content to other peers. However, peer churn, *i.e.*, peers joining and leaving the system, and varying network conditions such as bandwidth, delay and delay jitter present challenges to delivering high quality video streams to peers participating in the multicast session. Many P2P multicast systems are designed to counter these challenges through buffering, augmenting available bandwidth via seeds or super nodes, and redundant downloading of the stream at the peers.

In this paper, we analyze logs from a large-scale video multicast session (comprising approximately 120,000 distinct IP addresses) conducted via a commercial P2P multicast system. The analysis yields several quantitative measures of the analyzed P2P video streaming multicast session. We also map the peers' IP address information into Internet BGP routing logs from the day of the multicast session and present several measurements that quantify the effects of large-scale P2P video multicast traffic on Internet service providers (ISPs). In addition, we also present quality measurements in terms of blocks of the video stream lost and channel startup times.

We dwell upon explaining some of the quality-related metrics in light of the network and user behavior analysis. Our analysis may be useful to P2P system providers and to ISPs for better understanding the issues and the improvements possible in large scale P2P video multicast applications. In addition, the analysis points to various research directions for improving delivered media quality and efficiency of P2P video multicast protocols.

A. Key Results

We state the key results of our analysis below in order to clarify the focus and contributions of this work.

- P2P video multicast is a bandwidth-intensive application. The peak aggregate bandwidth usage of the analyzed system was about 45 Gbps for a relatively modest 60,000 peers in the system.
- Peer behavior was highly dynamic, particularly at the end of the video multicast, with 100s of peers joining and leaving the system every second. We also report a low peer persistence - the median time for which peers stayed in the system was just 106 seconds for the 4-hour video multicast session.
- The total bandwidth contribution of peers fell short of the total required bandwidth, the difference being augmented by the system provider's bandwidth-injecting

server farms. Our analysis indicated that the injected bandwidth did not exceed 15% of the required bandwidth. In addition, the majority of peers were net receivers of bandwidth, receiving more than they contributed to the system.
- Most peers connected from behind Network Address Translation (NAT) devices; this corresponds to home broadband routers with NAT functionality. Effective NAT translation needs to be incorporated into P2P clients in order for them to participate successfully in the P2P system and to improve the delivered quality by making their upload bandwidth available to the system.
- The peers were geographically distributed in 100s of ISP and university networks: 51 countries were represented and 731 Internet autonomous systems (ASes) [2] were discovered in our analysis. This diversity calls for sophisticated QoS mechanisms to ensure acceptable delivered quality for users. Intelligent network-aware resource allocation within the P2P system would also help in improving delivered quality across the system.
- Half the peers correctly received 95% of the media stream, although there were bursts of consecutive stream losses on peers that may have resulted in lower media quality with frequent video frame-freezes and audio channel disturbances. Advanced audio and video codecs, such as H.264 [3], may be able to partially offset these high losses.
- There was significant variance in the fraction of correctly delivered stream (and hence, delivered quality) with different Internet ASes. This indicates that uniform QoS across the peers was not being achieved, and that P2P video multicast needs to be optimized further in order to address the network quality heterogeneity issue.
- Similarly, there is a large variance in the channel startup time across different ASes.
- The average channel startup time across all peers was 32 seconds, significantly slower than IP-multicast. This unusually high channel start-time is a significant challenge for live P2P video multicast and compares poorly to IP multicast.

B. Organization

The paper is organized as follows. In Section II we review related work in P2P systems, emphasizing recent work in P2P video multicast measurements. We present some background information about the analyzed P2P system and our analysis methodology in Section III. Section III presents the results of our analysis of the P2P multicast system in terms of the network characteristics of the system and several QoS related measurements. Section V summarizes our findings. Acknowledgements are provided at the end of this section.

II. Related Work

P2P systems are widely used to distribute audio and video content in the Internet, and it is estimated that a major portion of the bandwidth available on the consumer ISP networks carries P2P traffic [4, 5]. It has been shown through analysis [6, 7], simulations, and measurements [8–10] that the P2P content delivery model scales gracefully with user demands in heterogeneous P2P networks. A majority of popular P2P systems are built around file sharing applications. These applications typically distribute stored and not live content, and the downloaded content is played back only after the entire file has been downloaded. P2P live multicast applications, such as the one analyzed in this paper, are significantly more demanding in QoS because of their real-time content delivery requirements.

A number of P2P streaming multicast solutions are based around setting up an overlay tree with the content source being the root and terminals being the other nodes of the tree [11–13]. This approach has the disadvantage of limited robustness against peer disconnections that can disrupt stream reception on downstream neighbors in the overlay tree. Many techniques [14] for alleviating this problem based on redundant data paths, multiple overlay trees, and multiple description coding [15] have been proposed. An alternative approach of using mesh overlays is also popular and many P2P systems such as SopCast [16], PPLive [17], Coolstreaming [18, 19], Roxbeam [20], and Gridmedia [21] use mesh-based overlays. We refer the interested reader to a good survey of P2P multicast in [22].

The traffic characteristics of large scale P2P multicast systems have been a topic of interest since the first P2P multicast systems [11–13, 23]. A survey and comparison of the approaches and algorithms employed in various P2P multicast network overlays can be found in [24, 25]. There have also been recent studies of commercial P2P multicast systems (for example, [26–32]) that study networking characteristics of some commercial P2P systems such as SopCast [16, 33], PPLive [17], Coolstreaming [18, 19], Roxbeam [20], and Gridmedia [21]. The analysis in [34] was based on the same log data and may be considered a small subset of this work because it did not present results on peer churn dynamics, classification of the underlying IP network autonomous systems (ASes), the relation between delivered quality (stream quality and startup times) and AS, peer bandwidth contributions, type of Internet connections, geographical analysis of IP addresses, etc.

The key distinction of our analysis from the papers cited above is that we provide system-wide characteristics of a live, large-scale P2P multicast session, including the quality of the stream delivered to users. Our analysis benefits from the centralized logging facility in the system that records logs from each participating peer. This gives us a birds-eye view of the P2P system from the provider's point of view instead of being limited to only network measurements or end-user client measurements. We also combine information from Internet BGP routing logs collected on the day of the multicast session to highlight several network-related issues of the P2P multicast session.

III. Preliminaries

We analyze the large scale Internet usage characteristics of a P2P streaming system in Section IV. For this purpose, we use logs obtained from the P2P system provider for a 4 hour baseball game played in a developed country and streamed to over 120,000 peers all over the world. Commercial considerations and agreements restrain us from divulging the name of the particular P2P system provider and details of the underlying algorithms and protocols used by the particular provider, apart from mentioning that the P2P overlay network has a mesh-based architecture and that the P2P system provider injects additional bandwidth into the system to improve the QoS, as discussed in Section IV-A. We emphasize that the session was an instance of a pre-planned multicast and not the streaming of some user-generated content that had suddenly become popular. This prior information of the event allowed the P2P system provider to preemptively allocate additional bandwidth and seeds/supernodes in the overlay network in anticipation of the large number of peers.

The streamed content comprised of a constant bit rate (CBR) windows-media encoded stream at 759 kbps (64 kbps audio + 695 kbps video), encoded at 29 frames per second and 640×480 pixels (VGA) video resolution - a resolution comparable to standard definition TV. The stream was delivered to the P2P clients via the P2P overlay network of the P2P system provider. The software client of this P2P streaming system was available to users as a web-browser plug-in, allowing them to easily download and integrate the peer client into their web-browsers. The client used windows media player (WMP) software embedded in the web-browser to decode and playback the received video.

Each P2P client logged data to a centralized logging server at the P2P system provider's control center. The information presented in this work was extracted from these logs. The log data was saved as a single file with each client-reported information string indexed by the UTC time of receipt at the logging server and the IP address of the client. One of the key challenges in analyzing the data was to re-build and aggregate the performance of the whole P2P system on the basis of strings reported by individual P2P clients over the run-time of the P2P video multicast session.

Each client-reported information string included a unique identifier of the client because using the IP address as a unique identifier may be inaccurate given that the DHCP [35] IP addresses of clients may change during the course of the multicast session. The information string, periodically reported to the logging server every 5 minutes by each client, contained information about the state of the media buffer in the P2P client, total bytes sent and received by the client, running time of the client, channel id, connection firewall type, recently lost packets, etc. In addition, special events such as client start and stop conditions were reported to the logging server as soon as they occurred.

We extracted several relevant parameters of the P2P streaming system from the logs, such as peer dynamics, number of clients, temporal variation of aggregate upload and download bandwidth, peer persistence in the system, net bandwidth contributing peers versus net bandwidth receiving peers, and the type of NAT/firewalls. In addition, we parsed the IP address data of the clients and report on their geographical and Internet-wide spread. We also extracted a classification of the underlying ASes taking part in the overlay network based on BGP [36] routing information and derive the correlation between a peers' AS and its received QoS.

We also derived several QoS measures from the log data, such as the video stream available to the media decoder of clients as a percentage of the required video stream, channel startup time of the clients, and average number of consecutive stream blocks lost by clients. Then, we correlated this QoS information to the characteristics of the underlying IP network and ASes. This linkage highlights some of the issues and the improvements possible in P2P video multicast if the underlying IP networks are taken into account while designing and running P2P video multicast systems.

IV. Data Analysis

We first present the peer dynamics, bandwidth usage and other system-wide characteristics in Section IV-A. Then, we analyze the IP addresses and underlying IP networks of the participant peers in Section IV-B. The quality of multimedia stream delivered in terms of fraction of the stream correctly received and number of consecutive stream blocks lost are presented in Section IV-C, where we also present the channel startup time analysis for all the peer clients in the system.

A. Peer Behavior and Network Impact

Fig. 1 shows the number of peers in the streaming system for an 11 hour interval including the 4 hour long streamed baseball game. The rate of users joining and leaving the system is also plotted. The steep rise in the number of peers at hour 4 marks the start of the baseball game that ends just before hour 8 as characterized by the sharp drop in the number of peers in the system before hour 8. These measurements indicate that live P2P streaming system architects need to design their protocols to withstand high peer churn at the beginning and end of video programs. Moreover, ISPs can expect to see such large and rapid fluctuations in network usage corresponding to video start and end events. We can also expect to see "flash crowd" effects when content suddenly becomes very popular for relatively short durations of time (such as the end of the baseball game).

As peers continued joining and leaving the multicast session during the game, the number of peers kept increasing to a maximum of about 60,000 concurrent peers. There were steep rises in the joining and leaving rate of peers toward the end of the game because people joined to see the climax and then stoped watching the game immediately thereafter. The maximum joining and leaving rates were 80 and 328 peers per second respectively. Some users stayed beyond the end of the game, this corresponds to people who do not close their web-browser after the game, thus leaving the P2P streaming

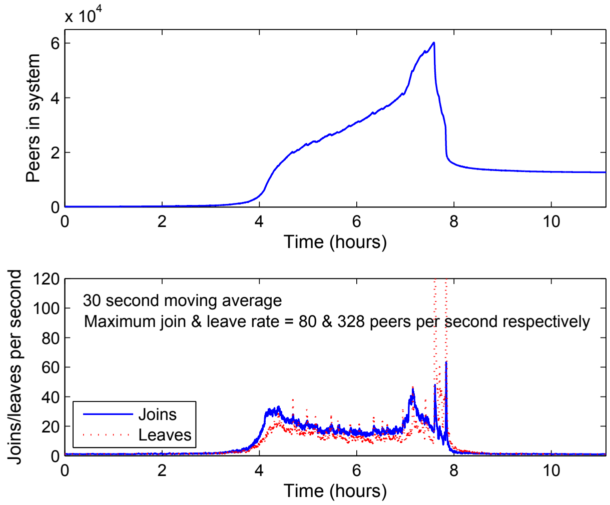

Fig. 1. Peer churn: The total number of peers in the system vs. time. Peers joining and leaving rates during the course of the baseball game is also plotted. The highest peer-churn of 328 peer leaves/second occurred at the end of the game.

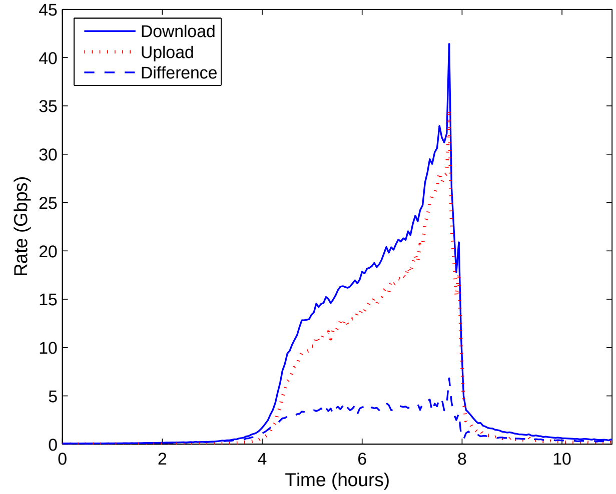

Fig. 2. Network bandwidth load: The aggregate download and upload data-rates summed over all peers in the P2P streaming system are plotted as a function of time. The difference between them is the bandwidth supplied by the system provider/content provider using server farms or super-nodes to inject bandwidth into the system.

client active on their computer, which periodically reported back to the logging server.

Fig. 2 shows the aggregate peer bandwidth being used in the P2P streaming system. Two important features are noticeable. First, the sheer volume of network bandwidth required - the aggregate download rate shot to over 42 Gbps at the end of the game and the aggregate upload rate from peers rose steeply as well. From Fig. 1 the total number of peers just before the end of the game was about 60000, and given the bitrate of the delivered video stream (759 kbps), this translates to about 45 Gbps. Second, the upload and download rates rise and fall steeply at the beginning and the end of the game respectively, corresponding to the number of peers in the system as shown in Fig. 1. In this discussion it is worthwhile to note that the peer clients in the access (edge) network served the video stream to other peer clients. This presents a challenge to ISPs because scaling up the access network is usually very expensive.

A key observation in Fig. 2 is that the peers by themselves do not provide all the bandwidth they consume because the total upload rate is consistently lower than the total download rate. This difference is also plotted in the Figure. In our discussions with the P2P system provider, we learnt that this shortfall is compensated by the system provider through "server farms" that inject bandwidth to make up for the difference in aggregate upload and download bandwidth. Still, the lions share of the bandwidth is provided by peers. It is absolutely remarkable that the amount of provider-injected bandwidth remains almost constant for increasing number of users during the multicast session, suggesting that P2P overlays are scalable and cost-effective video delivery platforms.

We observed that the persistence of many peers in the system is quite short as many of them choose to join and leave the system at short intervals. Table I shows the most important statistics of this behavior. Half of the peers stayed for 106 seconds or less in the system. This characteristic is also one factor responsible for the peer churn of Fig. 1. We also observed that many peers joined, left, and re-joined the P2P streaming system multiple times.

TABLE I
STATISTICS OF PEER PERSISTENCE IN THE SYSTEM (SECONDS). THE STREAMED CONTENT WAS ABOUT 4 HOURS, OR 14400 SECONDS IN LENGTH.

Minimum	1
1st Quartile	18
Median	106
Mean	1307
3rd Quartile	1205
Maximum	68738

Fig. 3 plots the histogram of the logarithm of the ratio of number of bytes downloaded to the number of bytes uploaded by a peer. Most peers download more than they upload. However, since many peers only stay connected to the overlay network for a brief time, a net contributing (receiving) peer of Fig. 3 may not contribute (receive) much in terms of actual number of bytes. The obvious asymmetry in number of net downloading and uploading peers calls for building incentives in the P2P algorithms for peers to contribute bandwidth more willingly.

A likely explanation for the lower upload from most peers is that many peer clients were connected through asymmetric broadband connections with limited upload capacity. In some cases the existence of firewalls and NAT routers may prevent users from actively contributing bandwidth. Table II shows the classification of peer IP connections based on the type of Internet connections. Just under 9% of the users have direct connections to the Internet. A majority of users are behind

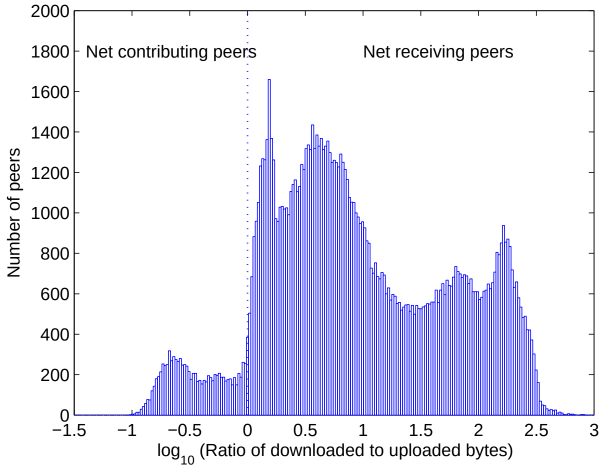

Fig. 3. Histogram of the logarithm of the ratio of bytes downloaded to bytes uploaded by peer clients. Most peer clients receive more than they transmit; these results may also indicate the need for bandwidth-sharing fairness algorithms in P2P streaming.

TABLE II
STATISTICS OF THE TYPE OF INTERNET CONNECTIONS (PERCENT). MOST PEERS WERE BEHIND A NAT; EFFECTIVE P2P STREAMING SYSTEMS NEED TO HAVE EFFECTIVE NAT TRAVERSAL CAPABILITIES.

Direct connections	8.98
UPNP	6.53
Firewall	6.80
NAT	70.24
Other/Unknown	7.44

NAT devices; these are mostly commercial broadband Internet users behind home routers that implement NAT.

B. Underlying IP Network Considerations

The log data recorded the IP addresses of all clients participating in the P2P video multicast session. We identified $120,275$ unique peer IP addresses in the logs. Table III shows the characteristics of the IP addresses of peers in the system. A high degree of similarity of the IP address prefix of any two IP addresses is often a good indicator of the two belonging to the same ISP. We observed 118 (of 255 possible) different most significant 8-bit IP address prefixes, 1000s of 16-bit address prefixes, and 10s of thousands of 16-bit IP address prefixes, indicating that peers from many different IP networks participated in the multicast session. This diversity also highlights the advantage of P2P multicast over IP-multicast in the former being able to traverse multiple Autonomous Systems (ASes) of the Internet and serve niche content to geographically distributed users.

A "whois" lookup on each of the IP addresses showed that peers from 637 different ISPs in 51 countries connected to the P2P video multicast session. We conjecture that this large spread in IP addresses in multiple ASes will become more pronounced when more globally popular content is streamed through P2P streaming systems. Hence, ISPs and

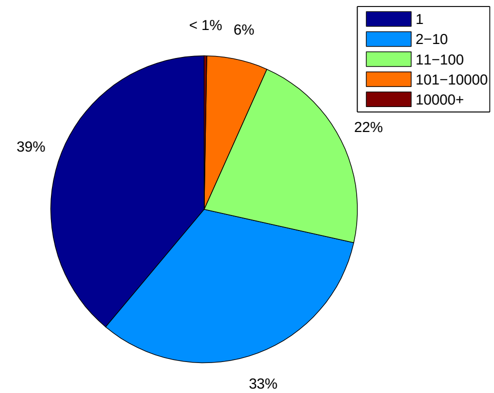

Fig. 4. Peer IP addresses per AS. Although just two ASes accounted for more than half of all peer IP addresses, many other ASes contained very few IP addresses. This sparse peer population in many ASes results in significant inter-AS traffic.

TABLE III
IPV4 ADDRESS DISTRIBUTION IN TERMS OF IP-SUBNET MASK PREFIXES.

Unique Prefix	8 bits	16 bits	24 bits
Total	118	6314	69523
Statistics of IP addresses per prefix			
Minimum	1	1	1
1st Quartile	2.25	1	1
Median	9	2	1
Mean	1019.36	19.05	1.73
3rd Quartile	85	9	2
Maximum	17666	1017	78

P2P streaming systems could mutually benefit by sharing information about the underlying mapping of the IP network onto the physical network and traffic localization of P2P streaming algorithms, as has been proposed in [4].

We sorted the IP addresses by entities, or the autonomous systems (ASes) [2], which owned these IP addresses. We matched the IP addresses of the peers with the data from the Cymru service [37] to obtain this classification. Our analysis showed that 731 ASes were represented in the data; Fig. 4 shows the distribution of the number of IP addresses connecting to the P2P system in ASes. Two ASes located in the country of origin of the baseball game contained more than half (63563) of the IP addresses in the system. But most other ASes had very few peers, in fact 72% of the ASes had 10 or fewer peers. This is significant because the lack of peer neighbors in the same AS results in higher inter-AS P2P traffic as peers upload (download) the video stream to (from) peers in other ASes.

We then used the AS classification tool from CAIDA [38, 39] to report Table IV about the types of ASes in the logs. The classification of ASes is based upon data extracted from Internet routing registries. The classifier [40] uses information such as the AS number, the organization description record, the

TABLE IV
CLASSIFICATION OF THE 731 PARTICIPATING AUTONOMOUS SYSTEMS (ASES) BASED ON [38, 40].

AS type	Number found
T-2 (Small ISPs)	314
COMP (Customers)	90
NIC (National Information Centers)	87
EDU (Universities)	55
T-1 (Large ISPs)	20
IX (Internet Exchange Points)	1
Not Classified	198

number of inferred customers, and the number of /24 prefixes covering the advertised IP space to infer the type of AS.

The important message from the AS classification shown in Table IV is the heterogeneity in the type of ASes. Although we do not have individual peer-to-peer traffic information, we conjecture from our analysis that the P2P video multicast session creates high volume streaming traffic that traverses different types of ASes, paricularly since most peer clients had very few peer neighbors in their own AS.

This may have economic consequences for ISPs, for example, heavy traffic from a 'customer' AS that buys Internet connectivity increases the operating costs for the ISP. Similarly, many small ISPs buy whole-sale bandwidth from large ISPs and pay for the usage in terms of the volume of data transmitted. Large amounts of P2P streaming data traversing their networks will correspondingly increase their costs. In addition, many ISPs presently have reciprocal peering agreements [41, 42] with their peers for transmitting each other's traffic. Changes in the traffic patterns due to P2P video multicast may trigger the re-negotiation of such peering agreements.

C. Quality of Service

1) Stream Quality: The P2P system delivers blocks of the video stream to peer clients who assemble these blocks into the video stream and serve them to the media decoder. The logs provide video quality information about the first 5 minutes video quality reported by each peer client. In particular, each peer reports the number and the order of stream blocks correctly served to the media decoder (in this case, Windows Media Player - WMP). In order for a block to be correctly served to the media decoder, it should arrive before the playout deadline of frames contained in it, and it should be free of errors. In this analysis, we assume that partially received blocks are not decodable by the media decoder.

The ratio of the number of stream blocks correctly served to the number of blocks required for perfect decoding is indicative of the percentage of video stream correctly downloaded by the host. Each block's playout length is of the order of at least 1 second, and therefore, specific information about which type of video frames are lost is less important given that each block contains multiple frames of different types (recall that the frame rate of the transmitted video is 29 fps).

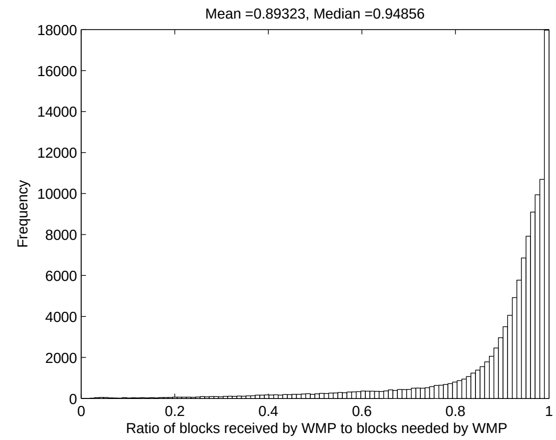

Fig. 5. Histogram of ratio of blocks received by the media decoder to blocks needed for perfect playback for the first 5 minutes on peers. A higher ratio corresponds to higher user-perceived video quality with fewer frame-freezes.

TABLE V
STATISTICS OF THE VIDEO QUALITY RATIO FOR THE FIRST 5 MINUTES AT THE PEERS (1 IS BEST).

Minimum	0.02
1st Quartile	0.88
Median	0.95
Mean	0.89
3rd Quartile	0.98
Maximum	1

Fig. 5 shows the histogram of the ratio of blocks received to the total blocks needed for perfect decoding and Table V shows the statistics of this video quality ratio. The median value of 0.95 indicates that half of the peers correctly receive atleast 95% of the required stream blocks. If the missing blocks were uniformly distributed and each block contained just a few frames then 95% reception could be considered favorably. Howev errors in block reception at a peer client were not uniformly spread over the incoming stream and each block contained at least 1 second of video (hence, atleast 29 frames). We therefore analyzed the length of consecutive missing blocks on the peer clients and present our findings below.

Many video decoders use "copy previous" error concealment to hide missing frames in the video stream from users. This means that in the event of not receiving a certain frame, the last correctly rendered frame is displayed on the screen, resulting in the frame-freezes that are sometimes seen in video playback. The effectiveness of this error concealment quickly decreases with consecutive frames because the correlation in the video information of frames that are apart from each other is usually small. Hence, long bursts of lost blocks are very detrimental to overall video quality and so measuring their length is an important metric of delivered video quality.

Fig. 6 shows the histogram of the average number of consecutive blocks lost during a disruption on the peers. There

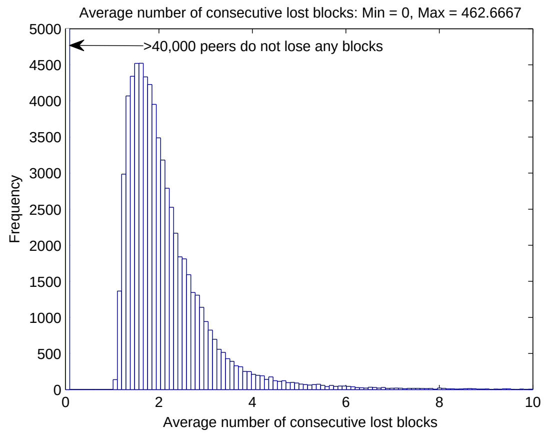

Fig. 6. Histogram of the average number of consecutive blocks never delivered to peers during the first five minutes.

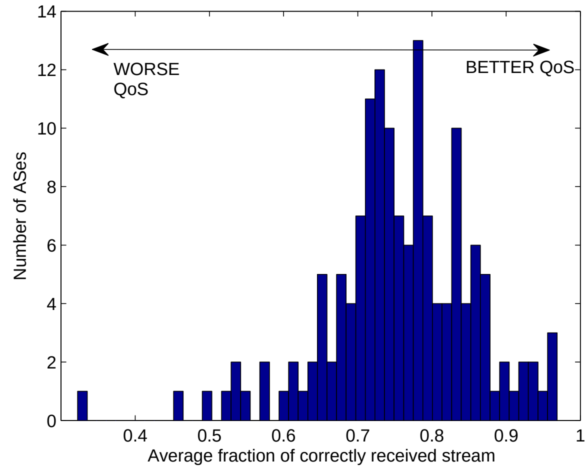

Fig. 7. Histogram of the average fraction of correctly delivered stream blocks to ASes. Note that a large number of peer clients were in ASes that had a large average fraction of correctly delivered stream blocks, as indicated by the relatively large median value of 0.95 for the video quality ratio in Table V.

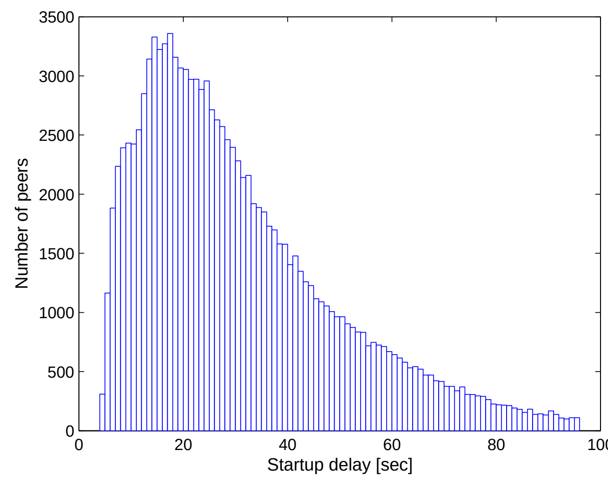

Fig. 8. Histogram of the "channel startup time": the time a user has to wait before video playback starts. The average channel startup time was 33 seconds.

are two important insights. First, there is a significant number of peers that lose multiple consecutive blocks. Second, the video quality in terms of the length of frame freezes differs widely across the peers.

An important point to note about the 40000 peers in Fig. 6 marked as not having lost any blocks is that 18650 of these peer clients receive no stream at all and hence, do not lose any blocks. We believe that this lack of reception of blocks on peers may be due to the inability to successfully connect to the overlay network due to firewall or connectivity issues and/or due to the low peer persistence as shown in Table I.

We classified the ASes according to the average fraction of correctly received blocks on the peers in the ASes with the aim of studying the variation in delivered quality across different ASes. In this analysis, we ignored those ASes whose peer clients did not receive any part of the video stream, leaving us with 478 ASes (out of the total 731 ASes). In addition, because the fraction of correctly received blocks also depends on several peer client factors such as the type of Internet connection, peer client hardware, etc., we only considered the 149 ASes with 10 or more peers, and averaged the fraction of correctly received blocks on the peers in these 149 ASes.

Fig. 7 reports a large variance in the delivered video quality across the ASes, and indicates that the delivered quality to a peer is highly dependent on the AS of the peer. P2P video streaming system providers may need to optimize their algorithms in order to improve video quality to different ASes - for example - by placing bandwidth injection server farms closer to ASes with low average fractions of correctly delivered stream blocks in order to improve their the delivered QoS. The ultimate goal of P2P video streaming applications - having an effective, scalable, distributed, application layer content delivery protocol to serve niche content across the whole Internet - will depend on whether these quality problems can be satisfactorily addressed.

2) Channel Startup Time: Another important quality parameters for Internet video streaming is the amount of time it takes to buffer the stream before video playback. This time is the aggregate delay in connecting to the P2P streaming system, contacting peers, receiving video, and buffering the video stream until it is ready for playback on the video media decoder. We plot the histogram of the channel startup time of the peer cliens in Fig. 8. The minimum, median, and average startup times are 4, 27, and 32.55 seconds respectively. Fig. 9 reports a large variance in the channel startup times across ASes, further highlighting the difference in the quality of user experience in terms of channel startup times across ASes.

The channel startup time can be attributed to minimum buffering requirements of media decoding software such as Windows Media Player. The variation in startup time can be attributed to different network conditions and streaming protocol dynamics. In addition, most of the peers do not have direct IP connections to the Internet but are behind firewalls and NAT routers as noted earlier in Table II. These

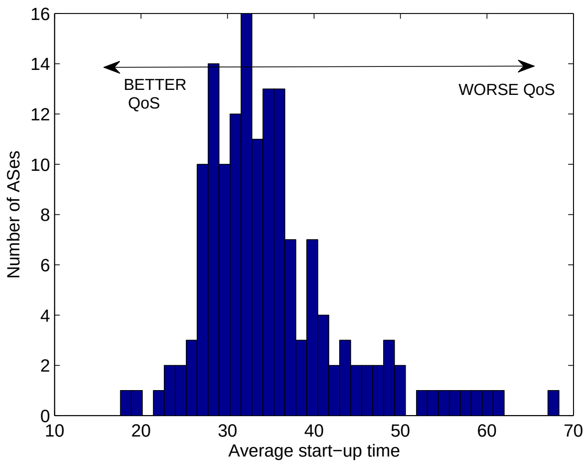

Fig. 9. Histogram of the average channel startup time on peers in ASes. Note the large variance in the channel startup time across ASes.

heterogeneous characteristics of the Internet connections add further variability to the startup time.

We emphasize that the analyzed P2P system was streaming a live event to the peer clients and so preemptive caching of content on clients was not possible. In case of pre-recorded content, the first few minutes of video content can be pro-actively cached on peer clients for a much quicker channel startup time on the peers.

V. Conclusion and Future Work

P2P video streaming is now a mainstream Internet application with several commercial grade systems having large user bases. The traffic characteristics of P2P streaming systems are significantly different from other P2P applications given the real-time requirements for video streaming. Our analysis of a large scale commercial P2P streaming system in this paper reveals these network characteristics such as large aggregate network load, high peer churn, flash-crowd like effects when peers rapidly join the P2P streaming system, and end-game effects when peers leave the system en mass as soon as the live content becomes uninteresting or ends.

Our analysis indicates that large scale P2P video streaming systems may need to explicitly augment bandwidth because the total bandwidth that peers contribute (upload) is often less than the bandwidth used by peers (download). We report a very large traffic foot-print of the P2P video multicast session at the network edge (access network). Moreover, high churn and low peer persistence in the system are inherent characteristics of P2P video streaming, perhaps more so than P2P file-sharing because peers usually stay connected to download the whole data file in the latter case.

Another interesting aspect we report is that the P2P clients were well separated in the Internet IP address space spanning 100s of ISPs and Internet ASes. This may fuel the debate about the effects of large scale P2P video multicast on ISP networks. Quality measurements indicate that several peers experienced slow channel startup times of the order of dozens of seconds and that the quality of the delivered video stream varied widely across the peers participating in the system. The dependence of received video quality and channel startup time on the AS clearly indicates that effective P2P video streaming systems cannot be agnostic of the underlying networks, instead, ISPs and P2P video streaming providers could work together to improve delivered quality.

The less-than-ideal delivered quality observed in this work points to several research and development opportunities. We enlist some of these below.

- Creation of P2P video multicast systems that are resilient to peers joining and leaving frequently.
- Efficient placement of provider-injected bandwidth servers in carefully chosen parts of the Internet to assist poorly performing ASes.
- Inclusion of incentive mechanisms in the P2P algorithms for inducing peers to share more upload bandwidth and improve the QoS delivered to other peers, especially in light of the large variance in delivered video quality between ASes, as seen in our analysis.
- Building effective NAT and fire-wall traversal algorithms and services into peer clients.
- Better traffic localization algorithms in order to reduce the burden on the Internet core by reducing the amount of traffic moving across the Internet.
- Better inter-ISP real-time traffic monitoring tools that flag inefficient and excessive inter-AS traffic due to P2P video multicast.
- Use of adaptive-QoS techniques such as channel coding and intelligent packet (re-)scheduling to improve the fraction of correctly received stream blocks on peers.
- Using powerful seeding servers that supply initial bursts of stream blocks to connecting peers in order to reduce the channel startup time.
- Optimization of audio/video codecs for the high packet loss of P2P video multicast networks.
- Suitable buffer management and preemptive content caching algorithms that reduce channel startup times.

We believe that large scale P2P video multicast systems will have to address many of the issues listed above in order to become the video content distribution platforms of choice on the Internet.

Acknowledgment

The authors thank the P2P system provider for providing logs. They also thank Anja Feldmann and Vinay Aggarwal for useful discussions.

References

[1] Akamai, “Akamai Inc.” http://akamai.com.
[2] J. Hawkinson and T. Bates, “Guidelines for creation, selection, and registration of an autonomous system (AS),” http://rfc.net/rfc1930.html, 1996.
[3] *Advanced Video Coding for Generic Audiovisual services, ITU-T Recommendation H.264 - ISO/IEC 14496-10(AVC)*, ITU-T and ISO/IEC JTC 1, 2003.

[4] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should internet service providers fear peer-assisted content distribution?" in *IMC'05: Proceedings of the Internet Measurement Conference 2005 on Internet Measurement Conference*. Berkeley, CA, USA: USENIX Association, 2005, pp. 6–6.

[5] J. Glasner, "P2P fuels global bandwidth binge," Wired, April 2005.

[6] R. Ma, S. Lee, and D. Yau, "A game theoretic approach to provide incentive and service differentiation in P2P networks," in *ACM SIGMETRICS/Performance Evaluation Review*, 2004.

[7] Z. Ge, D. R. Figueiredosn, S. Jaiswal, J. Kurose, and D. Towsley, "Modeling peer-peer file sharing systems," in *IEEE Infocom*, April 2003.

[8] S. Saroiu, P. Gummadi, and S.Gribble, "A measurement study of peer-to-peer file sharing systems," in *Multimedia Computing and Networking*, 2002.

[9] D. Qiu and R. Srikant, "Modelling and performance analysis of bittorrent-like peer-to-peer networks," in *ACM Sigcomm*, August 2004.

[10] T. Ng, Y. Chu, S. Rao, K. Sripanidkulchai, and H. Zhang, "Measurement-based optimization techniques for bandwidth-demanding peer-to-peer systems," in *IEEE Infocom*, April 2003.

[11] Carnegie Mellon University, "End System Multicast," http://www.cs.cmu.edu/ esm.

[12] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai, "Distributing streaming media content using cooperative networking," in *ACM NOSSDAV*, 2002.

[13] Y. Chawathe, "Scattercast: an adaptable broadcast distribution framework," *Multimedia Syst.*, vol. 9, no. 1, 2003.

[14] J. Silber, S. Sahu, J. Singh, and Z. Liu, "Augmenting overlay trees for failure resiliency," in *IEEE Global Telecommunications Conference*, 2004.

[15] V. K. Goyal, "Multiple description coding: Compression meets the network," *IEEE Signal Processing Magazine*, vol. 18, no. 5, pp. 74–93, September 2001. [Online]. Available: citeseer.ist.psu.edu/goyal01multiple.html

[16] SopCast, "Sopcast." [Online]. Available: http://www.sopcast.org/

[17] PPLive, "PPLive." [Online]. Available: http://www.pplive.com/

[18] X. Zhang, J. Liuy, B. Liz, and T. Yum, "Coolstreaming/DONet: A Data-Driven Overlay network for efficient Live Media Streaming," *In Proc. of IEEE INFOCOM*, March 2005.

[19] Coolstreaming, "Coolstreaming." [Online]. Available: http://www.coolstreaming.us/

[20] Roxbeam Media Networks, "Roxbeam." [Online]. Available: http://www.roxbeam.com

[21] Gridmedia, "Gridmedia." [Online]. Available: http://www.gridmedia.com.cn/

[22] C. Abad, W. Yurcik, and R. H. Campbell, "A survey and comparison of end-system overlay multicast solutions suitable for network centric warfare," citeseer.ist.psu.edu/735587.html.

[23] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "Splitstream: high-bandwidth multicast in cooperative environments," in *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*, 2003.

[24] E. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A Survey and Comparison of Peer-to-Peer Overlay Network Schemes," *IEEE Communications Surveys and Tutorials*, vol. 7, no. 2, pp. 72–93, Oct. 2005.

[25] N. Magharei, R. Rejaie, and Y. Guo, "Mesh or Multiple-Tree: A Comparative Study of Live Peer-to-Peer Streaming Approaches," *In Proc. IEEE INFOCOM*, May. 2007.

[26] T. Silverston and O. Fourmaux, "P2P IPTV Measurement: A Comparison Study," 2006. [Online]. Available: http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0610133

[27] S. Ali, A. Mathur, and H. Zhang, "Measurement of commercial peer-to-peer live video streaming." [Online]. Available: citeseer.ist.psu.edu/ali06measurement.html

[28] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross, "A Measurement Study of a Large-Scale P2P IPTV System," *IEEE Transactions on Multimedia*, vol. 9, no. 8, December 2007.

[29] X. Hei, C. Liang, J. Liang, Y. Liu, and K. Ross, "A Measurement Study of a Large Scale P2P IPTV System," Department of Computer and Information Science, Polytechnic University, New York, USA, Tech. Rep., 2006.

[30] X. Zhang, J. Liu, and B.Li, "On Large Scale Peer-to-Peer Live Video Distribution: Coolstreaming and its Preliminary Experimental Results," *In Proc. of IEEE International Workshop on Multimedia Signal Processing, MMSP*, Oct. 2005.

[31] M. Zhang, L. Zhao, Y. Tang, J. Luo, and S. Yang, "Large-Scale Live Media Streaming over Peer-to-Peer Networks through the Global Internet," *In Proc. of ACM Multimedia 2005, Workshop on Advances in Peer-to-Peer Multimedia Streaming*, Nov. 2005.

[32] Y. Tang, J. Luo, Q. Zhang, M. Zhang, and S. Yang, "Deploying P2P networks for large-scale live video-streaming service," *IEEE Communications Magazine*, vol. 45, no. 6, pp. 100–106, June 2007. [Online]. Available: http://www.comsoc.org

[33] G. A. Fowler and S. McBride, "Newest Export from China: Pirated Pay TV," *Wall Street J.*, Sept. 2 2005.

[34] S. Agarwal, "A case study of large scale P2P video multicast," in *IMSAA'07: Proceedings of the IP Multimedia Subsystems Architecture and Applications Conference*, 2007.

[35] R. Droms, "Dynamic host configuration protocol," http://www.faqs.org/rfcs/rfc2131.html, 1997.

[36] Y. Rekhter and T. Li, "A border gateway protocol 4 (bgp-4)," http://www.faqs.org/rfcs/rfc1771.html, 1995.

[37] Cymru, "The Team Cymru IP to ASN Lookup Page," http://www.cymru.com/BGP/asnlookup.html.

[38] CAIDA, "Cooperative Association for Internet Data Analysis," http://www.caida.org/.

[39] K. C. Claffy, "Caida: Visualizing the internet," *IEEE Internet Computing*, vol. 5, no. 1, p. 88, 2001.

[40] X. Dimitropoulos, D. Krioukov, G. Riley, and K. C. Claffy, "Revealing the autonomous system taxonomy: The machine learning approach," in *In Proceedings of Passive and Active Measurements (PAM)*, 2006.

[41] G. Huston, "Interconnection, Peering, and Settlements," *in Proc. INET*, June 1999.

[42] Wikipedia, "Peering," http://en.wikipedia.org/wiki/Peering.