

# Diagnosing Network-wide P2P Live Streaming Inefficiencies

Chuan Wu

Department of Computer Science

The University of Hong Kong

Hong Kong

cwu@cs.hku.hk

and

Baochun Li

Department of Electrical and Computer Engineering

University of Toronto

Canada

bli@eecg.toronto.edu

and

Shuqiao Zhao

UUSee Inc.

China

zhaosq@uusee.com

---

Large-scale live peer-to-peer (P2P) streaming applications have been successfully deployed in today's Internet. While they can accommodate hundreds of thousands of users simultaneously with hundreds of channels of programming, there still commonly exist channels and times where and when the streaming quality is unsatisfactory. In this paper, based on more than two terabytes and one year worth of live traces from UUSee, a large-scale commercial P2P live streaming system, we show an in-depth *network-wide* diagnosis of streaming inefficiencies, commonly present in typical mesh-based P2P live streaming systems. As the first highlight of our work, we identify an evolutionary pattern of low streaming quality in the system, and the distribution of streaming inefficiencies across various streaming channels and in different geographical regions. We then carry out an extensive investigation to explore the causes to such streaming inefficiencies over different times and across different channels/regions at specific times, by investigating the impact of factors such as the number of peers, peer upload bandwidth, inter-peer bandwidth availability, server bandwidth consumption, and many more. The original discoveries we have brought forward include the two-sided effects of peer population on the streaming quality in a streaming channel, the significant impact of inter-peer bandwidth bottlenecks at peak times, and the inefficient utilization of server capacities across concurrent channels. Based on these insights, we identify problems within the existing P2P live streaming design and discuss a number of suggestions to improve real-world streaming protocols operating at a large scale.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed Applications*

General Terms: Measurement

Additional Key Words and Phrases: Peer-to-peer streaming, streaming inefficiency

## 1. INTRODUCTION

Large-scale live peer-to-peer (P2P) streaming systems have been successfully and commercially deployed in the Internet, delivering channels of live multimedia content to hundreds of thousands of users at any given time [PPLive ; UUSee ; PPStream ; TVAnts ]. Practical experience with these streaming systems, however, has shown that they may not be able to provide a satisfactory viewing experience to all the users at all times. Since the streaming quality experienced by the peers is of pivotal importance to the success of a P2P streaming system, it is critical to study existing commercial systems to look for any possible streaming *inefficiencies*, *i.e.*, low streaming qualities experienced by participating users (or *peers*). If such inefficiencies exist, it is also important to find reasons that may have caused them.

To this end, we have collected more than two terabytes and one year worth of live traces from a large-scale commercial P2P live streaming system, UUSee [UUSee ], one of the top three commercial P2P streaming systems in mainland China, along with PPLive and PPStream. This paper presents our in-depth investigation of the network-wide streaming inefficiencies in this large-scale real-world P2P streaming system, that may adversely affect user experiences. Our first objective is to look for patterns that may show evidence that such streaming inefficiencies exist. If so, we wish to identify influential factors that lead to such inefficiencies, in order to gain useful insights towards further improvements of P2P live streaming protocols.

As the first highlight of our study, we have identified an evolutionary pattern of low streaming qualities in the system, and the distribution patterns of streaming inefficiencies across various streaming channels and in different geographical regions. We then conduct an extensive study of the causes to such inefficiencies, by investigating the impact of factors such as the number of peers, peer upload bandwidth, inter-peer bandwidth availability, server bandwidth consumption, the level of reciprocity among peers, and many more. In particular, we explore the correlation between the evolutionary pattern of streaming inefficiencies and the variation of various factors over time, and also zoom into snapshots of the entire system to investigate causes to the low streaming qualities experienced in specific channels and regions at each specific time.

A number of original discoveries have been brought forward in our study. *First*, the population of peers has two-sided effects on the streaming quality in a streaming channel: while peers in a larger channel typically enjoy a better streaming quality, the streaming quality degrades as the peer population increases in the channel during peak hours of a day. *Second*, the dedicated server capacity still plays an indispensable role in P2P streaming over today's Internet, and the inefficient supply of server capacity has largely led to the low streaming quality at specific times and regions. Third, smaller channels tend to obtain less server capacity on a per-peer basis than large channels, which contributes significantly to their low streaming qualities. *Fourth*, inter-peer bandwidth availability represents a more significant bottleneck than the upload bandwidth at peers during daily peak hours. *Finally*, increasing peer indegree in the topology does not help enhancing the peer streaming quality. Based on these insights, we discuss a number of suggestions on improving the design of real-world P2P streaming systems operating at a large scale.

The remainder of this paper is organized as follows. In Sec. 2, we outline trace

collection methodologies in the UUSee streaming platform, and present the measurements that constitute the basis of our study. In Sec. 3, we investigate distribution of streaming inefficiency across different regions and channels, as well as their evolutionary patterns over a long period of time. We then search for the causes to the temporal pattern of streaming inefficiencies in Sec. 4, and diagnose the distribution of streaming inefficiencies across various streaming channels and in different regions in Sec. 5. We discuss related work in Sec. 6. Finally, we summarize our discoveries and discuss suggestions on improving current P2P streaming protocols in Sec. 7.

## 2. BACKGROUND

We first present an overview of the P2P streaming solution used in UUSee, our methodology for collecting traces, and the scale of the P2P streaming application.

### 2.1 UUSee P2P streaming solution

Starting September 2006, we have continuously monitored the performance statistics of a real-world commercial P2P streaming platform, offered by UUSee Inc., a leading P2P streaming solution provider with legal contractual rights with mainstream content providers in China. Similar to all state-of-the-art P2P streaming applications, UUSee P2P streaming employs a hybrid framework consisting of dedicated streaming servers and mesh-based P2P streaming overlays: a collection of about 100 dedicated streaming servers are deployed, which simultaneously broadcast over 800 live streaming channels (mostly encoded to streams around 500 Kbps) to hundreds of thousands of peers, and provide a certain level of stable upload bandwidth for the streaming channels.

As a typical mesh-based P2P streaming system, UUSee organizes peers in each streaming channel into a mesh overlay topology and utilizes “pull-based” design on the mesh P2P topologies, that allows peers in each channel to serve other peers (“partners”) by exchanging media blocks in their playback buffers. The playback buffer at each peer represents a sliding window of the media channel, with a buffer size of 500 media blocks, where each block represents 1/3 second of media playback.

When a new peer joins a channel in UUSee, the initial set of a number of partners (up to 50) is supplied by one of the tracking servers, by randomly selecting from all the existing peers in the channel. The peer establishes TCP connections with these partners, and buffer availability bitmaps (*i.e.*, buffer maps) are periodically exchanged. Based on their partners’ block availability, each peer schedules the retrieval of new blocks from each partner, prioritizing the block retrieval sequence based on the playback deadline of the blocks.

To maximally utilize peer upload bandwidth and alleviate server load, UUSee has incorporated a number of algorithms in the peer selection process, beyond the initial random partner assignment from the tracking servers. Each peer applies an algorithm to estimate its maximum upload capacity, and continuously estimates its aggregate instantaneous sending throughput to its partners. If its estimated sending throughput is lower than its upload capacity for 30 seconds, it will inform one of the tracking servers that it is able to receive new connections. The tracking servers keep a list of such peers, and assign them upon requests of partners from other peers. In addition, the number of consecutive blocks received and cached

in the playback buffer, starting from the current playback time, is used in UUSee protocols to represent the current streaming quality of a peer, which is referred to as the *buffering level*. During the streaming process, neighboring peers also recommend known partners to each other based on their current streaming quality. When a peer  $i$  finds that another peer  $j$  has a low buffering level, *i.e.*, an insufficient number of consecutive blocks in its buffer, peer  $i$  will recommend its known partners with larger buffering levels. As a last resort, when a peer has a low buffering level for a sustained period of time, it will contact the tracking server again to obtain additional partners with better streaming qualities.

Each dedicated streaming server in UUSee utilizes a similar P2P protocol as deployed on regular peers, and is randomly selected to serve the peers by the tracking servers when it still has available upload capacity.

## 2.2 Trace collection methodologies

To inspect the run-time behavior of UUSee P2P streaming, we have implemented detailed measurement and reporting capabilities within its P2P client software, based on collaborative efforts with the UUSee development team. Each peer collects a set of its vital statistics, and reports them to dedicated trace servers every 5 minutes via UDP. The statistics include its IP address, the channel it is watching, its current buffer map, the number of consecutive blocks in its current playback buffer (*i.e.*, the buffering level), its instantaneous aggregate sending and receiving throughput to and from all its partners, as well as a list of all its partners, with their corresponding IP addresses, TCP/UDP ports, and current sending/receiving throughput to/from each partner. As each dedicated streaming server in UUSee utilizes a similar P2P protocol as deployed on regular peers, it reports its related statistics periodically as well. Details on the measurement methodologies of the above metrics can be found in our previous work [Wu et al. 2007].

## 2.3 Trace summary

During September 2006 to September 2007, we have collected more than two terabytes of traces with more than 700 million unique IP addresses and 3000 million P2P flows, representing time-continuous snapshots of the live P2P streaming system every five minutes. At any time, there are more than 100,000 concurrent peers and over 1 million simultaneous P2P flows in the entire system. As the basis of our study, we summarize the distribution of peers and streaming servers in different regions and channels in UUSee.

**2.3.1 Geographical distribution of peers and servers.** Fig. 1 shows the span of UUSee network with respect to ISPs/ASs in China and continents/countries in other parts of the world. The AS number and country each IP address in the traces belongs to are derived using the Whois service provided by Cymru [Cymru Whois service]; to map each AS number in China to its affiliated ISP, we make use of the official mapping data provided by CERNET, China [Cernet BGP View]. The number of peers in each region is derived as the average concurrent UUSee peer population in September, 2007. As compared to the China ISP statistics provided by CNNIC [CNNIC], we find that the UUSee streaming network spans all the major ISPs (*i.e.*, Telecom, CNC, etc.) and most of the small regional ISPs (XAOnline,

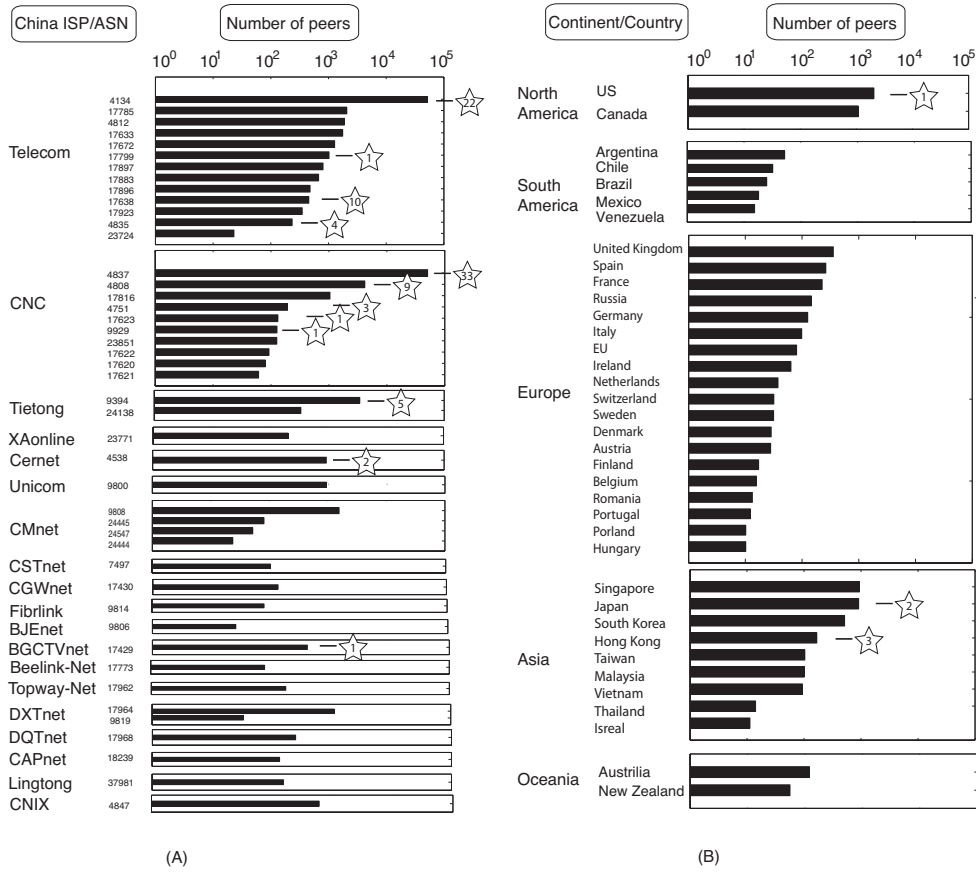


Fig. 1. Geographic distribution of UUSee peers and servers: (A) in China ISPs/ASNs; (B) in other continents/countries.

CAPNet, BJEnet, etc.) in China, and also has a substantial number of users in most other parts of the world as well.

Fig. 1 also shows the distribution of streaming servers in UUSee network, where the number of streaming serves deployed in a region is marked inside a star shape beside the region. We observe that the servers are deployed across different regions in a very ad-hoc fashion, which only slightly corresponds to the peer population in the regions.

2.3.2 *Channel-wise distribution of peers and servers.* Fig. 2 plots the distribution of concurrent peer population in different streaming channels in UUSee, as computed as the average simultaneous number in each channel in September, 2007. Across the streaming channels, the popularity differs significantly: There are a small number of most popular channels ( $\approx 2\%$ ) with an average peer population of more than 5000, a small percentage of channels ( $\approx 12\%$ ) have fewer than 100 peers, and the majority of UUSee channels accommodate a peer population in the range of 500 to 3000 ( $\approx 51\%$ ).




Fig. 2. Distribution of UUSee peers in different streaming channels.

For channel deployment on streaming servers, each server typically hosts around 40 channels, while each channel is deployed on 5-6 servers.

### 3. DISTRIBUTION OF STREAMING INEFFICIENCY

We now start our characterization of streaming inefficiencies commonly existing in the P2P streaming application, by investigating the low streaming quality scenarios across different channels, in various regions, and over different times.

We first define how we evaluate the streaming quality of a channel/region. The *streaming quality of a channel/region* at each time is defined as the percentage of high-quality peers in the channel/region, where a high-quality peer is defined as a peer with a buffering level (i.e., the number of consecutive blocks received and cached in the current playback buffer of the peer) of more than 80% of the total size of its playback buffer. The criterion of buffering level has been extensively used in UUSee streaming protocols to evaluate the current streaming quality of a peer; and the 80% buffering level benchmark has empirically been shown to be effective in reflecting the playback continuity of a peer in the following few minutes, based on observations made on an internal performance monitoring system deployed in UUSee. Accordingly, we also use peer buffering level as our basic streaming quality criterion, and the streaming quality of a channel/region in all the figures hereinafter refers to the definition above.

#### 3.1 Distribution across different channels

We first investigate the distribution of streaming inefficiencies in UUSee, by plotting the distribution of streaming quality across all the channels at four representative snapshots of the system on September 21, 2007.<sup>1</sup> In Fig. 3, each point  $(x, y)$  shows the streaming quality  $y$  of the channel with index  $x$ , and each of the curves is plotted in the descending order of the channel streaming qualities. Therefore, the indices on the x-axis may not correspond to the same channels along different curves. In addition, the total number of channels with non-zero viewers is different at different times (500, 595, 704, and 758, respectively), which explains the different ranges of

<sup>1</sup>All the times in the paper are based on Beijing time, GMT+8. We note that the majority of UUSee users reside in this time zone, and the possibly different viewing behavior of users in other time zones has little impact on our observations to appear hereinafter in the paper.

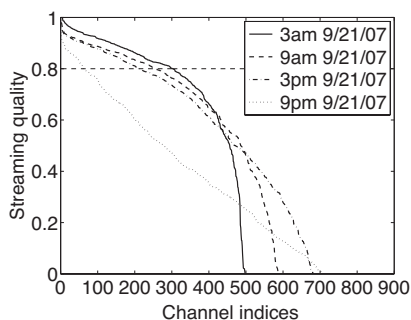


Fig. 3. Distribution of streaming quality across all the channels in UUSee.

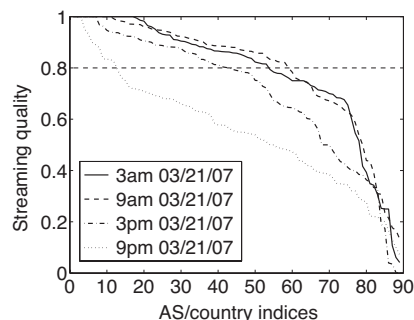


Fig. 4. Distribution of streaming quality among all the regions in UUSee.

the channel indices for different curves. We can observe that the streaming quality varies significantly among channels at each time. We also notice that 50 – 60% of all the channels have a streaming quality higher than 0.8 (*i.e.*, 80% of the peers in the channel enjoy smooth playback) at earlier times of a day, the percentage decreases towards later times of the day, and it is significantly low in the evenings (667 out of 754 channels at 9pm have a streaming quality lower than 0.8). These observations exhibit the existence of streaming inefficiency (low streaming quality) in channels at each time, which becomes more evident at later times of a day.

### 3.2 Distribution among different regions

We next explore any difference of streaming quality experienced by peers located in different geographic regions. To investigate such spatial distribution of streaming inefficiency, we group peers inside China based on their ASs, and those outside China according to their countries. The streaming quality of each region (*i.e.*, referring to an AS in China, and a country outside China) at each time is evaluated as the percentage of high-quality peers in the region.

Fig. 4 plots the distribution of streaming quality among all the regions at four representative times on March 21, 2007. Similar to Fig. 3, each curve is plotted in the descending order of the streaming quality of the regions, and indices on the x-axis may not correspond to the same regions for different curves. Largely different streaming qualities are observed among different regions in UUSee, and a temporal variation of the distribution, similar to that in Fig. 3, can be observed: 60 – 70% of all the regions at early hours of the day have a streaming quality higher than 0.8, 50% in the afternoons, and only 17% in the evening.

In addition, Table I presents a closer investigation of the regional distribution of streaming inefficiency, by listing the number of ASs with streaming inefficiency (*i.e.*, a streaming quality lower than 0.8) in each ISP in China (with the percentage of such ASs inside the ISP in the parentheses), and that of countries in each continent, at a morning time and an evening time on March 21, 2007. We observe streaming inefficiency in the morning time mainly occurs in countries outside China, and in the evening, streaming inefficiency generally exists within all the ISPs in China and all the overseas countries, with the percentage of low streaming quality ASs slightly lower in the major ISPs — Telecom and CNC — and almost all 100% in the other

Table I. Distribution of regions with streaming inefficiency across ISPs/continents on Mar. 21, 2007

ISP	No. of regions with streaming inefficiency at 9 am (Percentage)	No. of regions with streaming inefficiency at 9 pm (Percentage)
Telecom	1(7%)	12(92%)
CNC	2(20%)	8(80%)
Tietong	0(0%)	2(100%)
XAOnline	0(0%)	1(100%)
Cernet	1(100%)	1(100%)
Unicom	0(0%)	1(100%)
CMNet	0(0%)	4(100%)
CSTNet	0(0%)	1(100%)
CGWNet	0(0%)	1(100%)
BGCTVNet	0(0%)	1(100%)
Beelink-Net	0(0%)	1(100%)
Topway-Net	0(0%)	1(100%)
DXTNet	0(0%)	2(100%)
DQTNet	0(0%)	1(100%)
CAPNet	0(0%)	1(100%)
Lingtong	0(0%)	1(100%)
CNIX	0(0%)	1(100%)
FiberLink	0(0%)	1(100%)
BJENet	0(0%)	1(100%)
North America	0(0%)	2(100%)
South America	4(80%)	5(100%)
Europe	13(68%)	15(79%)
Asia	3(33%)	8(89%)
Oceania	2(100%)	2(100%)

ISPs and overseas countries.

These observations show the existence of streaming inefficiency in regions at all times in a day, most of which are distributed in smaller ISPs in China and overseas countries. In addition, more regions experience streaming inefficiency on evenings than on mornings.

### 3.3 Evolutionary pattern over the trace period

The above observations on the distribution of streaming inefficiencies across channels and regions have both revealed a time of the day effect on the streaming quality. To inspect any evolutionary pattern of streaming quality, we plot in Fig. 5 and Fig. 6 the evolution of the average streaming quality over all the channels or regions, respectively, over the 12 months of trace span (Note that April 2007 is skipped due to lack of traces in the month caused by the trace server upgrade). The average streaming quality across all the channels/ASs/countries in these figures is computed as the weighted average of streaming qualities of individual channels/ASs/countries, weighted by the population in each channel/AS/country, respectively.

To reveal any daily evolutionary pattern over the entire trace span, we have chosen to plot in the figures the evolution of streaming quality on a representative date in each month. From all the plots, we can clearly observe a daily evolutionary






Fig. 5. Evolution of the average streaming quality of all the channels.




Fig. 6. Evolution of the average streaming quality among (A) all ASes in China; (B) all countries outside China.

pattern: no matter in terms of channels or regions, the average streaming quality is generally better at early hours of a day, degrades to a low value around the noon time, then improves slightly, and drops to its daily lowest value before midnight. The daily lowest value of streaming quality can be as low as 0.1 – 0.3, meaning that only 10% – 30% users in UUSee have a high buffering level on evenings, while the majority may experience playback jitters.

The interesting patterns of streaming inefficiencies pose us an intriguing question: what has caused such low streaming qualities at peers at evening times and in specific channels/regions? In what follows, we conduct an extensive investigation of the influential factors, that may have effected the network-wide temporal pattern of streaming qualities and the distribution of inefficiencies across channels/regions.

#### 4. DIAGNOSING STREAMING INEFFICIENCY: THE DAILY EVOLUTION

We first investigate the cause to the daily evolutionary pattern of streaming qualities in UUSee, by exploring the correlations between the evolution of streaming quality and the variation of various influential factors. The factors we are to investigate include:

*Number of peers in a channel.* It directly decides the level of bandwidth demand and supply in a channel, affecting the achievable channel streaming quality.

*Average upload bandwidth per peer.* The bandwidth contribution at the peers constitutes the main source of upload bandwidth in a P2P streaming channel.

*Average server bandwidth consumption per peer.* The server capacity represents an indispensable source of upload bandwidth in cases of insufficient peer bandwidth contribution and volatile peer dynamics.

*Intra-ISP/Inter-ISP per-link bandwidth availability.* We investigate the impact of average bandwidth availability along the P2P links within a same ISP (*i.e.*, the *intra-ISP* case where both peers are in the same ISP) and across ISP boundaries (*i.e.*, the *inter-ISP* case where the two peers are in different ISPs), which may reflect possible inter-peer bandwidth bottlenecks.

*Average peer indegree.* It is related to the level of download bandwidth at the peers.

*Average buffer map difference at a peer.* It represents the level of content availability along each incoming P2P link at a peer, and is computed in the following fashion: The total number of non-duplicated new blocks a peer's partners can supply to it (which also fall in the peer's window of interest) is counted by comparing the buffer maps, and then divided by the number of partners. In a mesh streaming system featuring block exchanges, a larger buffer map difference represents less chance of content bottleneck at the peers during streaming, and a higher probability to saturate the available bandwidth on each P2P link.

#### 4.1 Evolution of the influential factors

With the example of a popular channel CCTV1 and a less popular channel CCTV12, Fig. 7 plots the evolution of the streaming quality and the above factors in the channels. A daily variation pattern can be observed in all the evolution series. A closer look reveals that the drop of streaming quality on evenings in each streaming channel is generally accompanied by the increase of peer population, and the decrease of available server capacity per peer and per-link bandwidth availability, with respect to both the intra-ISP and inter-ISP link bandwidths. On the other hand, a bit surprisingly, we notice that the peer upload bandwidth contribution, the buffer map difference, and the peer indegree may have increased at the evening peak hours.

In order to verify our observations on the possible correlations between the evolution series of each factor and that of the streaming quality in a streaming channel, we compute the *cross-correlation* between the series over the trace period. The cross-correlation between two time series,  $x_t, t = 1, 2, \dots$  and  $y_t, t = 1, 2, \dots$ , at a delay of  $d$  is evaluated as the correlation between the corresponding members of  $x_t$  and  $y_{t+d}$ ,  $t = 1, 2, \dots$ , in the two series. For example, we may compute the cross-correlation between the time series of the streaming quality and the number of peers for CCTV1 and CCTV12 at the delay of 0. For CCTV1, a significant negative cross-correlation,  $-0.71$ , validates a strong negative correlation between the peer population and the achieved streaming quality in this channel over time. On the other hand, a weaker negative correlation in the case of CCTV12 is also established, based on the less significant negative cross-correlation of  $-0.11$  at the delay of 0.

#### 4.2 Cross-correlation with the influential factors

To investigate whether our observations made for channel CCTV1 and CCTV12 generally exist among all the channels, we have extensively investigated the cross-correlations between evolution series of streaming quality and each of the influential

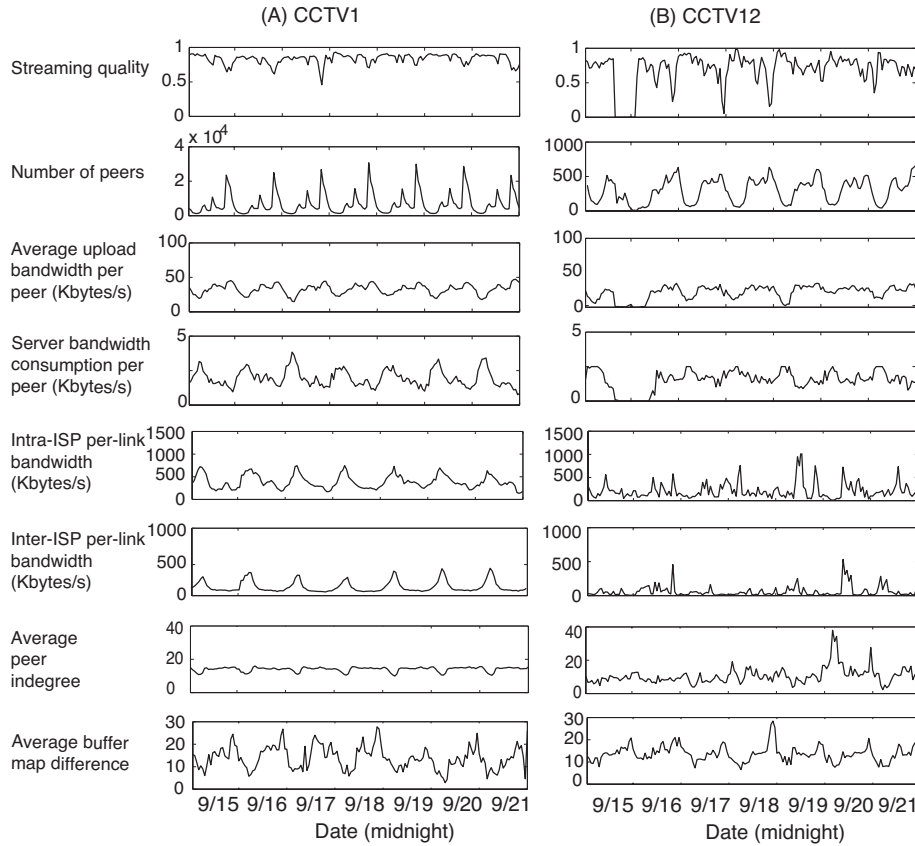


Fig. 7. Evolution of streaming quality and influential factors in September 15 — 21, 2007.

factors for all the channels in UUSee. As we have observed difference levels of cross-correlation for channels of different sizes, *i.e.*, different peer population, we show our results about the cross-correlation in each streaming channel by plotting it against the size of the channel, as in Fig. 8, in order to discover any possible relationship between the two. In each sub figure in Fig. 8, the y-axis represents the cross-correlation at the delay of 0 between the two series of the streaming quality and one respective influential factor (we have used “cross-correlation with the respective factor” for short in the y labels), and each sample in the plot denotes one streaming channel. The size of each channel is derived as the average number of peers in the channel over the trace period.

Fig. 8(A) exhibits that for most channels in UUSee (as denoted by the dense areas), the cross-correlation between the evolution of streaming quality and that of the peer population is negative, meaning that the streaming quality is worse with more peers in the channels. Such a negative cross-correlation is more significant for popular channels with larger peer population; only in channels with fewer than a few hundreds of peers, can the cross-correlation be positive. Such a negative impact of peer population in most channels is *controversial* to the common belief that the P2P system may scale better with more peers in it. But why is it so?

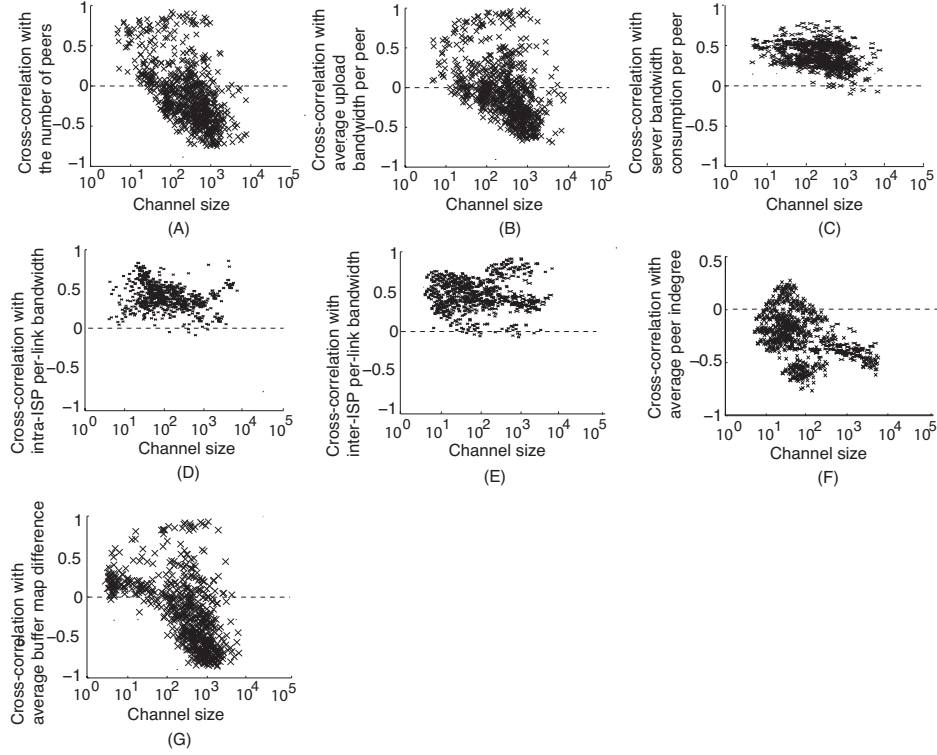


Fig. 8. Cross-correlation between the series of streaming quality and each influential factor for all the streaming channels in UUSee.

Let’s look at the results in Fig. 8(B), which reveal a negative cross-correlation between the streaming quality and average peer bandwidth contribution for most channels. This observation has made the question even more interesting: why shall the streaming quality in most sizable channels still degrade in the daily peak hours with more peers and larger bandwidth contribution from each peer?

The plots in Fig. 8(C), (D) and (E) bring us useful insights towards the answer. A positive cross-correlation is exhibited in these figures for almost all the channels, revealing that the evolution of streaming quality in each channel is largely caused by the evolution of the three respective factors. In another word, at daily peak hours, the less available server capacity to each peer and the lower level of available bandwidth along the P2P links, have produced a negative impact on the streaming quality. However, why would the available server capacity become less sufficient with more peer bandwidth contribution at the peak hours? We identify the reason to be as follows: even when peers contribute more at the peak hours (as shown in Fig. 7), the level of peer bandwidth contribution is still lower than the required streaming bandwidth (*i.e.*, 500Kbps),<sup>2</sup> and the insufficient amount has to be covered by the upload capacity from the streaming servers. As the total amount of

<sup>2</sup>The majority of UUSee users are in China, with a typical upload capacity below 512 Kbps. However, such an upload capacity still represents a theoretical upper bound, which users may

server capacity is limited and the increase of peer population at peak hours is more significant than the increase of peer bandwidth contribution, the average server capacity available to each peer becomes less with more peers in the channel.

Other than the bandwidth availabilities, we further investigate the impact of peer indegree and buffer map difference. Fig. 8(F) plots a negative cross-correlation between streaming quality and the average peer indegree in most channels. This brings another interesting discovery: at daily peak hours, peers tend to obtain more partners (possibly as a result of trying to boost their degraded streaming quality), but nevertheless the streaming quality is still poorer, due to the low level of bandwidth availability along each link (as shown by the intra- and inter-ISP per-link bandwidth plots in Fig. 7).

From Fig. 8(G), we observe a negative cross-correlation for most of the large channels. It shows that at peak hours in large channels, there is no content bottleneck present among the peers, as the number of available blocks to exchange among peers is actually larger than that at the other times. Therefore, the content availability does not constitute a significant cause to the streaming quality downgrade in large channels. On the other hand, in small channels, the positive cross-correlations reveal the existence of content bottlenecks among peers, when the low streaming quality occurs.

In summary, the above observations bring us intriguing insights with respect to the causes to the streaming inefficiencies commonly existing on a daily basis:

- ▷ In large channels with hundreds or thousands of peers, the lack of available server capacity per peer and the downgrade of inter-peer bandwidths contribute significantly to the downgrade of streaming quality with more peers in the channel at daily peak hours, while peers have contributed better than usual and there is no content bottleneck among them.
- ▷ In small channels with fewer than a few hundreds of peers, streaming inefficiency usually occurs at times with smaller peer population, due to a general lack of server, link and peer bandwidths and the available content blocks to exchange (which may essentially be a result of the lack of bandwidth as well).
- ▷ The increase of partner numbers in both large and small channels may not help boosting the peer streaming quality, as the available bandwidth along the links is the key.
- ▷ In a practical P2P streaming system, server capacity is still indispensable to guarantee the streaming quality of mid-quality streams around 500Kbps, given the realistic level of possible peer bandwidth contributions over today's Internet.

Based on these insights, we derive a number of suggestions for the enhancement of P2P streaming protocols:

- 
- (1) Given the importance of server capacity and its cost, it is desirable to dynamically provision server capacity in the system, catering for the different demand over different times, *e.g.*, to temporarily enable more servers at peak hours.

---

not be able to achieve in practical scenarios. The average peer upload bandwidth contribution in UUSee network over time is around 350 Kbps, as derived from our trace measurements.

(2) As streaming inefficiency may occur at different times among large channels and small channels, we may make use of such demand difference in different channels at each time to smartly allocate server capacity across channels over time.

(3) While both intra-ISP and inter-ISP link bandwidths degrade at daily peak hours, the intra-ISP bandwidths are always larger than inter-ISP bandwidths (as shown in Fig. 7). Therefore, it is more desirable for peers at the peak hours to get more intra-ISP partners, which should not be difficult with many peers in the system at peak times.

(4) To find a few good partners with good bandwidths is better than having a lot of partners. The current random peer selection protocol in UUSee-like streaming systems should include more judgement of inter-peer bandwidth availability and replace bad partners (with low bandwidth) with good ones over time.

## 5. DIAGNOSING STREAMING INEFFICIENCY: FOCUSING ON ONE SNAPSHOT

Other than the temporal pattern of streaming inefficiencies over time, if we zoom into the snapshot of UUSee network at any specific time, we have observed that there commonly exist channels and regions which regularly experience a lower streaming quality than the others, as shown in Fig. 3 and Fig. 4. We now inspect the causes to such distribution of streaming inefficiency, by focusing on one representative snapshot of UUSee network at 9pm September 21 2007.<sup>3</sup> We first investigate the influential factors leading to the different streaming qualities achieved by different channels, and then explore the distribution of streaming quality in different regions.

### 5.1 Channels with low streaming quality in one snapshot

The influential factors we investigate for their impact on the distribution of streaming quality across channels include: *number of peers* (channel size), *average peer upload bandwidth*, *average server bandwidth consumption per peer*, and *average buffer map difference at a peer* in each channel. The average peer indegree and per-link bandwidth availability are not included as their values do not differ significantly among channels at one same time, since the same peer selection and streaming protocols are employed in each channel.

Fig. 9 plots the correlation between channel streaming quality and each of the factors. Each sample in the plots represents one streaming channel in UUSee. A positive correlation is observed in each of the plots with the computed Pearson product-moment correlation coefficients,  $\rho$ , marked at the upper right corner of each plot. It exhibits that the streaming quality is generally better in channels with a larger peer population, more peer bandwidth contribution on average, higher server capacity availability per peer, and larger buffer map differences. To dig deeper on the rationale behind the positive effect of peer population on the streaming quality across channels, we explore the relationship between channel peer population and the other factors in Fig. 10.

It is interesting to discover from Fig. 10(A) that the average peer bandwidth

<sup>3</sup>We choose such a representative snapshot in the evening, in order to focus on influential factors to low streaming qualities commonly existing at the peak times.

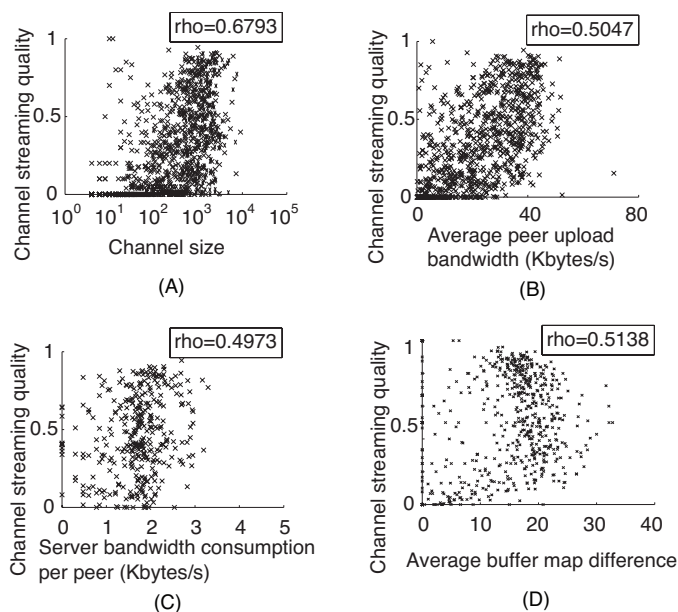


Fig. 9. Correlation of channel streaming quality with influential factors: for all the channels in UUSee at 9pm September 21 2007.

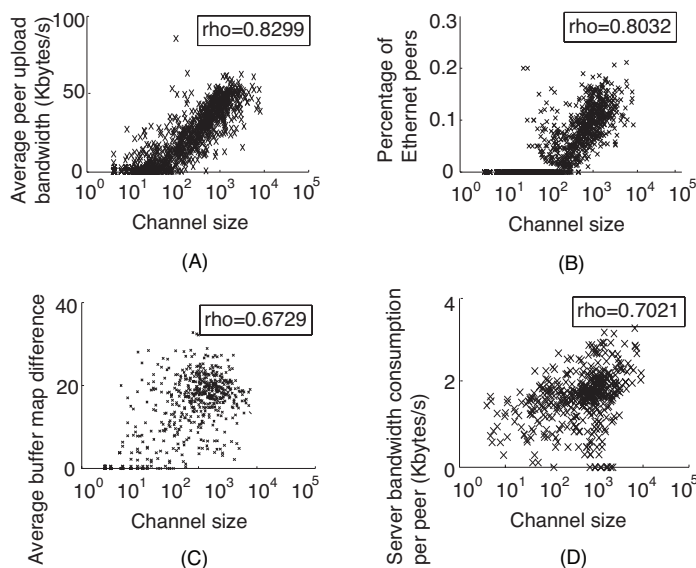


Fig. 10. Correlation of channel size with the other influential factors: for all the channels in UUSee at 9pm September 21 2007.

contribution is larger in a larger channel. We further identify the reason to be (1) the higher percentage of Ethernet peers in larger channels, as revealed by Fig. 10(B), and (2) the relative larger buffer map difference, as shown in Fig. 10(C), such that

the available peer upload bandwidth is more efficiently utilized. The phenomenon of larger buffer map differences in larger channels can be explained by the more diversified block distribution, when the peers have more diversified sets of partners in a larger mesh overlay.

In addition, Fig. 10(D) exhibits that the average server capacity used by each peer is larger in the larger channels as well. In the current UUSee protocol, there is no explicit allocation of server capacity to the different channels, and the streaming servers can be selected to serve peers from any channels in a similar random way as any regular peers are selected. This poses us an interesting question: Why is the per-peer server bandwidth share larger in large channels than that in small channels, since they may have been similar given such a random way of server capacity usage? Our further exploration brings the following discovery: Though the servers are randomly selected to serve peers in the entire system, the total number of upload connections each server can support is limited (to a constant number). Peers in a large channel usually obtain a better streaming quality and may typically stay online longer. They therefore have a higher chance to obtain a connection to a server upon their partner requests during streaming (a peer requests partners from a tracking server when it first joins and when its total partner number falls below a threshold). Therefore, the available connections to servers are mostly occupied by peers in large channels. When the peers from small channels request partners, they may not be able to obtain such a connection to a server.

Our investigations above show that popular large channels may grasp not only more server bandwidth in total, but also more server capacity per peer as well, which is not a fair situation for the less popular channels.

We have investigated snapshots at other specific times and made similar observations. In summary, we can derive the following insights towards the distribution of streaming quality across channels at each specific time:

- ▷ At each specific time across different channels, less popular channels with a smaller number of peers tend to experience lower streaming qualities, due to a lower percentage of high-bandwidth peers, less efficient utilization of peer upload bandwidth due to largely overlapping buffer maps among a small population, and the lower level of obtained server capacity in the competition with large channels, based on the current random way of provisioning server capacity in UUSee.

Our suggestions for the improvement of the streaming quality in less popular channels are:

---

As smaller channels are shown to be inferior in the competition for server capacity with large ones, an explicit allocation of server capacity across channels is desirable. In combination with suggestion (2) made at the end of Sec. 4, such an allocation should consider the different bandwidth demand and supply in different channels at each time, and aim to maintain a good streaming quality for both large and small channels.

---

## 5.2 Regions with low streaming quality in one snapshot

Inside each channel, the streaming quality of peers may differ from one geographic region to another. We now investigate the effects of the following influential factors



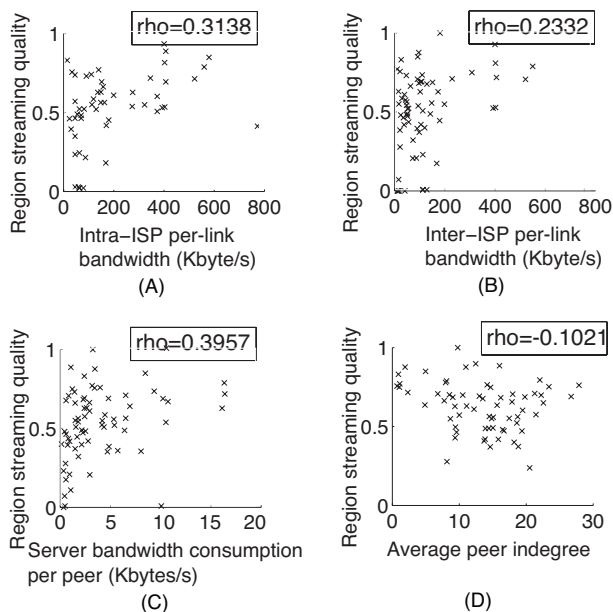


Fig. 11. Correlation of region streaming quality with influential factors: for all the regions in UUSee at 9pm September 21 2007.

on such regional distribution: *Intra-ISP/Inter-ISP per-link bandwidth availability*, *average server bandwidth consumption per peer*, and *average peer indegree* at the peers in each region. In our study, we have investigated the regional distribution of streaming qualities inside individual channels and over the entire UUSee network, and have made similar observations on the influential factors. Therefore, in Fig. 11 where we plot the correlation between regional streaming quality and each of the influential factors, the streaming quality and factors of a region are computed using measurements from all the peers in the region, spanning multiple channels. Each sample in the plots represents one region (*i.e.*, referring to an AS in China, and a country outside China).

We can observe a positive correlation between the streaming quality and the respective factor in Fig. 11(A), (B) and (C), respectively, while no evident correlation is shown in Fig. 11(D). This again shows that server and link bandwidths are important factors to decide the streaming qualities across different regions, while peer indegree is not. In our study, we have carefully investigated the regions with low streaming quality, as shown in Fig. 11, and made the following discoveries: (1) Overseas countries constitute a large portion of those regions. To explore the reason, we have observed that peers overseas are largely streaming all the way from peers or servers in China, with low-bandwidth inter-continental P2P links. (2) The regions inside China with low streaming quality are mostly within small regional ISPs. Since peers in these regions do not have much choice of partners within the same ISP, most of their incoming links are from peers or servers in other ISPs, thus suffering from the inter-ISP link bandwidth bottlenecks.

In summary, we derive the following insights for regions with low streaming

quality at each specific time:

- ▷ Overseas countries and regions inside small China ISPs represent lower streaming qualities in UUSee, due mainly to the lack of access to servers and peers in adjacent regions.

Possible measures to enhance the streaming quality in those regions include:

---

(1) A fundamental way is to deploy a minimum but sufficient amount of server capacity overseas and in China ISPs with apparent peering bandwidth bottlenecks with other ISPs.

(2) In addition to the locality suggestion (4) made at the end of Sec. 4, locality of partners should be explicitly explored at all times in peer selection at peers in these regions, in order to minimize the significant impact of link bandwidth bottlenecks on these peers.

---

## 6. RELATED WORK

There have recently emerged a number of measurement studies targeting commercial P2P live streaming applications. In most studies, traces were collected using a *crawler* or the *passive sniffing* technique from a few dedicated monitoring computers. In comparison, our trace collection methodology can obtain snapshots of the system not only at a much larger scale (from essentially all the peers in the system), but also with less distortion over the time domain.

Using both crawling and passive sniffing, Hei *et al.* [Hei et al. 2006; 2007] have studied PPLive, with respect to peer distribution, fractions of control and video data at each peer, peer partnership, etc. In a recent work [Hei et al. 2007], they have further exploited buffer maps harvested from PPLive peers to monitor the streaming quality, and shown the correlation between buffer map information and viewing continuity at the peers. This result validates our way of evaluating peer streaming quality using the number of consecutive blocks in the buffer.

Also using traces collected with a crawler, Vu *et al.* [Vu et al. 2007] investigated two graphical properties of PPLive overlays, namely node outdegree and graph clustering coefficient. They concluded that small overlays in PPLive are similar to random graphs in structure. Using passive sniffing with one dedicated computer, Gao *et al.* [Gao et al. 2008] studied the packet size and self-similarity of the traffic in PPLive. We have extensively explored the causes to the low streaming qualities, which has not been the focus of the existing studies.

Jia *et al.* [Jia et al. 2007] investigated PPStream, using a dedicated PPStream crawler they have developed. They revealed some characteristics on the geographic clustering, connection stability, arrival/departure pattern, and the ratio of upload over download in the system.

Ali *et al.* [Ali et al. 2006] have studied PPLive and SopCast with traces collected in 2-3 hour durations using passive sniffing. They analyzed the resource usage in the systems, *i.e.*, the total sending/receiving throughput at each peer, number of children supported by each peer, and stability of the data distribution. Silverston *et al.* [Silverston and Fourmaux 2006; 2007] compared the traffic patterns among four P2P streaming applications, PPLive, PPStream, SopCast, and TVAnts, using

passive sniffing during the broadcast of the 2006 FIFA world cup. They compared their total upload/download traffic, the types of video and control traffic involved (UDP or TCP), the number of serving peers, and the peer life time distribution among the applications. In a recent paper, Lu *et al.* [Lu et al. 2009] also studied the packet size and traffic pattern in SopCast, using a small-scale emulation network that runs SopCast clients.

One of the few studies that use a similar trace collection methodology as ours is by Li *et al.* [Li et al. 2008; Li et al. 2007]. In their study of CoolStreaming, an ActiveX component is deployed at each peer and reports periodically to an internal logging server. They collected traces only for a few days, instead of over a long period of time as we did. They investigated the type and distribution of users, the distribution of peer uploading contributions, and the playback continuity in the system. They pointed out that the excessive start-up time and high failure rates during flash crowds remain critical problems in CoolStreaming, but have not extensively explored the causes.

Using data logged by P2P clients onto a centralized logging server, Agarwal *et al.* [Agarwal et al. 2008] have evaluated the performance of a large-scale P2P live video multicast session (without specifying the name of the streaming application). They investigated the peer connection types and network distribution, streaming quality, and peer behavior. They have mentioned that peer streaming quality differs across different ASs, but have not explored the detailed causes.

In a recent study by Alessandria *et al.* [Alessandria et al. 2009], four P2P live streaming applications, PPLive, SopCast, TVants and TVUPlayer, were compared using passively sniffed traces collected in a controlled environment. They focused on investigating the behavior of these applications under adverse network conditions, such as dynamically changing delay, loss rate and available bandwidth. They concluded that all applications employ adaptive mechanisms to cope with packet loss and congestion in the network core.

To the best of our knowledge, our work represents the first extensive and in-depth measurement study that explores causes to network-wide low streaming qualities in a large-scale P2P streaming system. A preliminary report of this work appeared in INFOCOM 2009 mini-conference [Wu et al. 2009]. This paper represents a substantial revision and extension, with complete diagnosis of streaming inefficiencies commonly existing over different times, across various streaming channels, and in different geographic regions.

## 7. CONCLUDING DISCUSSIONS

This paper presents the first effort in the research community to extensively explore network-wide streaming inefficiencies that regularly exist in large-scale mesh-based P2P streaming applications, with abundant traces from a real-world P2P streaming system. As one of the highlights in the paper, we have identified an evolutionary pattern of streaming inefficiencies in the system over different times, and its distribution in channels with a smaller population and in geographic regions including small ISPs in China and countries overseas.

Our extensive investigations of possible causes to such streaming inefficiency patterns have brought forward the following intriguing discoveries: (1) The population of peers has two-sided effects on the streaming quality in a streaming channel: while

peers in a larger channel typically enjoy a better streaming quality than those in small channels, the streaming quality degrades with the dynamic peer population increase in the channel during peak hours of a day. (2) The dedicated server capacity still plays an indispensable role in P2P streaming over today's Internet, and the inefficient supply of server capacity has largely led to the low streaming qualities at peak times, in less popular channels, and across regions overseas and in small China ISPs. (3) Smaller channels tend to obtain less server capacity on the per-peer basis than large channels, based on the random server provisioning algorithm that does not explicitly allocate capacity among channels. (4) Inter-peer bandwidth availability represents a more significant bottleneck than the upload bandwidth at the peers during peak hours. (5) Increasing peer indegree does not help enhancing the peer streaming quality, where it is the per-link bandwidth availability that matters.

These discoveries have pointed us to a number of possible directions to enhance such a typical mesh-based P2P streaming protocol:

*First*, the available server capacity should be explicitly and dynamically allocated across different streaming channels, based on their current bandwidth demand and supply at each time; and it is desirable to dynamically adjust the total amount of server capacity deployed in the system, both to meet the demand at peak hours and to save the cost at off-peak times. One possible approach for server capacity allocation is to assign channels different priorities at different times, and maximally guarantee the bandwidth supply to high-priority channels at each time (*e.g.*, a channel with high popularity, or a small channel currently broadcasting some important content), while providing best-effort bandwidth supply to low-priority ones.

*Second*, to enhance the streaming quality in small ISPs and overseas, more server capacity may have to be deployed in these regions, and the key question is to decide the minimum but sufficient amount of server capacity to cover the insufficient portion beyond what the peers can obtain from their remote partners.

*Third*, the locality of partners should be explored especially during daily peak hours and in overseas regions and small ISPs, when and where the inter-ISP bandwidth bottleneck is more evident, *i.e.*, a peer should select more partners within the same AS or ISP in these cases.

*Fourth*, the peer selection protocol should make the acquirement of good partners with high inter-peer bandwidth as its main goal, instead of finding many more partners.

We believe that our findings bring important insights towards complete understanding and fundamental improvement of streaming qualities in large-scale practical P2P streaming. As ongoing work, we are working on the design of better streaming protocols that concrete our suggestions.

## REFERENCES

- AGARWAL, S., SINGH, J. P., MAVLANKAR, A., BACCICHET, P., AND GIROD, B. 2008. Performance and Quality-of-Service Analysis of a Live P2P Video Multicast Session on the Internet. In *Proc. of the 16th International Workshop on Quality of Service (IWQoS 2008)*.
- ALESSANDRIA, E., GALLO, M., LEONARDI, E., MELLIA, M., AND MEO, M. 2009. P2P-TV Systems under Adverse Network Conditions: a Measurement Study. In *Proc. of IEEE INFOCOM*.
- ALI, A., MATHUR, A., AND ZHANG, H. 2006. Measurement of Commercial Peer-To-Peer Live Video Streaming. In *Proc. of Workshop in Recent Advances in Peer-to-Peer Streaming*.

- CERNET BGP VIEW. <http://bgpview.6test.edu.cn/bgpview/curana/ipv4cn/chinaasnlst.shtml>.
- CNNIC. <http://cnnic.cn/en/index/index.htm>.
- CYMRU WHOIS SERVICE. <http://www.cymru.com/BGP/asnlookup.html>.
- GAO, P., LIU, T., CHEN, Y., WU, X., EL-KHATIB, Y., AND EDWARDS, C. 2008. The Measurement and Modeling of a P2P Streaming Video Service. In *Proc. of the Second International Conference on Networks for Grid Applications (GridNets'08)*.
- HEI, X., LIANG, C., LIANG, J., LIU, Y., AND ROSS, K. W. 2006. Insight into PPLive: Measurement Study of a Large-Scale P2P IPTV System. In *Workshop on Internet Protocol TV (IPTV) Services over World Wide Web, in conjunction with WWW 2006*.
- HEI, X., LIANG, C., LIANG, J., LIU, Y., AND ROSS, K. W. 2007. A Measurement Study of a Large-Scale P2P IPTV System. *IEEE Trans. on Multimedia* 9, 8 (December), 1672–1687.
- HEI, X., LIU, Y., AND ROSS, K. W. 2007. Inferring Network-Wide Quality in P2P Live Streaming Systems. *IEEE Journal on Selected Areas in Communications, Special Issue on Advances in Peer-to-Peer Streaming Systems* 25, 9 (December), 1640–1654.
- JIA, J., LI, C., AND CHEN, C. 2007. Characterizing PPStream across Internet. In *Proc. of 2007 IFIP International Conference on Network and Parallel Computing Workshops (NPC 2007)*.
- LI, B., QU, Y., KEUNG, Y., XIE, S., LIN, C., LIU, J., AND ZHANG, X. 2008. Inside the New Coolstreaming: Principles, Measurements and Performance Implications. In *Proc. of IEEE INFOCOM*.
- LI, B., XIE, S., KEUNG, G. Y., LIU, J., STOICA, I., ZHANG, H., AND ZHANG, X. 2007. An Empirical Study of the CoolStreaming+ System. *IEEE Journal on Selected Areas in Communications, Special Issue on Advances in Peer-to-Peer Streaming Systems* 25, 9 (December), 1627–1639.
- LU, Y., FALLICA, B., KUIPERS, F. A., KOOLJ, R. E., AND MIEGHEM, P. V. 2009. Assessing the Quality of Experience of SopCast. *International Journal on Internet Protocol Technology* 4, 1, 11–23.
- PPLIVE. <http://www.pplive.com/>.
- PPSTREAM. <http://www.ppstream.com/>.
- SILVERSTON, T. AND FOURMAUX, O. 2006. P2P IPTV Measurement: A Case Study of TVants. In *Proc. of the 2nd Conference on Future Networking Technologies (CoNEXT '06)*.
- SILVERSTON, T. AND FOURMAUX, O. 2007. Measuring P2P IPTV Systems. In *Proc. of the 17th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV 2007)*.
- TVANTS. <http://www.tvants.com/>.
- UUSEE. <http://www.uusee.com/>.
- VU, L., GUPTA, I., LIANG, J., AND NAHRSTEDT, K. 2007. Measurement and Modeling a Large-Scale Overlay for Multimedia Streaming. In *Proc. of the Fourth International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness (QShine 2007)*.
- WU, C., LI, B., AND ZHAO, S. 2007. Characterizing Peer-to-Peer Streaming Flows. *IEEE Journal on Selected Areas in Communications* 25, 9 (December), 1612–1626.
- WU, C., LI, B., AND ZHAO, S. 2009. Diagnosing Network-wide P2P Live Streaming Inefficiencies. In *Proc. of IEEE INFOCOM 2009 Mini-Conference*.