Dissecting User Behaviors for a Simultaneous Live and VoD IPTV System

NING LIU and HUAJIE CUI, Sun Yat-sen University
S.-H. GARY CHAN, The Hong Kong University of Science and Technology
ZHIPENG CHEN, Sun Yat-sen University
YIRONG ZHUANG, Guangdong Research Institute, China Telecom Co. Ltd.

IPTV services deployed nowadays often consist of both live TV and Video-on-Demand (VoD), offered by the same service provider to the same pool of users over the same managed network. Understanding user behaviors in such a setting is hence an important step for system modelling and optimization. Previous studies on user behavior on video services were on *either* live TV *or* VoD. For the first time, we conduct an in-depth large-scale behavior study for IPTV users offering simultaneously live TV and VoD choices at the same time. Our data is from the largest IPTV service provider in China, offering hundreds of live channels and hundreds of thousands of VoD files, with traces covering more than 1.9 million users over a period of 5 months. This large dataset provides us a unique opportunity to cross-compare user viewing behaviors for these services on the same platform, and sheds valuable insights on how users interact with such a simultaneous system.

Our results lead to new understanding on IPTV user behaviors which have strong implications on system design. For example, we find that the average holding time for VoD is significantly longer than live TV. live TV users tend to surf more. However, if such channel surfing is discounted, the holding times of both services are not much different. While users in VoD tend to view HD longer, channel popularity for live TV is much less dependent on its video quality. In contrast to some popular assumptions on user interactivity, the transitions among live TV, VoD, and offline modes are far from a Markov model.

Categories and Subject Descriptors: C.2 [**Computer-Communication Network**] Network Architecture and Design, Network Operations, Distributed Systems

General Terms: Measurement, Experimentation, Human Factors

Additional Key Words and Phrases: IPTV, live TV, VoD, user behavior

ACM Reference Format:

Ning Liu, Huajie Cui, S.-H. Gary Chan, Zhipeng Chen, and Yirong Zhuang. 2014. Dissecting user behaviors for a simultaneous live and VoD IPTV system. ACM Trans. Multimedia Comput. Commun. Appl. 10, 3, Article 23 (April 2014), 16 pages.
DOI: http://dx.doi.org/10.1145/2568194

This work was supported in part by Fundamental Research Funds for the Central Universities (under grant no. 2010620003161035), Hong Kong Research Grant Council General Research Fund (610713), HKUST (FSGRF12EG05 and FS-GRF13EG15), and Hong Kong Innovation and Technology Fund (UIM/246).
Author's addresses: N. Liu (corresponding author), H. Cui, and Z. Chen, School of Software, Sun Yat-sen University, Guangzhou, China; email: liuning2@mail.sysu.edu.cn; S.-H. G. Chan, Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong; Y. Zhuang, Guangdong Research Institute, China Telecom Co. Ltd., Guangzhou 510630, China.

DOI: http://dx.doi.org/10.1145/2568194

1. INTRODUCTION

With the penetration of broadband Internet to the home, many service providers have offered large-scale IPTV (Internet Protocol Television) services. In IPTV, video content is digitized and distributed to receivers (often set-top boxes) by means of IP multicasting or unicasting. In contrast to the traditional broadcasting media, this new delivery approach provides users with much richer and more flexible viewing experience for both *live TV* and *Video-on-Demand* (VoD) services.

Many IPTV deployments nowadays have already offered their users both live TV and VoD services. These simultaneous services are offered to the same pool of users, by the same service provider or operator, and in the same infrastructure network. As the video streams share the same network resources, understanding user behavior in such a setting has important implications on its system design and optimization.

Based on realistic measurement data from a deployed commerical system, we seek to understand in this article user behavior for a large-scale IPTV system offering both live TV and VoD services for users to navigate and with which to interact. To the best of our knowledge, this is the first piece of work exploring user behavior and access pattern of a *joint* IPTV system offering simultaneous live TV and VoD services. Our results serve as an initial step toward properly planning, maintaining, and analyzing such a system of growing worldwide deployment.

Though there has been much behavior study on IPTV services, the previous works are conducted independently on a system offering either live TV *or* VoD. As these live or VoD systems are studied separately on different network infrastructures, under different network settings, and from different sets of users, it is impossible to draw meaningful comparison on user behaviors of the two services. Studying a joint system where the same pool of users is free to switch to either services any time not only leads to new user behavior issues, but also allows us to compare on the same system setting how user behavior differs for the live TV and VoD services. Comparison questions we may ask are: how do user access patterns (arrival process, user holding time, etc.) differ for live TV and VoD in such a setting? How does the popularity model for live channels compare with that of VoD files? What is the process like for users to switch modes (from VoD to live streaming, and vice versa)? What is the extent of user holding time affected by the video quality (HD and SD) for live TV and VoD?

We study user behavior based on real data provided by China Telecom, the largest joint IPTV service provider in China. The data, of size more than 3TB, was collected over a period of more than 5 months. The IPTV system is a large-scale deployed CDN (Content Delivery Network) in a province of China with more than 1.9 million of subscribers. It offers both live TV and VoD services with similar number of concurrent users at any time. With hundreds of live channels and hundreds of thousands of VoD titles, we trace user behavior for both viewings. Our large set of data from a large number of users, collected from a typically deployed IPTV system, provides us a valuable opportunity and an excellent vantage point to make meaningful observations on user behaviors of such a setting. (Because the IPTV system is run on a managed private network with little packet loss and delay jitter, the impact of network anomalies on user viewing behavior is not of interest here. Readers interested in such topics may refer to Dobrian et al. [2011].)

The salient and unique feature of this work is that our data is from the same pool of users being offered joint live TV and VoD services, over the same infrastructure network, from one service provider, and from a typically deployed large-scale IPTV system. The data enables us to trace users engaged in both services, allowing us to compare their behaviors in these services. With the growing popularity and increasing deployment of joint live TV and VoD IPTV systems in many countries, our study sheds light and insights on the system design for service providers. Our results and their design implications are highlighted next.

—*Holding time.* The average holding time of VoD is significantly longer than that of live channels. This is because there is considerably higher channel surfing activity in live TV than VoD, as users in live TV are in general curious about what is being played on other channels. If we discount the surfing time, live TV has similar viewing time as VoD. The result means that optimizing channel zapping latency for live channels is important.

VoD users have lower surfing activities as compared with live TV, because they often know what a video is about by its description. However, only a small fraction of them watch the entire videos from the beginning to the end, even for videos of length of only a few minutes. Such browsing behavior means that it is beneficial to replicate the "prefixes" (the leading few minutes) of movie files and store them close to users [Sen et al. 1999; Chan 2001; Wang et al. 2002], which significantly reduces access traffic in the network and user access delay. This low interactivity also implies that system design for VoD can be greatly simplified by leveraging network broadcasting/multicasting to save network resources, for example, hot movies may be broadcasted/multicasted and "patching" can be used to support limited interactivity [Chan and Tobagi 2001; Chan and Yeung 2002].

—*Request traffic.* The request traffic for live TV or aggregate (i.e., live TV and VoD) is hardly Poisson. However, the VoD traffic may be much better approximated as a pseudo-stationary Poisson process on a time scale (or interval) of around 30 minutes. This implies that some of the queuing analysis based on pseudo-stationary Poisson arrival may be applicable to VoD only for such a time scale. However, for live TV or its shared network with VoD, conclusions drawn based on Poisson assumptions may not hold well. Using such an assumption to design the system may lead to suboptimal performance [Bommaiah et al. 2000].

—*Access popularity.* While the popularity of VoD files may be better modeled as a Zipf distribution, such distribution does not fit well with the popularity of live channels. In comparison, live channels are much better modeled by a geometric distribution. Therefore, Zipf distribution as commonly assumed in literatures should be used judiciously [Yu et al. 2006; Qiu et al. 2009b; Choi et al. 2012]. VoD file popularity is much more "volatile" as compared with the popularity of live channels, especially for those popular files which may drop quite fast over a month. Such finding has strong implication in a replication policy where video popularity often assumes to be known beforehand and remain quite stable.

—*Video quality effect.* IPTV services are currently predominantly of SD quality; HD is still not a major quality offering. For live TV, there is not much correlation between the holding time and video quality (SD or HD). On the other hand, given similar movie length in VoD, users tend to view HD videos longer. This means that for VoD, besides streaming bandwidth, the change in user holding time due to video quality should also be considered in resource planning. We also find that the access popularity does not depend much on the video quality, no matter it is live channel or VoD. This means that content, rather than quality, plays a more important role in video popularity, so long as a certain quality (SD) is reached.

—*Mode switching behavior.* Despite much use of the Markov process in user behavior modeling (see, for examples, Branch et al. [1999], Gopalakrishnan et al. [2010], and Lee et al. [2010]), we find that the user transition behavior between the modes live TV, VoD, and offline is hardly Markov. Due to complex non-Markovian characteristics, it is not easy to model interactive behavior of a joint IPTV system. More sophisticated mathematical tools beyond Markov analysis are called for in designing such a system.

The remainder of the article is organized as follows. In Section 2 we discuss the related work. In Section 3 we describe the IPTV system and the data elements under study. We define behavior parameters and how they are collected and calculated in Section 4. The comparison results on user behavior for live TV and VoD are presented in Section 5. We conclude in Section 6.

2. RELATED WORK

There are two major delivery architectures for video streaming: CDN (Content Delivery Network) and P2P (Peer-to-Peer). Telecommunication operators often provide IPTV service over a managed CDN, hence providing paid entertainment service with QoS guarantee. On the other hand, P2P Internet TV services are often delivered over the public Internet with little QoS guarantee.

For P2P VoD, much attention has been paid to characterize user interactive behavior [Chang et al. 2008; Luo et al. 2009; Gopalakrishnan et al. 2010; Grsun et al. 2011; Ying et al. 2011; Zhang et al. 2011; Ullah et al. 2012; Lee et al. 2012]. Based on the trace data collected by the Vanderbilt University media service, Chang et al. [2008] studied the feasibility of P2P VoD streaming by empirically analyzing user behaviors, including video popularity, request inter-arrival time, and user online duration, etc. Luo et al. [2009] estimated the scalability of P2P VoD service via a substantial trace-driven simulation, which was based on the analysis results of a large number of real traces of user behaviors provided by a famous television station in China, CCTV. There has also been measurement work on P2P live streaming [Hei et al. 2007; Silverston and Fourmaux 2007; Wu et al. 2009a; Wu et al. 2009b]. Hei et al. [2007] studied the packet traces sniffed from a popular P2P IPTV system, PPLive, and presented detailed measurement results on it, including the variance of users watching a channel over time, user arrival and departure patterns, etc.

While the aforesaid works are impressive, they are for either VoD or live applications and have not studied how users behave in a joint system. Moreover, the results in P2P study cannot be directly extended to a CDN IPTV service model which we consider here. For example, it has been shown that users in P2P systems often have long startup delays and playback lags, and the videos are often of lower quality. These lead to lower interactive activities and holding time. In contrast with regular IPTV users, P2P users may be more unwilling to switch channels or videos due to join latency. Such reluctance in channel/video switching is not expected in an IPTV service. Moreover, the main concerns of P2P Internet TV studies have been mainly on the level and amount of content sharing among peers, and how upload bandwidth can be taken exploted. In a CDN-based IPTV system, such peer sharing is minimal and less of a concern. As compared with the previous P2P studies, the user pool under our study is also much larger, and the video quality is much higher. Due to the differences in video quality and service architecture, it is expected that the results from P2P study cannot be extended to CDN IPTV systems.

The work in Smith [2007], Annapureddy et al. [2007], Mu et al. [2012], and Abrahamsson and Bjorkman [2013] presents some possible models for user interactive behavior. However, they are not based on real data. There has also been work on modeling user behaviors for some specific premises such as schools [Chang et al. 2008; Won et al. 2008]. In contrast, we model user behavior and access patterns for a large-scale IPTV system using large volume of real user data. Recent studies on user behaviors of CDN IPTV networks have focused independently on either live TV [Sripanidkulchai et al. 2004; Veloso et al. 2006; Agrawal et al. 2007; Cha et al. 2008; Qiu et al. 2009b, 2009a] or VoD [Cherkasova and Gupta 2004; Yu et al. 2006; Cheng et al. 2008; Choi et al. 2012; Abrahamsson and Bjorkman 2013]. One of the first measurement studies is Yu et al. [2006], where the authors studied the traces of a large CDN VoD system and introduced a modified Poisson distribution that more accurately modeled the user arrival process. More recently, Cha et al. [2008] presented various new findings on user viewing behaviors through analyzing anonymized control messages in a CDN live TV system. Qiu et al. [2009b] extensively studied the channel popularity and captured its distribution and temporal dynamics on a large collection of user channel access data from a commercial CDN IPTV

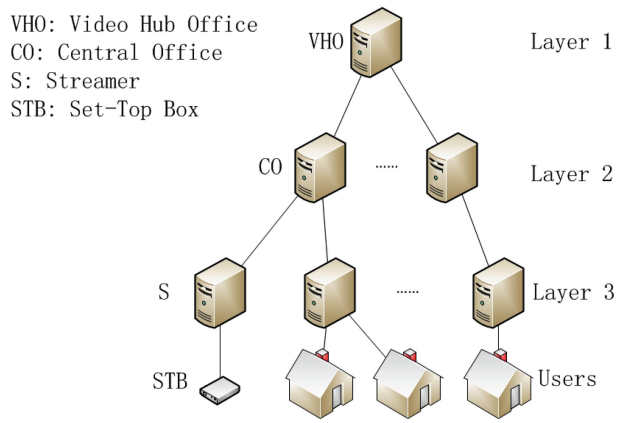

Fig. 1. Content Delivery Network (CDN) for the joint IPTV services.

network. In their subsequent work [Qiu et al. 2009a], they performed an in-depth and overall study on several essential characteristics of IPTV user behaviors but still only focused on live TV sessions. While we have some understanding on the user behaviors of these individual systems, the results cannot be directly extended to a joint system where users are offered both services offered by the same provider in the same network. A joint system also raises new behavior issues which need to be studied, such as the mode switching between VoD and live TV. Given the popularity of such IPTV systems in many countries, our current work is to understand user behaviors in such a setting, and to compare their behaviors in live TV and VoD services. By tracing users in their viewing of both live and VoD sessions, we explore many issues which have not been studied before (such as the behavior of the request traffic for SD and HD videos, cross-examination of user holding times for both services, interactive or switching behaviors between live and VoD services, etc.).

3. SYSTEM DESCRIPTION AND DATA COLLECTION

3.1 IPTV System Overview

The IPTV system of China Telecom provides live TV and VoD services over one seamless managed network. There are over 250 live channels and 250,000 videos available. The users in the system pay a flat monthly subscription fee to enjoy both live and VoD services. These services are mainly SD quality. The subscription fees are higher for more HD-quality videos or channels.

We show in Figure 1 the general hierarchical IPTV CDN of a province from which our data is collected. A user can use a remote control for set-top box to choose a live channel or search for a video to watch by means of menu selection. It is quite convenient for a user to switch from VoD to live streaming and back. The system uses various standard streaming protocols, such as RTSP, RTP, and http streaming, for streamers to communicate with set-top boxes. The communication among servers is based on http. As the network is privately managed, congestion and hence packet loss seldom occurs.

The IPTV network provides two simultaneous services to users.

—*Live TV*. They are popular public channels from cables such as CCTV, Hong Kong TVB Jade, Guangdong Satellite TV, Zhejiang Satellite TV, Jiangzu Satellite TV, etc. The contents are TV programs for the general public. There are no time-shift functions offered to users for these live channels.

The VHO (Video Hub Office) is the primary source of all the live television contents. The video streams are encoded and transmitted from VHO to the end-users via the CO (Central Office) and *S*

Table I. Record Samples

User ID	Content ID	Service Type	Start Time	End Time	Bitrate (kbps)
075540	3221256	Live	2012-06-19 01:58:06	2012-06-19 02:00:11	1500
066308	2691789	VoD	2012-06-20 11:40:12	2012-06-20 12:30:09	1300

(Streamer), where S is the nearest media server to the users. On the user side, there is a Set-Top Box (STB) accessing the IP network. Users can use the STB's remote controller to choose a channel from the channel list. The time that a user joins a channel is recorded, and a trace *record* is generated when the user leaves that channel.

—*VoD.* Regarding the videos, the short ones are mostly music videos while the long ones are mostly films. In our analysis, user interactivity can be detected by a change in video segment requests.

The VHO stores all video files, while CO and S only store a fraction (50%–80%) of videos due to their limited storage. The videos stored in CO is the superset of that in S, which is in turn the subset of that in VHO, namely $S \subseteq \text{CO} \subseteq \text{VHO}$. These servers are accessed in a hierarchical manner. When a user chooses a video, if S stores the video (local hit), it serves the user immediately. Otherwise (local miss), S redirects the request to its upper-tier server and lets it serve the user directly. Depending on the popularity of the video, it stores it locally. Such pulling may lead to some traffic cost, but through such storage update S can keep those most recently or frequently viewed movies locally. As in live TV service, when a user stops or pauses watching a video, a trace *record* is generated.

3.2 Data Description

In order to analyze and model the user behavior in live TV and VoD, we have collected a large amount of trace records from China Telecom's IPTV system from June 3rd 2012 to October 30th 2012, a total of 150 days. There were 1.9 million IPTV subscribers and the number has kept growing. The average number of daily records out of these users is 40 million (a user may generate multiple traces or records in a day), and about 25 million and 15 million records are from live TV and VoD, respectively. The data size of daily records is about 20GB, and we have got more than 3TB records for analysis. The amount of collected samples is large enough to be representative of general IPTV users.

A trace record contains several fields, including the user ID, content ID, service type (indicating which service the record belongs to, i.e., Live or VoD), user's IP, media server's IP, timestamps for the start and end of viewing a channel or video, video's size and length, the bitrate of a channel or video, etc. Table I shows the format of a trace record with the fields we consider in this article.

4. BEHAVIOR PARAMETERS AND ANALYSIS

Based on the data, we will study user behavior with respect to the following.

—*Bitrate.* It is the streaming bitrate of a channel or video delivered to users (in bits/s). We will study user behaviors for both SD and HD video quality. We are interested in understanding user holding time and access probability issues with respect to the bitrate of the live channels and VoD videos.

—*Popularity.* It is the probability of a random request accessing a certain channel or video, given by the access number of the channel or video divided by the total access number of all channels or videos. We are interested in how user holding time depends on the popularity, and how well they can be fitted by the Zipf or geometric distribution.

—*Video length.* It is for VoD, and is the video duration during normal playback without user interaction (in minutes). To understand user interactivity, we are interested in user holding time with respect to video length. Clearly, a holding time much different from video length means high interactivity, and vice versa.

—*Request rate.* It is the number of requests per second. We are interested in the arrival process for live channels, VoD, and the aggregate. We compare the distribution of the request process with a pseudo-stationary Poisson process of interval of T minutes. In other words, we cut the day into N periods of T minutes. For time period i, we calculate the average arrival rate λ_i (requests/s). We use Maximum Likelihood Estimation (MLE) to determine the parameter of the Poisson distribution. We calculate the fitting error E_i, defined as

$$E_i = \sum_{X=0}^{M_i} |P_i(X) - P_i^{'}(X)|, \tag{1}$$

where M_i is the maximum number of request arrivals in a second in the period i, $P_i(X)$ is the Poisson probability of X arrivals with the rate λ_i, and $P_i^{'}(X)$ is the actual probability for the X arrivals in the data. The average fitting error E of time granularity T is hence

$$E = \frac{1}{N}\sum_{i=1}^{N} E_i. \tag{2}$$

We are interested in the error E for different T to understand how well the request process may be modeled as Poisson.

—*Holding time.* It is the continuous duration of holding a channel or video without switching to other channels or videos. A trace record is a request record to a channel or video, and the difference between the end time and start time in a record is the holding time of the user. Depending on the length of the holding time, we further classify the behavior into two categories.

(i) *Browsing/Surfing.* If the holding of a request is less than some short time (say 5 to 30 seconds), we consider it as channel or video browsing/surfing. The user is "shopping" for a stream to watch, and there is no serious engagement in the stream.

(ii) *Viewing.* Otherwise, we consider the request in viewing state. The user is engaged in the channel and video.

Besides understanding the distribution of holding time (and correspondingly browsing and viewing time) of live channel and VoD, we are also interested in the impact of different parameters on holding time, such as video bitrate, popularity, or length.

—*User Mode*. Users in the system may be in any one of three modes as follows.

(i) *Live Mode*. When a user is watching or browsing live TV, he is in live mode.

(ii) *VoD Mode*. A user is in VoD mode if he is watching or browsing VoD videos.

(iii) *Offline Mode*. The user enters the offline mode when he leaves the IPTV system.

In an IPTV system, a user is in any one of the three states (modes) at any time: L (live), V (VoD), or O (Offline), and visits the states according to Figure 2. A user stays in these three states with some random (holding) time T_L, T_V, and T_O, respectively, and then transits to any of the other state according to some probability. Figure 2 also shows the possible process of mode transition behavior of a user. At time T_1, the user enters *live mode* and stays for a time T_L. Then he turns to the *VoD mode* and comes back to *live mode* after time T_V. At the end, the user leaves the system at time T_4, staying in the *offline mode* for time T_O. There are two live sessions, one VoD session, and one offline session in the figure.

User interactivity is often modeled as a Markov process. We will study, for each user, the distribution of time in each state and the transition probabilities from mode i to j denoted by P_{ij}, where $i, j \in \{L, V, O\}$ and $i \neq j$. We will study whether T_L, T_V, and T_O are exponential or not, whether they are independent, and hence whether the process can be modeled reasonably as a Markov process.

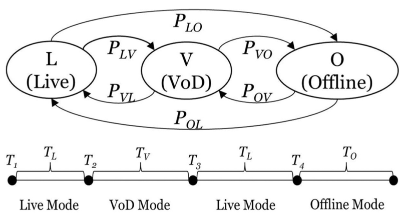

Fig. 2. Diagram of user mode transition.

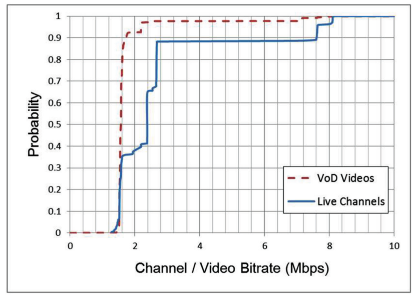

Fig. 3. Bitrate CDF of live channels and VoD videos.

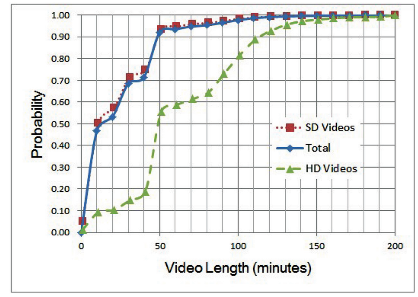

Fig. 4. Video length CDF.

5. COMPARISON STUDY ON USER BEHAVIOR FOR LIVE TV AND VOD

In this section, we take an in-depth cross-study of user behavior for a large simultaneous live TV and VoD system. We discuss and compare the characteristics of user behavior. Our study helps to better understand the different user behavior between live TV and VoD, which contributes to the IPTV system optimization for better user experience.

5.1 System Characteristics

The China Telecom IPTV system provides over 250 live channels and 250,000 VoD videos. Figure 3 shows the Cumulative Distribution Function (CDF) of bitrate for the live channels and VoD videos. The bitrate 1Mbps to 3Mbps belongs to Standard Definition (SD) quality while 7Mbps to 8Mbps is of High-Definition (HD) quality. The figure shows that TV channels and videos are predominantly of SD quality; only a small fraction are HD (around 10%).

We show in Figure 4 the CDF of video length for the total video pool, SD videos, and HD videos (for VoD only). From the total video pool, there is a substantial fraction of relatively short videos (less than 20 minutes). They are mostly news or music videos (MV). There are a considerable number of longer videos ranging from 20 to 50 minutes, which correspond to drama series and entertainment programs. Only a relatively small proportion (around 10%) of videos are long, which are motion pictures. In terms of average video length, HD videos are in general longer than SD videos. This is evident from the plot which shows that a signficant fraction (45%) of HD videos are longer than 50 minutes, while only a small fraction (less than 10%) of SD videos are more than that. The close match between the lines for SD and total video confirms that the system consists of mainly SD videos.

5.2 User Access Traffic

To better understand user traffic in the simultaneous system, we show in Figure 5 the typical hourly requests for live and VoD users over a week (taken from 3rd of June to 9th of June, 2012, where 3rd of June is Sunday). The pattern is typical over all weeks and hence we show here an arbitrary one for illustration purpose. It is clear that the daily access patterns are similar for all the days, both for live TV and VoD. There is no major difference in traffic between weekday and weekend; the hourly requests are only slightly higher at the weekend when there is no work. There is substantial difference between the peaks and troughs (differs by more than an order of magnitude), showing the variability in traffic (as discussed in many previous works).

Figure 6 shows the number of hourly requests and concurrent users over a typical day (on 5th of June, 2012) for live TV (the daily pattern for VoD is similar to live TV, and hence is omitted). In early morning (3AM to 5AM), the number of hourly requests and concurrent users reaches its minimum.

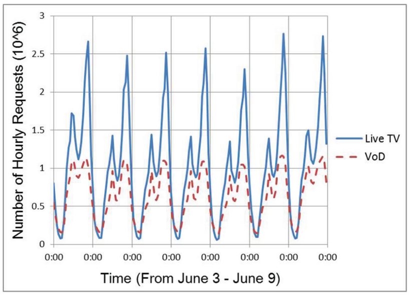

Fig. 5. Number of hourly requests across a week.

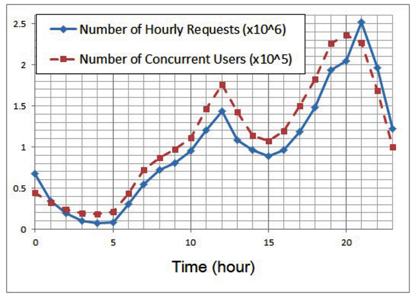

Fig. 6. Number of hourly requests and concurrent users in live TV.

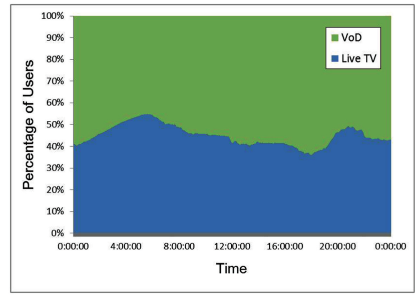

Fig. 7. Proportion of users in live TV and VoD.

There are apparently two peaks: one is in late morning between 11AM to 1PM, and the other one is at night between 8PM to 10PM. This conforms to the user's work-rest schedule. This finding suggests that early morning (say 3AM to 5AM) is the best time for the IPTV system's maintenance and videos' update. Note that the ratio between hourly requests and concurrent users is quite large, meaning that channel zapping frequency is high (channel surfing), that is, the holding time of users is in general short (this is further confirmed in Figures 15 and 16).

Figure 7 shows the proportion of concurrent users watching live TV and VoD in a typical day (5th of June, 2012). The numbers of concurrent users are similar, meaning that the live TV and VoD are equally popular. Between early morning (12AM to 3AM) and early evening (5PM to 7PM), there are more VoD users, because unpopular programs are often played in most live channels at that time.

It is clear from Figure 5 and Figure 7 that the ratio between hourly requests and concurrent users is much smaller in VoD, which means that a user's video switching frequency of VoD is lower.

5.3 Poisson Traffic Modeling

A Poisson distribution is often assumed or used in modeling the request arrival process [Zhang et al. 2011]. It is defined as

$$P\{X = k\} = \frac{\lambda^k e^{-\lambda t}}{k!}, k = 0, 1, 2, \ldots, \tag{3}$$

where k is the number of request arrivals in a time period t and λ is the arrival rate (in requests/s). In our analysis, $t = 1$ second, due to the high frequency of user arrivals and request records.

We study here whether the Poisson distribution is a good model for live TV, VoD, and their *combined* traffic. As the request rate is not stationary (refer to Figure 6), we divide the time into slots called time granularity. In each slot, we try to fit a Poisson distribution by assuming a constant λ over the slot (i.e., the so-called pseudo-stationary Poisson arrival process). Using the actual data, we show in Figure 8 the fitting error E given by Eq. (2) versus time granularity T for live TV, VoD, and the combined traffic. The error is quite high for live TV and combined arrival process, no matter what granularity it is. This shows that pseudo-stationary Poisson arrival is not a good model for them. For VoD, on the other hand, the error is the lowest when the granularity is around 30 minutes, meaning that the system may be approximated as pseudo-stationary with such a time scale.

We show in Figure 9 the best-fit Poisson curve for live TV (on 5-minute granularity) between 12:15AM to 12:20AM. We can see that it does not match the points very well. Besides visually showing the misfit, we have confirmed the result using a Kolmogorov-Smirnov (K-S) test. We first use the MLE to determine the parameter of the Poisson distribution. Then, we use the K-S test to examine the fit goodness.

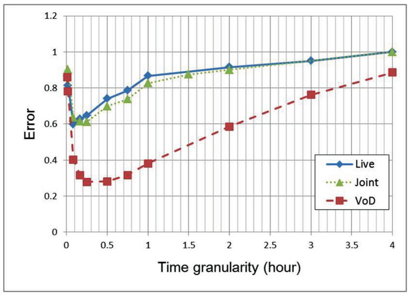

Fig. 8. Error of different time granularity.

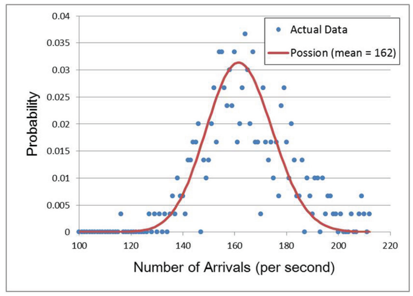

Fig. 9. Actual arrival versus Poisson in live TV.

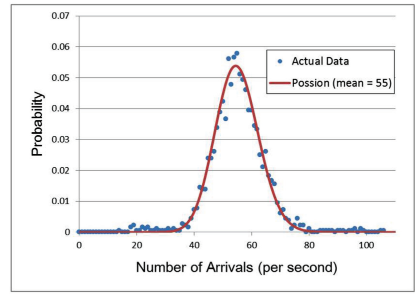

Fig. 10. Actual arrival versus Poisson in VoD.

Our result shows that the actual data cannot pass the K-S test at the 5% significance level, which means the actual arrival model doesn't obey the Poisson distribution well.

For VoD, on the other hand, when the time granularity T is about 30 minutes, the error E is very small. Figure 10 shows the best-fit Poisson curve for such granularity between 3AM and 3:30AM. The curve matches the points quite well. We have also used the MLE and K-S test to examine the validity of the Poisson distribution. The result shows that the model passes the K-S test at the 1% significance level, which indicates strongly that a Poisson distribution models the actual data well.

Our results show that VoD arrival may be modeled as a pseudo-stationary Poisson process if the time granularity is around 30 minutes.

5.4 Channel and Video Popularity

In this section, we model and compare the popularity of live channels and VoD videos. A Zipf distribution is often used to model file accesses and media accesses in the Internet. It is not clear whether the distribution can be used for live TV and VoD IPTV system. If we sort the video popularity in decreasing order and P_k is the popularity of the *kth* video, the Zipf distribution says that

$$P_k = \frac{\lambda}{k^{\alpha}}, \tag{4}$$

where λ is the normalizing constant and α is the skew factor.

Figure 11 shows the typical CDF of the channel popularity in a day. We also fit it with Zipf and geometric distributions, whose parameters are determined by MLE. It is obvious that a geometric distribution is a much better model for live TV channels. The actual data demonstrates that only a small fraction (20%) of the channels accounts for the majority (80%) of the accesses.

We show in Figure 12 the corresponding result for VoD in a day. Using the geometric curve to fit the video popularity is far from satisfactory. On the other hand, using Zipf is a much better fit for most videos except the tail (i.e., those unpopular movies), which drops more quickly than Zipf.

We next examine the change of popularity over time for live channels. We choose three groups of channels, namely popular (top 10%), medium (middle 10%), and unpopular (bottom 10%) on 5th of June 2012. We track their access frequencies over time.

We show in Figure 13 the popularity of the channels and videos over 4 months. We see that, for live TV, the channel popularity does not change much over such a long period of time. Our results show that the popular live channels stay popular, while the cold channels keep cold over a sustained period of time.

For VoD, the result is a little different. We show in Figure 14 how video popularity changes over a 4-month period. We can see that the accessing frequency of the popular videos declines with the

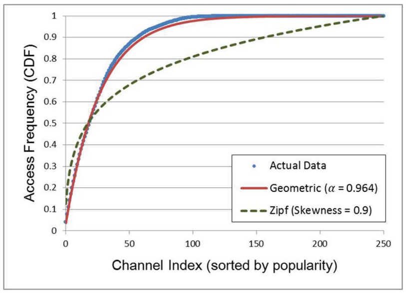

Fig. 11. Fitting the CDF of the channel popularity.

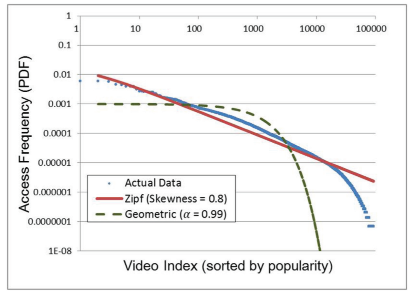

Fig. 12. Fitting the PDF of the video popularity.

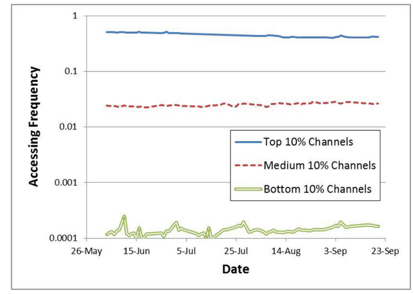

Fig. 13. The change of the channel popularity over 4 months.

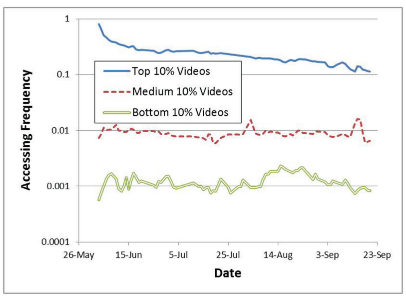

Fig. 14. The change of the video popularity over 4 months.

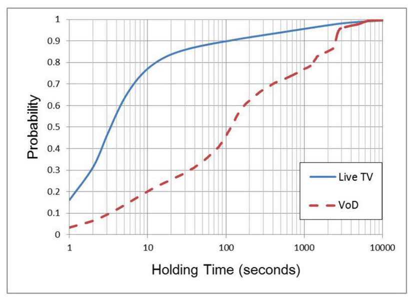

Fig. 15. CDF of holding time for live TV and VoD.

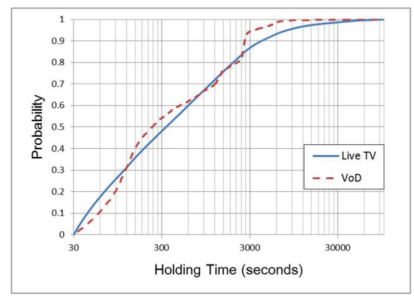

Fig. 16. CDF of holding time for live TV and VoD excluding those records shorter than 30 seconds.

passage of time, while the access frequencies of the medium and bottom ones do not change much. The popularity of the hot movies drops quite fast initially over a month, and keeps dropping over time. This is because new videos are added to the IPTV system every day, and they are apparently more attractive to the viewers than the old ones. That's why the accessing frequency declines for the popular videos.

In summary, a geometric distribution is more suitable to model channel popularity in live TV, while a Zipf distribution is more applicable for video popularity in VoD. Furthermore, channel popularity does not change much over time as compared with video popularity.

5.5 Holding and Viewing Time Comparison

We plot in Figure 15 the CDF of holding time. For live TV, a large fraction of the holding time is shorter than 10 seconds. This is due to *browsing behavior* of the live TV users. The figure shows that the frequency of user channel zapping is high; users tend to be curious about what programs are being played on other channels. The average holding time for live TV is only 6.45 minutes, mainly due to short browsing behavior. For VoD, such surfing behavior is significantly reduced (though there is still a substantial proportion of around 30%). The figure also shows that the average holding time of VoD is much longer than live TV's. That is because video switching frequency is much lower than channel zapping frequency. The reason is that the VoD users are well aware of the content of each video given its title or brief description. Therefore, they are less likely to "browse" through video files.

We next study the *viewing behavior* of live TV and VoD users, by discounting the browsing time. We show in Figure 16 the typical viewing time CDF of live TV and VoD, after filtering the browsing time of 30 seconds (we have studied many browsing values ranging from 15 seconds to 1 minute and the result

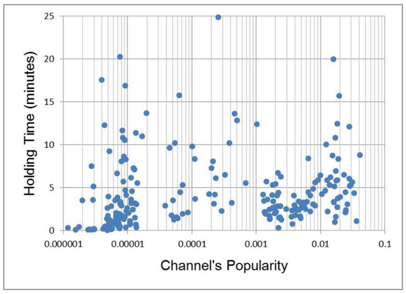

Fig. 17. Correlation between holding time and channel popularity.

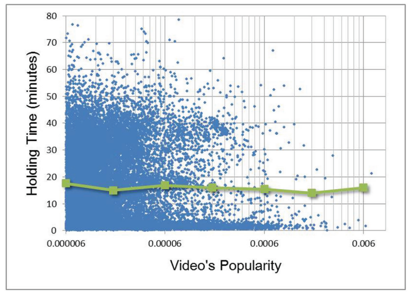

Fig. 18. Correlation between holding time and video popularity.

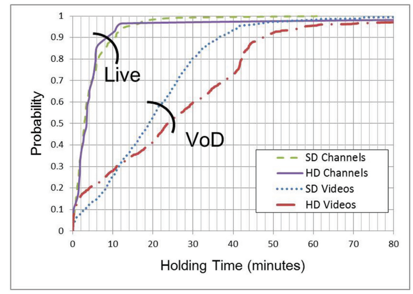

Fig. 19. Holding time CDF of SD/HD channels and videos.

is similar). The two curves are very close, meaning that live TV users and VoD users have similar *viewing behavior*. The average holding time of live TV increases markedly, close to the continuous playback time of a program (of about 22 minutes). What can be inferred is that live TV users prefer to switch the channel when a program ends or during commercial breaks (which aligns well with program time).

5.6 Correlation Study on User Holding Time

In this section, we focus on the correlation between the user holding time and other parameters of popularity, bitrate, or length.

Figure 17 shows the correlation between channel popularity for live TV and user holding time. Given the scattered nature of the plot, there is no strong correlation between channel time and its popularity. To confirm it, we calculate the correlation coefficient between them. Recall that the correlation coefficient ρ_{XY} between two parameters is defined by

$$\rho_{XY} = \frac{\sum_{i=1}^{N}(X_i - \overline{X})(Y_i - \overline{Y})}{\sqrt{\sum_{i=1}^{N}(X_i - \overline{X})^2}\sqrt{\sum_{i=1}^{N}(Y_i - \overline{Y})^2}}. \tag{5}$$

From the data, we obtain the correlation coefficient between holding time and channel popularity as $\rho(\textit{Holding Time, Channel Popularity}) = -0.035$, which is very close to 0. That confirms that their linear correlation is weak. This result may be a bit surprising, as we may tend to think that the holding time of more popular channels is longer. This is not the case due to channel surfing behavior in live TV: while users indeed visit the popular channels more often, they may also surf channels during commercial breaks.

As for VoD, we show in Figure 18 the popularity and corresponding holding time for files of nearly the same length. Due to the large number of videos of similar popularity, we use the average video holding time for each measure of popularity. It shows that the holding time is not sensitive to video popularity. With the data, we get $\rho(\textit{Holding Time, Video Popularity}) = -0.00022$, which means their linear correlation is very weak. The reason is that users choose videos to watch based on their preference. Therefore, the overall popularity of the videos does not affect one video's holding time because what they choose are just what they would like to watch.

We next study how bitrate affects the user's holding time. Figure 19 shows the CDF of holding time for live TV and VoD given HD and SD quality. For live TV, the curves of SD and HD channels are very close (the average holding times of SD channels and HD channels are 6.8 minutes and 7.2 minutes, respectively). This is because, for live TV, users pay more attention to the content, rather than picture

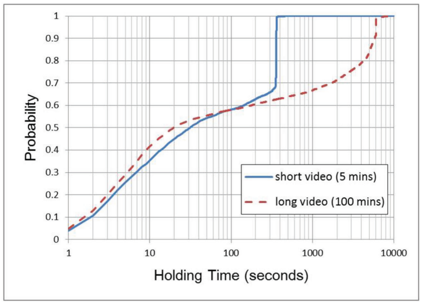

Fig. 20. Holding time CDF for a long video and a short video.

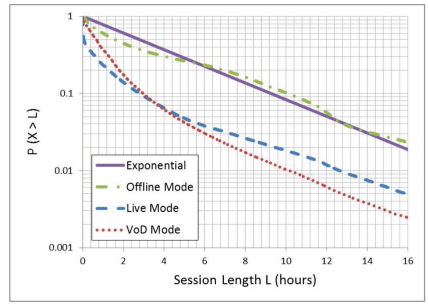

Fig. 21. CCDF of session length for the three user modes.

Table II. Transition Probability between Modes

From \ To	Live	VoD	Offline
Live	–	0.56	0.44
VoD	0.43	–	0.57
Offline	0.57	0.43	–

quality, of a channel (so long as it is SD or above). For VoD, we randomly choose videos from the SD videos and HD videos of similar video lengths (800 videos of 60 to 80 minutes). Then, we plot the holding time CDF for these two types of videos. The holding time for the HD videos is obviously longer than that of the SD ones (the average holding times of SD and HD videos is 20.6 minutes and 31.5 minutes, respectively). This means that VoD users prefer higher picture quality once a certain video content has been chosen.

China Telecom's IPTV service provides over 250 thousand videos with different length, ranging from minutes to hours. We next discuss the correlation between a video's holding time and its length. We consider two groups of randomly chosen videos of length around 5 minutes and 100 minutes, respectively. We plot the CDF of their holding time in Figure 20. The two curves are close to each other for the initial part (due to browsing), and separate when the holding time is close to 5 minutes. The initial match means that a user's browsing and viewing behavior for the two groups of videos are nearly the same. The figure also shows that the holding time for a video is rarely longer than the video length, showing that users do not interact often with the videos. Furthermore, only a small fraction of users watch the entire video from its beginning to the end, though the fraction is higher for the shorter ones.

The findings on user browsing behavior in VoD provides us some design guidelines. For example, given that the majority (around 60%) of the holding time is short (less than 200 seconds as in Figure 20), we may store the "leading" segment of a few minutes (i.e., prefixes) close to users. This would take advantage of the locality effect to improve the access speed and reduce network bandwidth.

5.7 Interactive Mode Switching Behavior

Based on our collected user records, we study the interactive behavior for each user. We show in Table II the transition probabilities among the live mode, VoD mode, and offline modes (refer to Figure 2). The table indicates that users are equally likely to switch to any other mode after they finish a session.

Given that much previous work has assumed user interactive behavior as Markov (see, for examples, Branch et al. [1999], Gopalakrishnan et al. [2010], and Lee et al. [2010]), we next study whether the sojourn times in each mode is independent. We collect the time a user is in a mode (session) before switching and calculate the correlation coefficients between them. The coefficients are all very close to 0, indicating that the sojourn times in the modes are statistically independent.

Regarding the distribution of the sojourn time in each mode, we plot the Complementary Cumulative Distribution Function (CCDF) of session length for the three modes in Figure 21. (The average length for live, VoD, and offline is 1 hour, 1.26 hours, and 3.54 hours, respectively.) Note that if it were Poisson, that is,

$$P(X > x) = 1 - P(X \leq x) = e^{-\lambda x} \tag{6}$$

on a log-y plot, it should be a straight line. Of all the modes, the curve of offline mode can be best fit by an exponential distribution. The other two curves do not match with an exponential distribution. Therefore, the mode switching process cannot be modeled as a continuous-time Markov chain or process, because the distribution of the waiting time until the next transition is not an exponential distribution. In other words, user interactive mode switching in the simultaneous system can hardly be modeled by a Markov process.

6. CONCLUSION

Nowadays widely deployed IPTV services generally offer both live TV and VoD services over the same network to the same pool of users. There has not been work studying user behavior for such a simultaneous system. Understanding such behaviors is an important step to properly engineer and optimize the system.

Using terabytes of trace records collected over five months from a leading IPTV system in China, we have compared and contrasted user behavior of live TV and VoD services in a simultaneous IPTV system. Our results have important implications on design and modeling of the system. We show that the average holding time of VoD is significantly longer than that of the live channels. However, if we discount the surfing time, live TV has similar viewing time as VoD. In VoD, most of the videos are not viewed fully, and there is much lower user interactive activities. This means that prefix storage would lead to good performance. When modeling the request arrival, a pseudo-stationary Poisson process may be used for the VoD traffic, while the arrival process for live TV is hardly Poisson. For live TV, the channel popularity is much better modeled by a geometric distribution, while a Zipf (or Zipf-Mandelbrot) distribution fits better for video popularity in VoD. Users tend to hold HD video longer for VoD than live TV where content plays a more important role than video quality. The process for the user interactive transition among live TV, VoD, and offline modes can hardly be modeled by a Markov process.

REFERENCES

Henrik Abrahamsson and Mats Bjorkman. 2013. Caching for iptv distribution with time-shift. In *Proceedings of the IEEE International Conference on Computing, Networking and Communications (ICNC'13).* 916–921.

Dakshi Agrawal, Mandis S. Beigi, Chatschik Bisdikian, and Kang-Won Lee. 2007. Planning and managing the IPTV service deployment. In *Proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management (IM'07).* 353–362.

Siddhartha Annapureddy, Saikat Guha, Christos Gkanssidis, Dinan Gunawardena, and Pablo R. Rodriguez. 2007. Is high-quality vod feasible using p2p swarming? In *Proceedings of the 16th International Conference on World Wide Web (WWW'07).* 903–912.

Ethendranath Bommaiah, Katherine Guo, Markus Hofmann, and Sanjoy Paul. 2000. Design and implementation of a caching system for streaming media over the internet. In *Proceedings of the 6th IEEE Real-Time Technology and Applications Symposium (RTAS'00).* 111–121.

Philip Branch, Greg Egan, and Bruce Tonkin. 1999. Modeling interactive behaviour of a video based multimedia system. In *Proceedings of the IEEE International Conference on Communications (ICC'99).* 978–982.

Meeyoung Cha, Pablo Rodriguez, Jon Crowcroft, Sue Moon, and Xavier Amatrianin. 2008. Watching television over an ip network. In *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement (IMC'08).* 71–84.

Shueng-Han Gary Chan. 2001. Operation and cost optimization of a distributed servers architecture for on-demand video services. *IEEE Comm. Lett.* 5, 9, 384–386.

Shueng-Han Gary Chan and Fouad Tobagi. 2001. Distributed servers architecture for networked video services. *IEEE/ACM Trans. Netw.* 9, 2, 125–136.

Shueng-Han Gary Chan and Siu-Hong Ivan Yeung. 2002. Client buffering techniques for scalable video broadcasting over broadband networks with low user delay. *IEEE Trans. Broadcast.* 48, 1, 19–26.

Bin Chang, Liang Dai, Yi Cui, and Yuan Xue. 2008. On feasibility of p2p on-demand streaming via empirical vod user behavior analysis. In *Proceedings of the 28th International Conference on Distributed Computing Systems Workshops (ICDCSW'08).* 7–11.

Bin Cheng, Lex Stein, Hai Jin, and Zheng Zhang. 2008. Towards cinematic internet video-on-demand. In *Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems (Eurosys'08).* 109–122.

Ludmila Cherkasova and Minaxi Gupta. 2004. Analysis of enterprise media server workloads: Access patterns, locality, content evolution, and rates of change. *IEEE/ACM Trans. Netw.* 12, 781–794.

Joonho Choi, Abu S. Reaz, and Biswanath Mukherjee. 2012. A survey of user behavior in vod service and bandwidth-saving multicast streaming schemes. *IEEE Comm. Surv. Tutorials* 14, 1, 156–169.

Florin Dobrian, Asad Awan, Dilip Joseph, Aditya Ganjam, Jibin Zhan, Vyas Sekar, Ion Stoica, and Hui Zhang. 2011. Understanding the impact of video quality on user engagement. In *Proceedings of the ACM SIGCOMM Conference (SIGCOMM'11)*. 362–373.

Vijay Gopalakrishnan, Rittwik Jana, Ralph Knag, K. K. Ramakrishnan, Deborah F. Swayne, and Vinay A. Vaishampayan. 2010. Characterizing interactive behavior in a large-scale operational iptv environment. In *Proceedings of the 29th Conference on Information Communications (INFOCOM'10).* 1–5.

Gonca Grsun, Mark Crovella, and Ibrahim Matta. 2011. Describing and forecasting video access patterns. In *Proceedings of the 30th Conference on Information Communications (INFOCOM'11).* 16–20.

Xiaojun Hei, Chao Liang, Jian Liang, Yong Liu, and Keith Ross. 2007. A measurement study of a large-scale p2p iptv system. *IEEE Trans. Multimedia* 9, 8, 1672–1687.

Choonhwa Lee, Euiyoung Hwang, and Dohoo Pyeon. 2012. A popularity-aware prefetching scheme to support interactive p2p streaming. *IEEE Trans. Consumer Electron.* 58, 2, 382–388.

Chae Young Lee, Chang Ki Hong, and Kang Yong Lee. 2010. Reducing channel zapping time in iptv based on user's channel selection behaviors. *IEEE Trans. Broadcast.* 56, 3, 321–330.

Ning Liu, Huajie Cui, Shueng-Han Gary Chan, Zhipeng Chen, Yirong Zhuang, Bing Wang, Subhabrata Sen, Micah Adler, and Don Towsley. 2002. Optimal proxy cache allocation for efficient streaming media distribution. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'02),* vol. 3. 1726–1735.

Jian-Guang Luo, Qian Zhang, Yun Tang, and Shi-Qiang Yang. 2009. A trace-driven approach to evaluate the scalability of p2p-based video-on-demand service. *IEEE Trans. Parallel Distrib. Syst.* 20, 59–70.

Mu Mu, William Knowles, and Nicholas Race. 2012. Understanding your needs: An adaptive vod system. In *Proceedings of the IEEE International Symposium on Multimedia (ISM'12).* 255–260.

Tongqing Qiu, Zihui Ge, Seungjoon Lee, Jia Wang, Jun Xu, and Qi Zhao. 2009a. Modeling user activities in a large iptv system. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement (IMC'09).* 430–441.

Tongqing Qiu, Zihui Ge, Seungjoon Lee, Jia Wang, Qi Zhao, and Jun Xu. 2009b. Modeling channel popularity dynamics in a large iptv system. In *Proceedings of the 11th International Joint Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'09).* 275–286.

Subhabrata Sen, Jennifer Rexford, and Don Towsley. 1999. Proxy prefix caching for multimedia streams. In *Proceedings of the 18th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'99).* 1310–1319.

Thomas Silverston and Oliver Fourmaux. 2007. Measuring p2p iptv systems. In *Proceedings of the International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'07).*

Donald E. Smith. 2007. IPTV bandwidth demand: Multicast and channel surfing. In *Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM'07).* 2183–2187.

Kunwadee Sripanidkulchai, Bruce Maggs, and Hui Zhang. 2004. An analysis of live streaming workloads on the internet. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement (IMC'04).* 41–54.

Ihsan Ullah, Guillaume Doyen, Gregory Bonnet, and Dominique Gaiti. 2012. A survey and synthesis of user behavior measurements in p2p streaming systems. *IEEE Comm. Surv. Tutorials 14*, 3, 734–749.

Eveline Veloso, Virgilio Almeida, Wagner Meira Jr., Azer Bestavros, and Shudong Jin. 2006. A hierarchical characterization of a live streaming media workload. *IEEE/ACM Trans. Netw.* 14, 133–146.

Young J. Won, Mi-Jung Choi, Byung-Chul Park, James W. Hong, Hee-Won Lee, Chan-Kyu Hwang, and Jae-Hyoung Yoo. 2008. End-user iptv traffic measurement of residential broadband access networks. In *Proceedings of the Network Operations and Management Symposium Workshops (NOMS'08)*. 95–100.

Di Wu, Chao Liang, Yong Liu, and Keith W. Ross. 2009a. View-upload decoupling: A redesign of multi-channel p2p video systems. In *Proceedings of the 28th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'09)*. 2726–2730.

Di Wu, Yong Liu, and Keith W. Ross. 2009b. Queuing network models for multi-channel p2p live streaming systems. In *Proceedings of the 28th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'09)*. 73–81.

Qiufang Ying, Yuchun Guo, Yishuai Chen, Xiaoying Tan, and Wei Zhu. 2011. Understanding users' access failure and patience in large-scale p2p vod systems. In *Proceedings of the 4th IET International Conference on Wireless, Mobile and Multimedia Networks (ICWMMN'11)*. 283–287.

Hongliang Yu, Dongdong Zheng, Ben Y. Zhao, and Weimin Zheng. 2006. Understanding user behavior in large-scale video-on-demand systems. In *Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems (EuroSys'06)*. 333–344.

Tieying Zhang, Zhenhua Li, Huawei Shen, Yan Huang, and Xueqi Cheng. 2011. A white-box empirical study of p2p-vod systems: Several unconventional new findings. In *Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN'11)*. 1–6.

Received September 2013; revised December 2013; accepted February 2014