On Next-Generation Telco-Managed P2P TV Architectures

Meeyoung Cha*† Pablo Rodriguez* Sue Moon† Jon Crowcroft‡
*Telefonica Research, Barcelona †KAIST, Korea ‡University of Cambridge, UK

***Abstract*—In recent years, Telcos worldwide have deployed IPTV networks to offer cable TV-like services over the IP backbones. Such walled garden IPTV networks are provisioned to guarantee the quality of service, fast channel switching, and user experience expected by TV viewers. A common key design element of these networks is the use of IP multicast within a single network domain to broadcast several hundreds of TV channels to millions of receivers. However, as the amount of content and channels increase and users demand more interactive and asynchronous viewing, current IPTV architectures are showing clear limitations. In this paper, we study next-generation Telco-managed IPTV architectures, where P2P distributed systems are integrated in Telco's TV set-top boxes or home gateways. We explore how P2P can complement existing Telco-managed IPTV architectures to support advanced rewind functionalities and whether P2P can substitute IP multicast solutions towards supporting a potentially unlimited number of live channels. To this extent, we analyze the TV viewing behavior of a quarter million users using real traces from one of the largest Telco-managed IPTV networks in the world and show the synergistic strengths and the potential for various P2P IPTV combined architectures.**

I. Introduction

Recent years showed a marked uptake of commercial grade live broadcast TV and video-on-demand (VoD) offerings a.k.a. IPTV [2]. Internet Service Providers (ISPs) such as PCCW, Telefonica, France Telecom, AT&T, China Telecom, and Korea Telecom have deployed well-provisioned "walled garden" IPTV architectures in their national backbones with efficient IP multicast as their core building element. As opposed to PC-based viewing, Telco's IPTV services target TV viewing environment integrated with set-top boxes (STBs), providing cable TV-like experience. The first generation IPTV architecture has shown itself to be scalable (serving nearly a million users), secure (providing pay-per view channels and VoD), profitable (as one of the key future sources of revenue), and even cost-effective (compared to other designs). Unfortunately, while existing architectures have proven successful in broadcasting a limited number of channels to large number of synchronous users, they have not been designed to withstand an almost infinite number of niche channels and video content with asynchronous viewing patterns.

To this extent, we study next-generation Telco-managed P2P live TV architectures in this paper. We explore how P2P distribution systems can be integrated into Telco's TV STBs or home gateways to support interactive rewind viewing functions and serve a potentially unlimited number of live channels. In a sense, Telco-managed networks (where content distribution, storage management, and bandwidth provisioning are all controlled and managed by the same ISP) have fewer bandwidth constraints and NAT or end-to-end connectivity challenges than those faced by other P2P-TV applications that are deployed over the public Internet (e.g., Zatoo, Joost, PPLive). However, despite the propitious conditions, the problem of scaling a distribution network to efficiently support millions of high quality TV and video programs with advanced interactive functionalities is not trivial.

For evaluation of designs, we have collected channel switching logs of a quarter million users from one of the world's largest Telco-managed IPTV networks. Our trace permits us to assess future IPTV architectures, using real data rather than speculation. In particular, we find that advanced rewind functions can be easily supported via a P2P distributed hard drive[1], where users can jump back to the beginning of any ongoing programs. We also show that P2P-based live TV architecture enhances scalability to support a potentially unlimited number of live TV channels, while imposing very little demand for equipment at the ISP's data centers.

As our trace is from a Telco network, we give a brief overview of its current IPTV multicast service architecture.

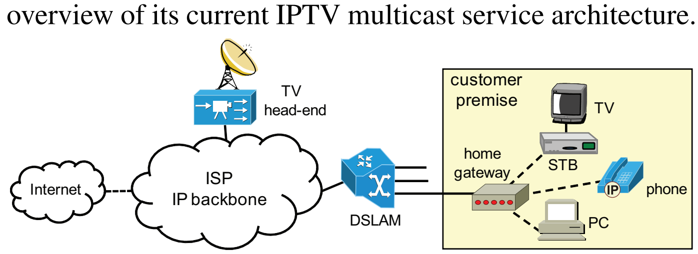

Fig. 1. Typical multicast IPTV service architecture

Figure 1 illustrates a typical IPTV service architecture, where a TV head-end sources IPTV contents to Digital Subscriber Line Access Multiplexers (DSLAMs) through an ISP backbone. Customers subscribe to IPTV, IP phone,

[1] We assume STBs or home gateways are equipped with large storage space. With 2-4 Mb/s of streaming rate, 21-43GB of storage space is required to store one-day-long content.

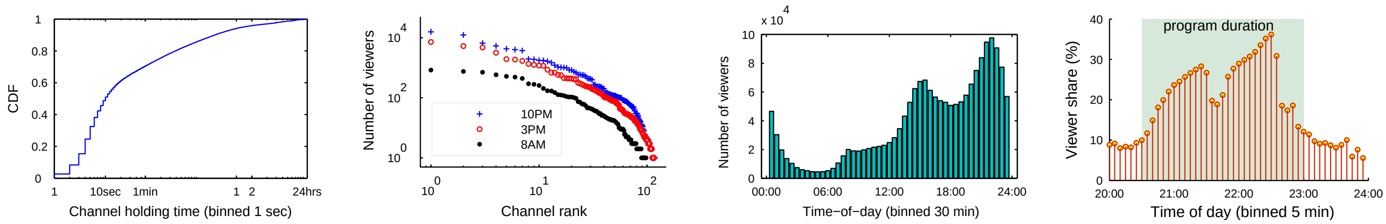

(a) Channel holding time (b) Zipf-like popularity (c) Number of viewers over time (d) Number of viewers over time

Fig. 2. User behavior in a nationwide, Telco-based IPTV system

and Internet access services. DSLAMs aggregate traffic from hundreds or thousands of users and connect to the high-speed IP backbone. For IPTV, the TV head-end streams live broadcast of all the channels across all DSLAMs through bandwidth-provisioned multicast trees [2]. Due to limited access bandwidth at the last mile (i.e., from a DSLAM to the home gateway at the customer premise), not all TV channels are delivered to customers simultaneously. DSLAMs extend or prune multicast trees dynamically toward the customers' STBs, based on the channel switching signals[2]. The STB translates a channel switching event from a user into a pair of Internet Group Management Protocol (IGMP) messages: one to alert the user's departure from the multicast group of a channel and the other to join another multicast group of a new channel. Last mile capacity is split between quality-assured IPTV, IP or circuit switched phone service, and the best-effort Internet access. The ratio of bandwidth for each service and that of upstream versus downstream is completely under the control of ISP.

The rest of the paper is organized as follows. In §II, we analyze diverse aspects of TV viewing habits based on a real trace. §III investigates a cooperative P2P and IP multicast system to provide rewind functions, and §IV further contrasts and compares the IP multicast design with alternate live TV architectures. Finally, we review related work in §V and in §VI, we conclude.

II. Analysis of TV Viewing Habits

IPTV provides a wide range of data and allows for a far more detailed analysis on TV viewing habits than ever before. We examine our trace to understand TV viewing habits that will later be fed into our architectural study.

A. Trace Methodology

We obtained a huge collection of IPTV channel switching logs from an operational backbone provider. Our 200GB trace spans from April 17 to May 30, 2007, and it contains IGMP messages from a quarter million users, including the event type, the timestamp in seconds, and the IP addresses of the DSLAM, the STB, and the multicast group. On any given day, we observe around 200,000 active users across 700 DSLAMs. This is because not all users turn on the TV everyday. The IPTV network serves 155 channels and the number of daily channel switchings clocks 13 million on the average.

B. Channel Holding Time

One difficulty in understanding the user behavior is that viewers leave the STB on, receiving data even when the television is off. In this paper, we focus on the channel holding times, i.e., the time between channel switchings. Channel holding time is affected by the program length, advertisement intervals, and viewer's interest, and it represents user's attention span. Figure 2(a) shows the CDF of channel holding times, where the slope of curve changes around 10 seconds and again around 1 hour, indicating different users' modes such as surfing, viewing, or away.

Accordingly, we consider channel holding times shorter than 1 hour as *active periods* and those longer than 1 hour as *inactive periods*, as active users will likely switch channels during the frequent advertising breaks offered in the channels monitored (normally once every 30 minutes). Based on our notion of user activeness, we observe that an average user watches TV 2.54 hours a day and browses 20 channels.

C. Channel Popularity

Figure 2(b) plots channel rankings against the number of viewers. The distribution is Zipf-like for popular channels, and the popularity decays fast for non-popular channels (ranked 30th or below). This is consistent across different times of day. Zipf-like popularity of top channels has also been reported in P2P-based IPTV services [5].

D. Temporal Correlation

Figure 2(c) shows the time-of-day trend on the number of viewers. We consistently observe strong peaks for particular times of day (e.g., 3PM, 10PM), indicating that a significant fraction of users are correlated about when they watch TV.

[2]We assume one channel is streamed to each customer. In broadband access, multiple channels are carried to support multiple TVs at home.

E. Zooming Into Per-Program Viewing Behavior

Now we focus on how viewers behave over a particular program. Figure 2(d) shows the viewer share of over a sample period of time, which includes a popular soccer game held on May 23, the final match of the European Champions League. The program had four semantic parts: first half, halftime, second half, and the award ceremony. A significant fraction of viewers leave the channel during half-time and re-join later, reflecting strong correlation in their behaviors. Also observe many viewers join or re-join the program after the beginning of each part. Motivated by this, we explore the possibility of feeding latecomers with missed scenes in a P2P fashion in §III. Other types of programs like soap opera and news also showed strong membership amongst viewers.

F. Summary

So far, we have analyzed TV viewing habits and found strong temporal, topical (or channel related), and behavioral correlation amongst users. In the following sections, we assume that any viewed content is stored at user's STB and later used for serving other users via P2P. We further assume that users turn off their STBs and do not store any content during inactive periods.

III. A Cooperative P2P and Multicast System

We consider a cooperative P2P and IP multicast architecture that provides rewind function through a P2P network of STBs (and by video servers if necessary). While our goal is to allow rewind function to any point back in time, for realistic demand, we limit ourselves to the case where users joining a TV program want to jump back to the beginning of it (e.g., a user wants to watch an ongoing movie from the beginning). Our rewind function is different from that provided by personal video recorders (PVR) such as TiVo. PVR provides rewind functions only to those scenes locally stored, while we support rewinding of any newly joined channel.

A. System Design and Assumptions

We make several assumptions. Video content is divided and requested in blocks of 1 minute. P2P is used only for patching passed scenes, while the ongoing scenes are received via IP multicast and stored at user's STB for later viewing. We assume a controlled environment where ISPs provision enough downstream and upstream bandwidth for users to receive both live TV and VCR rewind streams and also upload patching stream to one other peer.

Our P2P system runs as follows: All users joining a new multicast group (channel) rewind to the beginning of the ongoing program. Users request passed scenes to random peers, based on the ISP's tracker that provides an up-to-date viewer list of channels. If a peer possesses the requested blocks and is not serving others, one is deemed available to serve others as a seed peer. If a user fails to find any available seed peer within a given number of trials, then the video server immediately serves the user to assure the rewind quality. The user abandons the video server as soon as new seed peers become available. Our P2P patching stops when the user changes the channel, indicating lack of interest. For each video block downloaded, the user may act as a seed peer for those blocks. We also consider locality-aware P2P system where the ISP's tracker keeps information about which users are watching which channel at each DSLAM and matches local DSLAM peers whenever possible.

B. Trace-Driven Evaluation

For evaluation, we first obtain the TV program schedule and map the program schedule with channel switching logs. In our trace, some of the channels had varying schedules in different geographical regions. Unfortunately, we did not have information about the mapping of network devices to regions. Hence, we have manually identified 30 out of 155 channels with common program schedules and focused on only those channels in our evaluation. These 30 channels account for over 90% of traffic.

Figure 3(a) shows the total amount of traffic required to support VCR operations over time, the amount of traffic that could be served by peers in the local DSLAM, and the load that could not be served by the P2P system. The total traffic shows a time-of-day trend. Under a pure IPTV architecture, all VCR load would be directed to video servers (top-line). With an assistance of a P2P system, the video server load dramatically reduces to 5%. In P2P system, serving capacity for any particular block increases over time, as users requesting a block also become seeds for that block after the download. Hence latecomers in the program can easily find available seed peers.

With locality-aware peer selection, a remarkable 80% of traffic can be served from within a DSLAM. High locality reflects a high correlation in the channels local users watch. This is important, since a large fraction of requests can be handled within the DSLAM domain and the backbone cost can be avoided for this traffic. Naturally, the DSLAM size affects the level of locality. In Figure 3(b), we show the locality efficacy against the size of a DSLAM. The vertical axis represents, for a given DSLAM, the fraction of locally served traffic against the total traffic served for those users within the DSLAM. We observe that locality increases rapidly over the DSLAM size. Note that even for very small DSLAMs, the benefits are noteworthy. Figure 3(c) compares the locality against

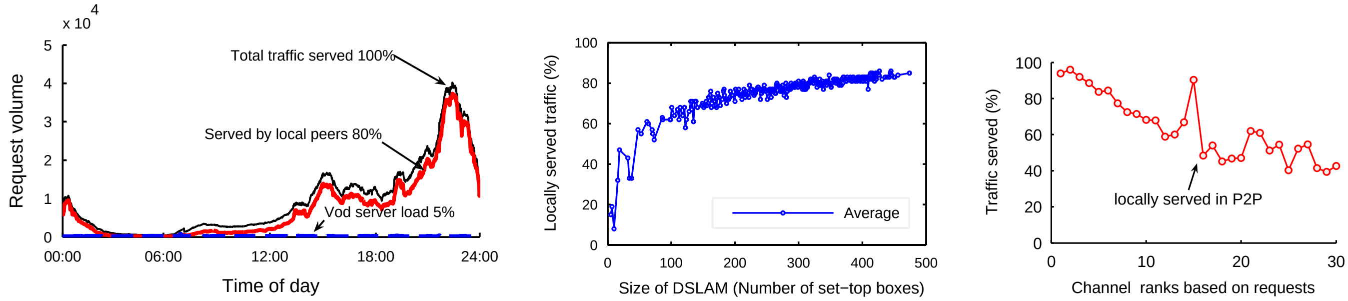

(a) Fraction of locally-carried traffic (b) Locality versus DSLAM size (c) Locality versus channel rank

Fig. 3. Performance of the locality-aware P2P

the channel popularity. We observe that popular channels have a higher chance of being served locally, and even less popular channels enjoy much locality.

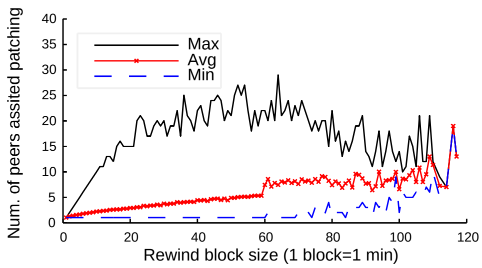

Fig. 4. Communication overhead versus request size

Another metric of interest in P2P system design is the number of peers required for patching. Figure 4 shows the maximum, the average, and the minimum number of peers that were required to patch all past scenes as a function a function of rewind period. We observe that more peers are needed for longer rewind periods. However, this number stays well below 30. The slope of the plot indicates that one has to switch peers once every 11 minutes on the average. This value is very similar to the average channel sojourn time of 12 minutes. The graph becomes noisy after 60 blocks, due to fewer such incidents.

IV. Exploring P2P Live TV Architectures

We now compare different delivery architectures for distributing a massive amount of live TV channels. Current IPTV architectures use static multicast distribution in the backbone, where all the channels are delivered to DSLAMs all the time, even when no user is watching them. As the number of channels increases from a few hundreds to a potentially unlimited number of channels, static multicast-based solutions waste too much backbone bandwidth and clearly do not scale. Herein, we consider the possibility of delivering live TV signals through a network composed of edge devices that cooperate to propagate the TV stream and compare it to an alternative approach which uses a hierarchical structure of servers.

A. *Contrasting Network Impact in Alternate Architectures*

We compare the following Telco-managed network architectures: 1) static multicast trees; 2) dynamic multicast trees; 3) server hierarchy; 4) topology-oblivious P2P over Telco's STBs or home gateways; and 5) locality-aware P2P. Compared to static trees, dynamic multicast trees prune back all the way up to the TV head-end if there are no receivers. In a server hierarchy design, the TV head-end sends out streams to regional servers that are located close to DSLAMs, then the regional servers replicate and forward the live broadcasts to individual users via unicast. We assume that regional servers are co-located with IP routers in the multicast IPTV network. In P2P designs, we assume that a single user per channel receives the TV signal directly from the TV head-end, while the rest of the users stream one another through ISP-provisioned upstream and downstream capacity. Peer selection is random in topology-oblivious P2P. In locality-aware P2P, peers first search for available seed peers within their DSLAMs and then search outside. Other details are similar to §III.

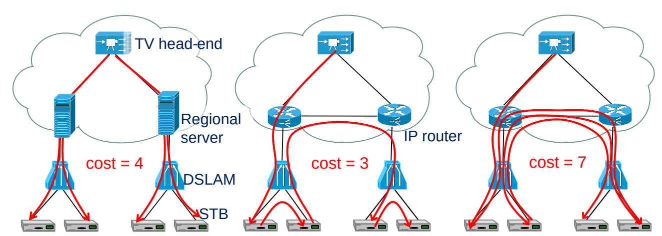

(a) Server hierarchy (b) Locality-aware P2P (c) Topology-oblivious P2P

Fig. 5. Illustrative example of alternative architectures

We contrast the differences across architectures using a toy example in Figure 5. We focus on the total amount of traffic that traverses the boundary of the ISP backbone (i.e. between the first IP router and the DSLAM). As multimedia contents dominate traffic volume, this access link is immediately affected and is likely a top candidate for bandwidth bottleneck. We do not consider the backbone crossing traffic, yet it is also very important; backbone cost depends on the network topology and ISP-dependent

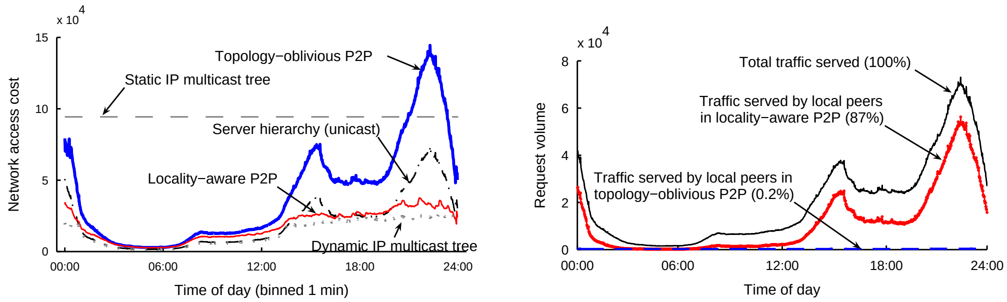

(a) Network cost comparison across all designs (b) Comparison of the two P2P designs

Fig. 6. Comparison of network cost of five different architectures

design choices, and it is orthogonal to our problem.

When all four users request the same channel in Figure 5, the network cost for both static and dynamic IP multicast is 2. In the server hierarchy design, regional servers send out redundant unicast streams to individual users as shown in Figure 5(a), resulting in the cost of 4. In locality-aware P2P, the leftmost peer receives the TV stream from the source, then serves another local peer. A third peer receives streaming from a remote peer, then serves its local peer. We assume that local traffic can be re-routed to peers of a common DSLAM and is not forwarded up to the IP router[3]. As a result, the network cost is reduced to 3 as in Figure 5(b). Figure 5(c) shows the worst-case scenario, where one peer receives TV stream from the source, while the rest of the peers receive from the remote peers. The resulting network cost is 7, which is more than three times higher than that of IP multicast.

B. Network Cost Comparison and Discussions

Now we perform a case study and apply the five architectures to our traces. Figure 6(a) compares the network cost of the five designs. We make several observations. First, the static IP multicast tree consistently requires high cost. This cost is calculated by multiplying the number of DSLAMs and channels. On the other hand, the dynamic multicast tree appears as the all-time most economical design, as expected. Network cost of dynamic tree is equal to the total number of multicast groups across all DSLAMS with at least one active user. Second, the cost of the server hierarchy design shows the time-of-day trend, as this design utilizes unicast transport for individual viewers. This implies a significant server load and scalability concerns during peak usage hours. Finally, a sophisticated locality-aware P2P dramatically reduces the network cost and is comparable to dynamic multicast tree. However, a typical random peer selection approach (which is the case for most existing P2P systems) results in much higher cost than any other designs, generating up to 6.2 times more traffic than the most economical solution (dynamic multicast) during peak usage hours.

To understand the drastic difference between the two P2P architectures, we further examine how much traffic is served by local DSLAM users in Figure 6(b). The average fraction of traffic served by local users shows a stark difference of 87% and 0.2%, respectively. It is interesting that the request volume served locally also follows the time-of-day pattern for locality-aware P2P, which explains how the cost remained relatively steady during peak usage hours in Figure 6(a). In summary, poorly designed P2P TV solutions can have a significant negative impact in the ISP network. Whilst the dynamic multicast solution always provides the best performance possible, locality-aware P2P solution has an overhead that is only slightly higher. In fact, dynamic multicast design represents the lower bound cost of P2P, when we assume that separate P2P systems are built for each DSLAM. This demonstrates high potential for the Telco-managed P2P live TV systems.

One may argue that dynamic multicast design should be the choice for next-generation IPTV services, scaling to massive number of channels. However, dynamic IP multicast is not yet well supported in current routers under massive number of channels. In fact, no IPTV provider has deployed it. This is because there lacks proper support in current routers and that aggregation of trees is not well understood. As a result, if P2P can provide comparable performance to a dynamic multicast solution, it can appear as an appealing alternative, mimicking a dynamic multicast infrastructure from the edge of the network.

Finally, we compare the current static IP multicast design with the dynamic tree design (i.e., P2P lower bound), where an order of magnitude difference exists in their costs. This suggests that the IPTV traffic carried over the backbone is often not consumed by any user and that an alternative design could significantly reduce the backbone traffic. The difference is largely determined by the DSLAM to STBs aggregation ratio (e.g., each DSLAM

[3]DSLAM can operate at layer 2 or 3. When IP layer functionalities are added, they are called IP DSLAM and can support IP routing.

covering few hundreds of users). Moreover, users are not active all the time and users' interest is often focused on few channels, attributing to difference in cost. We have indeed found that the average number of channels viewed per DSLAM is 17% of all channels at any given time. We expect that DSLAM to STB aggregation ratio will be even lower in the future, since more DSLAMs are deployed closer to the end user to reduce the attenuation over the last mile. Therefore, we expect the efficacy of Telco-managed P2P to be even higher in the future.

Several major issues need to be addressed before our cooperative P2P and multicast system becomes a reality. One is the increase of control traffic in the backbone that is used for distributed peer selection. This can be alleviated by providing trackers with a finer-grained view of peers (e.g., if a peer is serving others, when it becomes available again) and fully utilizing locality of peers (e.g., containing P2P traffic within a DSLAM or regions). One extreme solution is to adopt an "oracle" server that helps choose peers based on local proximity and availability. However, we mention that scheduling a globally optimal streaming across multiple channels is a complex open problem. Another concern is the low efficacy of P2P for unpopular channels and during low usage periods, where the role of video servers becomes critical. Yet, we expect large benefits from P2P due to highly skewed channel popularity distribution (see §II-C). According to a recent study, P2P can greatly reduce the overall VoD server load for contents with highly skewed popularity distribution, while benefitting only a small fraction of popular files [3].

V. Related Work

There exist many real implementations and research focusing on P2P live streaming [11] and P2P VoD [7], where diverse issues such as peer churn [6], tree or mesh topology, latency [5], adding DVD functions [10], upload capacity [4], and analytical modeling [8] have been studied. There have also been studies that assume that ISPs actively participate, upstream capacity is provisioned for IPTV, and ISPs make efficient peer data placement as well discussed in [1], [9]. Our work is different from previous work in that we evaluate next-generation P2P-based IPTV architectures under realistic viewing patterns, using a large trace from a live Telco-managed IPTV network.

VI. Conclusion and Future Work

This paper presents the first detailed analysis of traces of a commercial-scale IP multicast TV service and demonstrates the efficacy of Telco-managed P2P TV architectures that support rewind functions and an unlimited number of channels. The strength of our work is at the evaluation based on large-scale IPTV traces. In §III, we show that a significant portion (95%) of rewind requests can be served from P2P compared to video servers due to locality of users. In §IV, we ascertain the sweet spots and the overheads of server-based unicast, multicast, and serverless P2P and also show the empirical lower bound network cost of P2P (where cost reduction is up to 83% compared to current IP multicast distribution). We believe that our work provides valuable insights to service providers in designing the next-generation IPTV architecture. Especially, it highlights that dedicated multicast is useful for few of the extremely popular channels and that P2P can handle a much larger number of channels while imposing very little demand for infrastructure. Yet, further investigation is needed on effective support for niche TV channels.

There are a number of ways in which our work can be extended. First is to consider sophisticated identification of users' active and silent periods. One possible approach is to use program schedules and types, along with the channel switching behavior, and define the probability of user activeness over time. Second is to develop an analytical model that describes when the P2P live TV becomes effective for a given user population and the channel popularity distribution. Finally, we plan to repeat our study for a longer trace period and to use a real implementation to validate the results.

Acknowledgement

Meeyoung Cha did this work while visiting University of Cambridge, as an intern at Telefonica Research. Sue Moon was supported by KOSEF through AITrc.

References

[1] Aggarwal et al. Can ISPS and P2P Users Cooperate for Improved Performance? *ACM SIGCOMM CCR*, 2007.
[2] Cha et al. Case Study: Resilient Backbone Network Design for IPTV Services. In *Workshop on IPTV Services over WWW*, 2006.
[3] Cha et al. I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User Generated Content Video System. In *ACM IMC*, 2007.
[4] Douceur et al. Maximizing Total Upload in Latency-Sensitive P2P Applications. In *ACM SPAA*, 2007.
[5] Hei et al. A Measurement Study of a Large-Scale P2P IPTV System. *IEEE Trans. on Multimedia*, 9(8):1672–1687, 2007.
[6] Huang et al. Can Internet VoD be Profitable? In *ACM SIGCOMM*, 2007.
[7] Janardhan and Schulzrinne. Peer assisted VoD for set-top box based IP network. In *SIGCOMM P2P IPTV Workshop*, 2007.
[8] Kumar et al. Stochastic Fluid Theory for P2P Streaming Systems. In *IEEE INFOCOM*, 2007.
[9] Suh et al. Push-to-Peer Video-on-Demand system: design and evaluation. *IEEE JSAC*, To appear in 2008.
[10] Vratonjic et al. Enabling DVD-like Features in P2P Video-on-demand Systems. In *SIGCOMM P2P IPTV Workshop*, 2007.
[11] Zhang et al. CoolStreaming/DONet: A Data-driven Overlay Network for Peer-to-Peer Live Media Streaming. In *IEEE INFOCOM*, 2005.