Simulation of IPTV caching strategies

Henrik Abrahamsson
Swedish Institute of Computer Science
Box 1263, Kista, Sweden
Email: henrik@sics.se

Mats Björkman
Mälardalen University
Västerås, Sweden
Email: Mats.Bjorkman@mdh.se

***Abstract*—IPTV, where television is distributed over the Internet Protocol in a single operator network, has become popular and widespread. Many telecom and broadband companies have become TV providers and distribute TV channels using multicast over their backbone networks. IPTV also means an evolution to time-shifted television where viewers now often can choose to watch the programs at any time. However, distributing individual TV streams to each viewer requires a lot of bandwidth and is a big challenge for TV operators. In this paper we present an empirical IPTV workload model, simulate IPTV distribution with time-shift, and show that local caching can limit the bandwidth requirements significantly.**

I. INTRODUCTION

IPTV, where TV channels are distributed using IP multicast, has become popular and widespread. Many telecom and broadband companies have become TV providers and distribute TV channels using multicast over their backbone network. IPTV also means an evolution to time-shifted TV where viewers can choose to watch the programs at anytime.

When distributing the TV schedule using multicast there is only one TV stream per channel, while for time-shifted TV there can be one stream per customer. Distributing individual TV streams to each viewer requires a lot of bandwidth and this is a big challenge for TV operators. The operators now therefore only gradually introduce access to more and more time-shifted TV programs, and try out different technical solutions.

TV statistics show that there is a small set of very popular programs that most people are watching. A popular prime-time program that is scheduled and distributed with multicast at a given time, will most likely also have a lot of viewer that choose to watch the program time-shifted a bit later during the evening. For an operator with many hundred thousands of TV subscribers there can be a very large number of copies of the same popular content distributed and putting load on the network.

The question we address in this work is: *to what extent can we limit the bandwidth requirements from time-shifted TV by caching the most popular programs closer to the viewers?* The answer depends on several factors including cache size, caching strategy, and the viewers' request pattern for TV programs.

Caching is a well studied technique for web content and video [1], [2], [3], [4], [5], [6] and have started to attract interest also in the context of IPTV [7], [8], [9], [10].

In order to develop and evaluate IPTV caching strategies good workload models are needed. In this paper we use an empirical IPTV workload model to simulate IPTV distribution with time-shift and investigate the benefit of introducing a local cache closer to the TV subscribers. The simulations are based on real TV schedules, and statistics about TV program popularity and viewer activity. We simulate a large number of TV viewers that, when active, request scheduled or on-demand programs and we investigate the resulting bandwidth requirements on the down link for different cache sizes and caching strategies.

The contributions of this paper are: We present an empirical IPTV workload model. We simulate a realistic scenario for IPTV distribution and compare the Least Recently Used (LRU) and Least Frequently Used (LFU) caching strategies. We show that time-shifted TV can be very capacity demanding and that considerable amounts of bandwidth can be saved by caching the most popular programs closer to the viewers.

The rest of the paper is structured as follows: In section II we describe IPTV distribution and time-shifted TV. In section III we present the IPTV simulator, the workload model, and the simulation scenario. The caching strategies LFU, LRU and Clairvoyant are described in section IV. The simulation results and the evaluation of the caching strategies are presented in section V. Related work is in section VI, future work in section VII and we conclude the paper in section VIII.

II. IPTV AND TIME-SHIFTED TV

We consider IPTV distribution within in a single operator network, where the operator controls the network and how the TV content is distributed.

IPTV operators distribute traditional scheduled TV channels but also gradually introduce new TV services where the viewer can choose to watch a program later after its scheduled time. This service is called time-shifted TV (or sometimes TV on-demand or Catch-Up TV).

A typical IPTV architecture with a hierarchical tree-like network structure is illustrated in Figure 1. The TV content is delivered from content providers and comes into the network at a central distribution center from where it is transmitted to Video Hub Offices (VHO). A Video Hub Office has storage and video streaming equipment to serve a district or a city. Under the VHO there can be intermediate levels of storage and video servers. Different operators try and use different

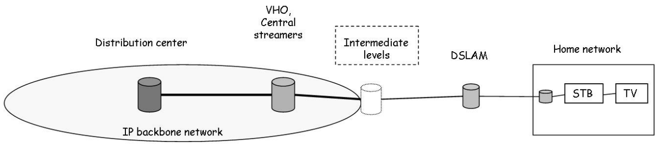

Fig. 1. IPTV network architecture.

structures of varying complexity. The figure also shows a TV subscriber with a home network where the TV and the set-top box (STB) is connected via a residential gateway to a Digital Subscriber Line Access Multiplexer (DSLAM).

The TV channels are distributed using IP multicast from the distribution center to the set-top boxes. TV programs requested outside the schedule are streamed with unicast from the VHO (or from an intermediate server if available) to the set-top box.

III. SIMULATION OF IPTV

A. *Workload model*

We want to investigate the load that IPTV with time-shift can put on a network and how caching can reduce the bandwidth requirements. For this we need a model of the network, a model of how TV is distributed and how TV viewers request programs and put load on the network. We need a TV schedule with channels and programs that is continuously updated, a set of on-demand programs, and a number of viewers that choose programs to watch (either live programs or time-shifted programs).

Our approach to this is to use an empirical model to simulate IPTV distribution. We have implemented a time-driven simulator that operates on the time-scale of minutes. We simulate TV distribution by stepping through real TV schedules and by using statistics about the TV programs' popularity and viewer activity.

B. *Data set*

We use a data set from traditional TV with 13 channels over 5 days from Mediamätning i Skandinavien (MMS) [11]. MMS together with Nielsen Audience Measurement [12] measure the viewing habits of the TV audience in Sweden. The measurements are done using a so called People Meter system where the viewing habits of sample households are logged using electronic meters connected to the remote control.

Our data set include 2225 TV programs from the most popular TV channels in Sweden. For each TV program we extracted the time it was scheduled, its length and the number of viewers. There are a few programs that have a very large number of viewers. The most popular program in this data set have more than 2.3 million viewers (26% of the population) while many of the programs only have a few thousand viewers. The graph in figure 2 shows the number of programs and their share of the total TV viewing time. The top 1% (22 of 2225 programs) most popular programs stand for 26% of the viewed TV time in this data set. The 10% most popular programs stand for 64% of the viewed minutes.

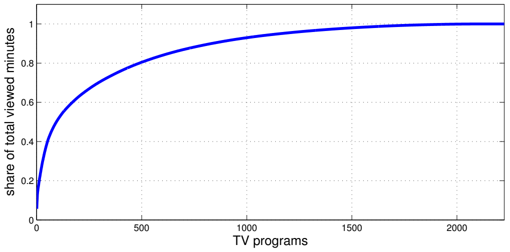

Fig. 2. TV program popularity.

The data set also gives us information about the total number of viewers that are active and watch TV at any given time. Figure 3 shows the fraction of the viewers that are active and how it varies over the five days (Monday to Friday). There are distinct peaks in the evenings when 40-48% of the viewers are active.

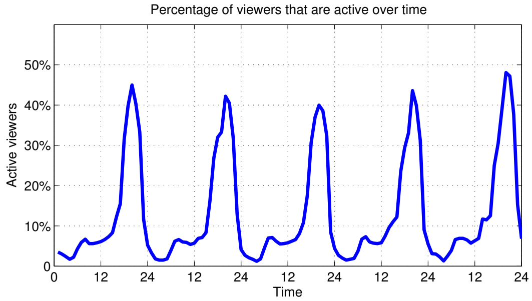

Fig. 3. Percentage of viewers that are active.

C. *TV programs*

In the simulator we represent and step through real TV schedules. For the data set described in III-B we have a schedule with 13 channels over 5 days. So, there is a set of at most 13 ongoing channel programs available at any point in time. The data set also gives us the number of viewers of each program.

In each time step in the simulation we move forward one minute in the schedule, update the set of programs (add new programs and delete the ones that ended), and re-calculate the relative popularity of each program. The latter sets the probability that a viewer will choose to watch a particular TV program.

There is also a set of time-shifted programs that is updated in each step of the simulation. All scheduled programs goes into the set of time-shifted programs and can be requested on-demand. The first minute of a program that is scheduled (and sent out with multicast) at time t is made available for time-shifted viewing at time $t + 1$. The time interval that the programs are available on-demand decides the size of the set of available programs. This a tunable parameter in the simulation. For the experiments described in section V we used a default value of 24 hours.

The popularity of the programs at a given time step in the simulation is illustrated in figures 4 and 5 for the scheduled and time-shifted programs respectively. Here no attempt is made to fit the data to well-known distributions. Instead we generate values from the empirical probability distributions using the inverse transformation method, for instance described in Jain [13]. To choose which program to watch a viewer generate a random number between 0 and 1 from a uniform distribution. For the example in figure 4, if the value is between 0 and 0.46 then program 1 is chosen, if the value is between 0.46 and 0.84 program 2 is chosen and so on.

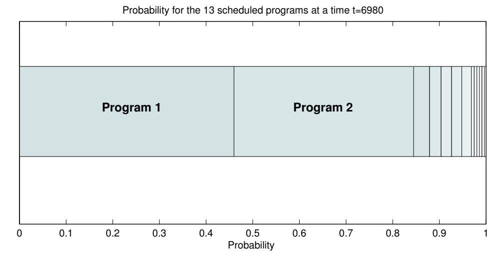

Fig. 4. Probabilities for the 13 scheduled programs at $t = 6980$.

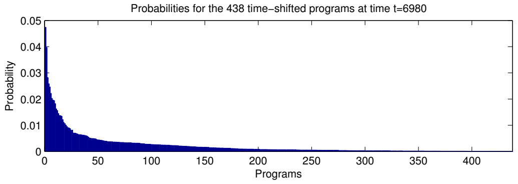

Fig. 5. Probabilities for the 438 time-shifted programs available at $t = 6980$.

D. TV viewers

Our TV viewers are either ON watching TV programs or OFF sleeping. In the simulator we follow the graph from the viewer statistics in figure 3 closely and in each time step adjust the fraction of the viewers that are active and watch TV. A viewer that is activated chooses a program to request. He chooses either to join the distribution of an ongoing scheduled program or to request one of the time-shifted programs that are available on-demand. The particular program to watch is then chosen randomly following the empirical probability distribution for the popularity of the currently available programs. Table I shows parameter settings for the simulations we present in this paper. The share of time-shifted viewing will most likely increase with time when more programs becomes available on-demand and the viewers get used to choosing programs outside the schedule. Here we use a 50% chance that an active viewer chooses to watch a time-shifted program.

E. Network model and simulation scenario

In our simulator we can represent different topologies with caching at different levels in the network including in the set-top boxes. But in this work we delimit the network structure to study the effect of introducing a local cache (in the DSLAM) and the importance of cache size and caching strategy used at this node.

TABLE I
SIMULATION PARAMETERS.

Number of viewers	1000
Number of TV channels	13
Number of TV programs	2225
Programs available time-shifted	24 hours
Simulated time	7200 minutes
Scheduled TV/Time-shifted TV	50/50
TV stream bit rate	2 Mbps

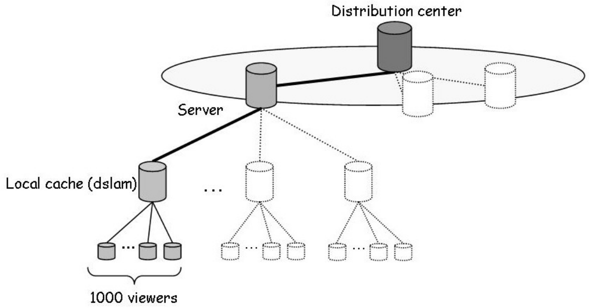

Fig. 6. IPTV network simulation scenario. Studying the effect of introducing a local cache.

For this, we study one branch of an IPTV network topology (as shown in figure 6) with one server, a single local cache, and one thousand viewers (TV set-top boxes). We assume that all programs are distributed to the TV server and that all programs are stored there as long as they are available for time-shifted viewing. For the local cache it is different: what is stored in a local cache at a given moment in time depend on the size of the cache (which is a parameter that we investigate in the simulations), the caching strategy in use; and what programs the viewers under the cache have chosen to watch (the request pattern). We monitor the load on the link from the server to the local cache and we investigate how the bandwidth requirements varies with cache size (including the case with no caching) and caching strategy.

The bit rate of a TV stream depends on the TV channel and codec used. For simplicity we here assume that all TV streams require 2 Mbps.

If a viewer requests a scheduled program, and none of its neighbors is watching this channel, then a multicast stream is added to the load on the link down from the server to the local cache. If someone is already watching the channel, then the new viewer joins the ongoing multicast distribution and no additional load is added to the down link. Requests for time-shifted programs first go to the local cache. If the program is not available there, it is instead transferred with unicast from the server adding 2 Mbps to the load on the down link to the local cache.

IV. CACHING STRATEGIES

When the cache is full, and a new program part arrives, a strategy is needed to decide what should stay in the cache and what to delete. In this work we simulate and compare three different strategies.

A. Least Recently Used

With the Least Recently Used (LRU) strategy we delete from the cache the program that has not been requested for the longest time.

B. Least Frequently Used

With Least Frequently Used (LFU) we discard the program that is requested least often. This could be done by counting the number of viewers that join the multicast distribution of a program and the number of on-demand requests. In the simulation we here use the known popularity of all programs, and delete the one with the least probability for being requested. In addition to that we only consider to delete *inactive* programs; that is programs that no one is watching at this moment in time.

C. Clairvoyant

In the simulation we also implement a clairvoyant strategy with the ability to look into the future and delete the program part that will not be needed for the longest time. This is done by running the simulations twice. In the first run all viewer requests are logged; and in the second run this information is used to determine which program part that will not be asked for for the longest time. The purpose of this strategy is to get an optimal caching strategy and a lower limit to which we can compare the LRU and LFU strategies.

V. EVALUATION

We simulate 5 days of TV distribution in a simple scenario as described in section III. There are 1000 viewers that, when active, request scheduled TV or time-shifted TV. The scheduled TV channels are distributed with multicast via a server; and all programs available for time-shifted viewing are also available from this node. The questions we address are: *how much bandwidth can we save by introducing a local cache (between the viewers and the server)? and how does the result depend on cache size and caching strategy in the local cache?*

Figure 7 shows the link load on the down link during the 5 simulated days. The top figure shows the case without a local cache, where all time-shifted TV are distributed in unicast streams from the core. The middle graph shows the link load when a 25 GB local cache is introduced. The bottom figure shows the result with a cache that is sufficiently large (250 GB) to hold all available time-shifted programs.

The graphs in figures 8 and 9 show the maximum and mean link loads for different cache sizes during the last three days of the simulations. Here we also see a comparison between different caching strategies. The steep slope of the curves show that introducing even a small cache can decrease the load on the down link considerably. The LFU strategy performs better than LRU and is also close to the lower limit set by the Clairvoyant caching strategy.

With a sufficiently large cache, with room for all available time-shifted TV programs, the traffic down from the core to the local cache is low, but not zero. Some viewers are watching the scheduled TV channels that are distributed with multicast from the core. There are also still some unicast transfers of time-shifted programs from the server. This is because, with the caching strategies investigated, a program is only distributed and cached if someone is requesting it. Scheduled programs that none of our 1000 viewers are watching (for instance during night when few viewers are active) are not distributed. If this program is later requested on-demand then, the first time, it is transfered with unicast from the server. This explains why, even with a sufficiently large cache, the link load down to the cache, can exceed that of 13 multicast channels (which would require 26 Mbps with the parameter setting used in the simulations).

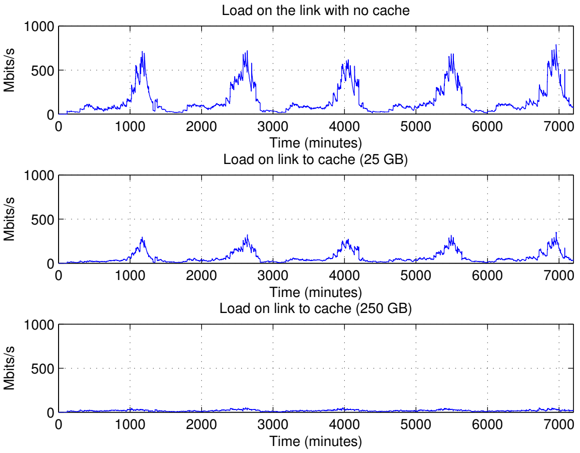

Fig. 7. Link loads over five days for the cases: no cache, 25 GB cache, and 250 GB cache. Here the LFU caching strategy is used.

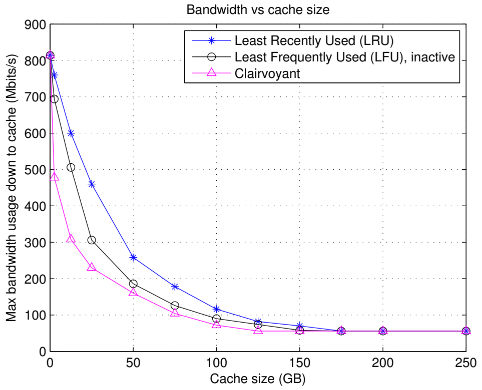

Fig. 8. Comparing maximum bandwidth usage on the down link for different caching strategies and cache sizes.

The use and efficiency of the local cache strategy depend on the request pattern i.e what programs the viewers request and in what order. To compare the impact on the LFU and LRU strategies we did 10 different simulation runs for each cache size. The error bars in figure 10 and 11 show lower and upper values for the resulting maximum and mean link loads on the down link.

The error bars are sometimes overlapping. A closer ex-

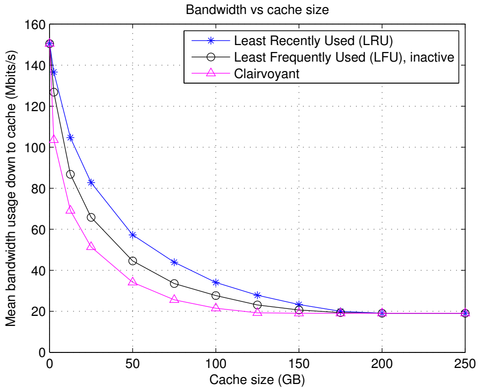

Fig. 9. Comparing mean bandwidth usage on the down link for different caching strategies and cache sizes.

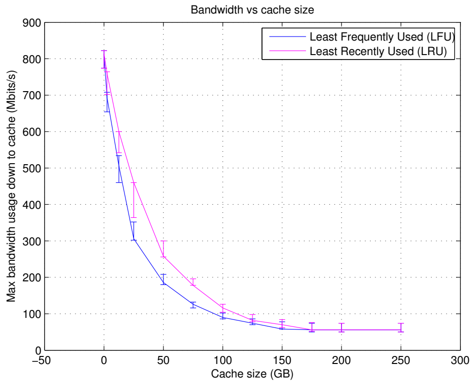

Fig. 10. Maximum bandwidth usage for LFU and LRU with error bars (10 simulation runs).

amination, as illustrated in figure 12, show that this is due to variation between simulation runs. With the same request pattern the LFU strategy always performs at least as good as the LRU strategy in our simulations.

VI. RELATED WORK

The recent growth and popularity of IPTV services have led to an increasing interest from researchers to measure and model IPTV viewing behavior. Cha et al. [14] present an extensive measurement study of viewing behavior including channel popularity and channel switching in an operational IPTV network. Ramos et al. [15] present work on constructing an IPTV workload model capturing the way viewers change channels and watch live TV. Qiu et al. model TV channel popularity [16] and user activities [17] in a large IPTV system and present the SimulWatch workload generator. Their model include set-top box on-times and off-times, channel popularity and channel switching. These studies are similar to ours in that they model IPTV viewer behavior – but they study traditional live TV, and model channel popularity and not the popularity of individual programs. We also simulate TV channels but our focus is on investigating time-shifted TV and caching, and for this the popularity of individual programs is a fundamental part of the model.

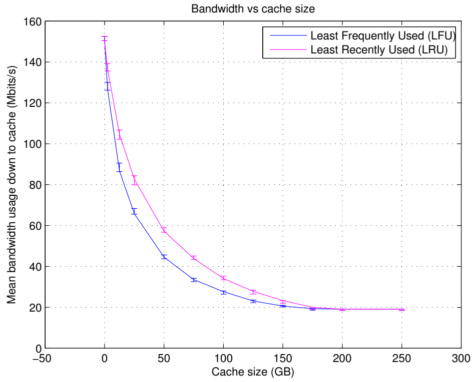

Fig. 11. Mean bandwidth usage for LFU and LRU with error bars (10 simulation runs).

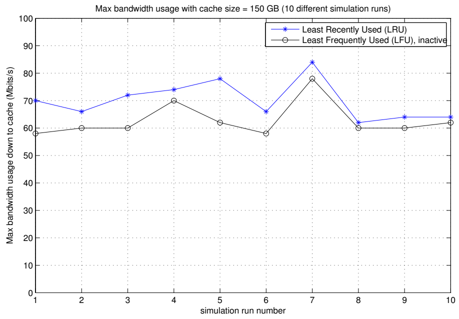

Fig. 12. Detailed look on the maximum bandwidth usage for 10 simulation runs with LFU and LRU and cache size = 150 GB. The LFU strategy is always the better.

Yu et al. [18] measure and model user behavior and content access patterns in a large video-on-demand system. There has been a vast amount of research on different server scheduling techniques and proxy caching strategies and combinations of the two for video-on-demand systems and content distribution networks [1], [2], [3], [4], [5], [6]. These works are similar to ours in that they study methods for minimizing bandwidth requirements for media content distribution and investigate the trade-offs between network bandwidth and cache storage resources. Time-shifted TV has many similarities to VoD but for time-shifted TV the broadcasters' schedules decide when programs are released and influence when and what people watch.

The work closest to ours are the studies by Wauters et al. [9], [10], Vleeschauwer et al. [8], and Krogfoss et al. [7]. They investigate the trade-off between bandwidth usage and storage in scenarios with time-shifted IPTV. But these studies have a more theoretical approach in that they do not use real TV schedules or TV statistics. Wauters et al. present an analytical model [9] and simulations [10] of time-shifted TV with a sliding-interval caching mechanism and co-operative caching. Vleeschauwer et al. [8] study a Catch-Up TV service where the viewers can select to watch the content at a time later than the original airing time. Based on observations from monitoring real TV program popularity they present a user behavior model and simulations where Poisson processes are used to generate the time when programs are aired (corresponding to the TV schedule) and the users' requests for a program. Their conclusions are consistent with ours that caching is needed to limit the otherwise enormous bandwidth requirements when new TV services are fully introduced. Krogfoss et al. [7] investigate several aspects of caching and optimization strategies for IPTV networks including network dimensioning and cache placement.

There are also related work that look at the larger picture of managing a whole IPTV deployment. Mahimkar et al. [19] present work on performance diagnosis in a large IPTV network. Agrawal et al. [20] develop a general framework for planning an IPTV service deployment. Much research has also focused on peer-to-peer techniques for TV distribution [21], [22], [23] and VoD [24], [25], [26]. Cha et al. [21] analyze the TV viewing behavior in an IPTV system and explore how P2P-techniques can complement existing Telco-managed IPTV architectures.

VII. FUTURE WORK

IPTV with time-shift and the use of caching for IPTV are still at an initial stage of development. We have here studied the basic LRU and LFU caching algorithms. For future work there are more complex IPTV scenarios and IPTV caching strategies to investigate including co-operative caches, pre-caching during low traffic and more. Furthermore, the monetary cost of introducing memory into the network versus providing the bandwidth needed is important for operators.

There are also many possible refinements of the simulation model including tuning parameters such as the popularity of time-shifted programs and introducing more complex viewer behavior.

As described in section VI much research has been done on caching for video-on-demand. Time-shifted TV is something different in that we have an initial multicast distribution of all programs and that the broadcasters' schedules decide when programs are released and influence when and what people watch. Also, much of the work on caching for video-on-demand, including sliding-interval, prefix- and segment caching surveyed in [1], assume that only a small part of a program can be kept in memory. The current trend with memory prices going down makes it possible to put much larger caches into the network today than just a few years ago.

In this work we have assumed that all parts of a TV-program have the same popularity. When more detailed viewing statistics become available for time-shifted TV; and if it shows that the popularity of different parts of programs differ a lot, then it could be interesting to re-visit and evaluate more fine-grained segment-based caching algorithms also in the context of IPTV.

VIII. CONCLUSIONS

IPTV is now popular and widespread. Many telecom and broadband companies have become TV operators and distribute TV channels using IP multicast in their network. Operators also gradually introduce new services like time-shifted TV where the viewers can choose to watch the programs later after their first airing.

With a centralized system, unicast distribution of time-shifted programs, and hundred-thousands of subscribers, time-shifted IPTV distribution can be very bandwidth demanding. And since TV statistics show that most people are watching the same programs there can be a very large amount of copies of the same content distributed and putting load on the network. Caching the most popular programs closer to the viewers can significantly reduce the network load, as we show in this paper.

The effectiveness of caching depend on several factors including viewing behavior, request patterns and program popularity. For the development and evaluation of good caching strategies it is therefore important to develop realistic IPTV workload models that include the new time-shifted TV services. In this paper we present a simple IPTV workload model, simulate IPTV distribution with time-shift, and show that caching can limit the bandwidth requirements significantly.

ACKNOWLEDGMENT

This work has been performed within the SICS Center for Networked Systems funded by VINNOVA, SSF, KKS, ABB, Ericsson, Saab Systems, TeliaSonera and T2Data.

REFERENCES

[1] J. Liu and J. Xu, "Proxy Caching for Media Streaming Over the Internet," *IEEE Communications Magazine*, vol. 42, pp. 88–94, August 2004.

[2] D. Eager, M. Ferris, and M. Vernon, "Optimized Regional Caching for On-Demand Data Delivery," in *Proceedings of Multimedia Computing and Networking (MMCN'99)*, San Jose, California, January 1999.

[3] S. Ramesh, I. Rhee, and K. Guo, "Multicast with Cache (Mcache): An Adaptive Zero-Delay Video-on-Demand Service," in *Proceedings of IEEE INFOCOM'01*, Anchorage, Alaska, April 2001.

[4] C. Venkatramani, O. Verscheure, P. Frossard, and K. Lee, "Optimal proxy management for multimedia streaming in content distribution networks," in *Proceedings of the 12th international workshop on Network and operating systems support for digital audio and video (NOSSDAV'02)*, Miami, USA, 2002, pp. 147–154.

[5] O. Verscheure, C. Venkatramani, P. Frossard, and L. Amini, "Joint server scheduling and proxy caching for video delivery," in *Proceedings of the Sixth International Workshop on Web Caching and Content Distribution*, Boston, USA, June 2001.

[6] B. Wang, S. Sen, M. Adler, and D. Towsley, "Optimal proxy cache allocation for efficient streaming media distribution," in *Proceedings of INFOCOM'02*, New York, USA, June 2002.

[7] B. Krogfoss, L. Sofman, and A. Agrawal, "Caching Architectures and Optimization Strategies for IPTV Networks," *Bell Labs Technical Journal*, vol. 13, pp. 13–28, 2008.

[8] D. D. Vleeschauwer, Z. Avramova, S. Wittevrongel, and H. Brueel, "Transport Capacity for a Catch-up Television Service," in *Proceedings of EuroITV'09*, Leuven, Belgium, June 2009, pp. 161–170.

[9] T. Wauters, W. V. de Meerssche, F. D. Turck, B. Dhoedt, P. Demeester, T. V. Caenegem, and E. Six, "Co-operative Proxy Caching Algorithms for Time-Shifted IPTV Services," in *Proceedings of 32nd EUROMICRO Conference on Software Engineering and Advanced Applications (SEAA)*, Dubrovnik, Croatia, September 2006, pp. 379–386.

[10] ——, "Management of time-shifted IPTV services through transparent proxy deployment," in *Proceedings of IEEE Globecom 2006*, San Francisco, USA, November 2006.

[11] "Mediamätning i Skandinavien (MMS)," On–line: http://www.mms.se.

[12] "Nielsen Audience Measurement," On–line: http://en-us.nielsen.com/.

[13] R. Jain, *The Art of Computer Systems Performance Analysis*. New York: John Wiley & Sons, 1991.

[14] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, “Watching Television Over an IP Network,” in *Proceedings of Internet Measurement Conference(IMC)*, Greece, October 2008, pp. 71–84.

[15] F. M. Ramos, F. Song, P. Rodriguez, R. Gibbens, J. Crowcroft, and I. H. White, “Constructing an IPTV Workload Model,” in *Proceedings of SIGCOMM, Poster session*, Barcelona, Spain, August 2009.

[16] T. Qiu, Z. Ge, S. Lee, J. Wang, Q. Zhao, and J. Xu, “Modeling Channel Popularity Dynamics in a Large IPTV System,” in *Proceedings of SIGMETRICS*, Seattle, USA, June 2009, pp. 275–286.

[17] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao, “Modeling User Activities in a Large IPTV System,” in *Proceedings of Internet Measurement Conference(IMC)*, Chicago, USA, November 2009, pp. 430–441.

[18] H. Yu, D. Zheng, B. Zhao, and W. Zheng, “Understanding user behavior in large-scale video-on-demand systems,” in *Proceedings of EuroSys2006*, Leuven, Belgium, 2006, pp. 333–344.

[19] A. Mahimkar, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and Q. Zhao, “Towards Automated Performance Diagnosis in a Large IPTV Network,” in *Proceedings of SIGCOMM*, Barcelona, Spain, August 2009, pp. 231–242.

[20] D. Agrawal, M. Beigi, C. Bisdikian, and K. Lee, “Planning and Managing the IPTV Service Deployment,” in *Proceedings of 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)*, Munich Germany, May 2007, pp. 353–362.

[21] M. Cha, P. Rodriguez, S. Moon, and J. Crowcraft, “On Next-Generation Telco-Managed P2P TV Architectures,” in *Proceedings of International workshop on Peer-To-Peer Systems (IPTPS)*, 2008.

[22] X. Hei, C. Liang, J. Liang, Y. Liu, and K. Ross, “A measurement study of a large-scale P2P IPTV system,” *IEEE Transactions on Multimedia*, vol. 9, pp. 1672–1687, December 2007.

[23] X. Zhang, J. Liu, B. Li, and T. Yum, “Coolstreaming/donet: A data-driven overlay network for efficient live media streaming,” in *Proceedings of IEEE INFOCOM’05*, Miami, FL, USA, March 2005.

[24] C. Huang, J. Li, and K. Ross, “Can Internet VoD be Profitable?” in *Proceedings of ACM Sigcomm 2007*, Kyoto, Japan, 2007, pp. 133–144.

[25] Y. Huang, T. Fu, D. Chiu, J. Lui, and C. Huang, “Challenges, Design and Analysis of a Large-scale P2P-VoD System,” in *Proceedings of ACM SIGCOMM 2008*, Seattle, USA, August 2008, pp. 375–388.

[26] K. Suh, C. Diot, J. Kurose, L. Massoulie, C. Neumann, D. Towsley, and M. Varvello, “Push-to-Peer Video-on-Demand System: Design and Evaluation,” *IEEE Journal on Selected Areas in Communications*, vol. 25, pp. 1706–1716, December 2007.