Caching for IPTV distribution with time-shift

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*—Today video and TV distribution dominate Internet traffic and the increasing demand for high-bandwidth multimedia services puts pressure on Internet service providers. In this paper we simulate TV distribution with time-shift and investigate the effect of introducing a local cache close to the viewers. We study what impact TV program popularity, program set size, cache replacement policy and other factors have on the caching efficiency. The simulation results show that introducing a local cache close to the viewers significantly reduces the network load from TV-on-Demand services. By caching 4% of the program volume we can decrease the peak load during prime time by almost 50%. We also show that the TV program type and how program popularity changes over time can have a big influence on cache hit ratios and the resulting link loads.**

I. Introduction

Many telecom and broadband companies have become TV operators. They distribute TV channels using IP multicast in their networks but also gradually introduce new services like time-shifted TV (or TV-on-Demand) where viewers can choose to watch the programs later, after its scheduled time. Distributing individual TV streams to each viewer requires a lot of bandwidth and server capacity resulting in a big challenge for TV operators.

One way to reduce the network load is to cache popular content closer to the users. Caching is a well studied technique for web content and video [1], [2], but TV is different in many ways.

The potential for caching depends on several factors including user behaviour and content popularity. The new TV-on-Demand services have many similarities with Video-on-Demand systems but there are also some clear differences. In many traditional Video-on-Demand systems there are only a few new releases of movies every week. For TV distribution with time-shift the TV schedule with many ongoing channels gives a constant inflow of new programs that become available for on-demand requests. The program popularity is also different. Many TV programs have a very short lifespan. For instance news programs and weather forecasts quickly become outdated and lose their popularity as soon as a more recent report is made available.

It is important to understand what impact these characteristics of the new IPTV services have on caching. For the development and evaluation of good caching strategies it is also important to develop realistic IPTV workload models that include the new time-shifted TV services and how popularity changes over time.

In this paper we use an empirical IPTV workload model to simulate IPTV distribution with time-shift. The simulations are based on real TV schedules, and statistics about TV program popularity and viewer activity.

The contributions of this paper are: we show that a comparatively small local cache can be used to significantly reduce the peak link loads for TV distribution with time-shift. We also show that in addition to cache size and cache replacement policy, TV program type and how program popularity changes over time can have a big influence on cache hit ratio and the resulting link loads.

The rest of the paper is structured as follows: In Section II we describe our data set and give examples of TV viewing behaviour both from traditional linear TV and time-shifted TV. In Section III we describe the IPTV simulator and the simulation scenario. The simulation results, showing the impact of TV program popularity, cache size, cache replacement policy, and program set size, are in Section IV. Related work is in Section V, and we conclude the paper with a discussion in Section VI.

II. On TV viewing behaviour

A. Traditional linear TV

We use a data set from traditional TV with 13 channels over 28 days from Mediamätning i Skandinavien (MMS) [3]. MMS together with Nielsen Audience Measurement [4] 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 includes 11635 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 is a large variation in the number of viewers between different programs. The median number of viewers of the programs in our data set is 27000. The maximum number of viewers of a program is almost 3.3 millions.

The data set also gives us information about the total number of viewers that are active and watch TV at a given time. Figure 1 shows the fraction of the viewers that are active and how it varies over the first two weeks of the data set. We can see predictable daily and weekly variations in viewer behaviour. There are small increases in viewer activity each morning and there are distinct peaks during prime time every evening when up to 50% of the population is watching TV.

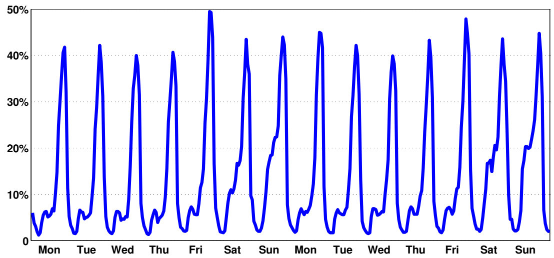

Fig. 1. Fraction of active TV viewers over two weeks. Predictable daily and weekly variations in viewer behaviour.

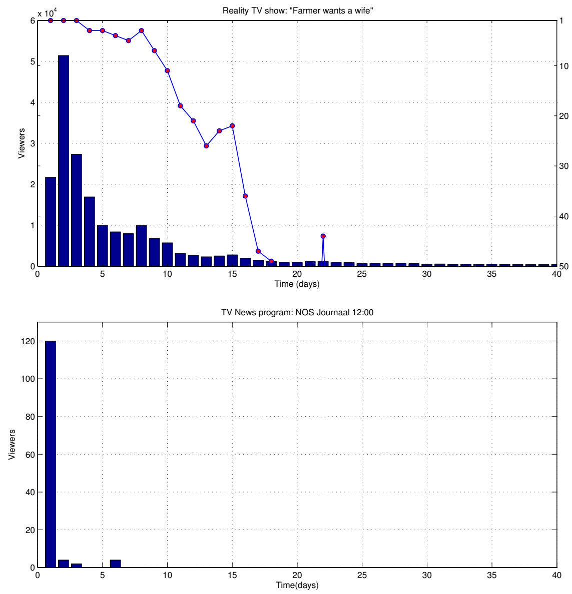

Fig. 2. The top figure shows the number of viewers per day and the rank for an episode of the reality TV show *Farmer wants a wife*. The bar graph shows viewers per day with the scale on the Y-axis shown to the left. The plotted line shows the rank with the scale on the Y-axis shown to the right. The bottom figure shows the number of viewers per day for a TV news program. The figures show how the popularity changes over 40 days after the live broadcast. The popularity of a news program quickly declines.

As expected, Friday and Saturday evenings are the times when most people are watching TV, and we can also see that during the weekend more people are watching TV during daytime.

B. Time-shifted TV

For TV programs that are available on-demand, popularity declines with time. In a TV-on-Demand system, there is also a constant inflow of scheduled TV programs that become available on-demand. Therefore it is not the same programs that are the most popular day after day. Figure 2 shows two examples of how the number of viewers and the rank of programs decrease with time. The examples come from the Dutch TV-on-Demand site Uitzending Gemist [5]. Many programs such as news programs and weather forecasts quickly become outdated and lose their popularity when available on-demand. Other programs, typically drama TV-shows, retain interest from some viewers even a long time after their first release and initial peak in popularity. This categorization is for instance described by Avramova et al. [6] that study and model the popularity evolution of on-demand programs.

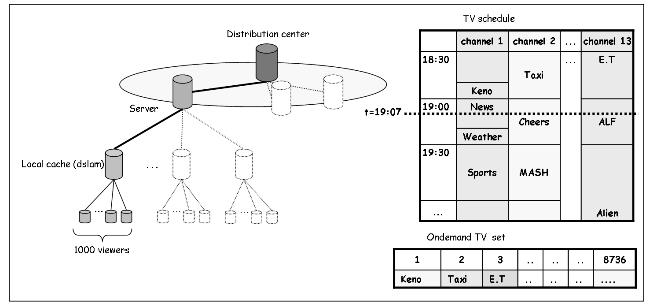

Fig. 3. IPTV simulation scenario.

III. Simulation of IPTV with time-shift

In order to simulate IPTV distribution and evaluate caching strategies we use an empirical model based on the data set described in Section II-A. We simulate IPTV distribution on the time scale of minutes. We have a TV schedule with 13 channels over 4 weeks and statistics about viewer activity and the popularity of the TV programs.

An earlier version of the simulator was presented in [7], and the reader is referred to that paper for details of the simulator that are not included here.

In the earlier version in [7] the TV programs were only available on-demand for a short time (24 hours) and they kept their original popularity. An important extension in the current version is that we simulate how program popularity changes over time and how this depends on program type. This is described in Section III-C.

A. Network model and simulation scenario

We have a scenario with one branch of an IPTV network topology (as shown in Figure 3) with one server and one thousand viewers (TV set-top boxes). We simulate viewer requests for TV-programs and study the effect of introducing a local cache (in the DSLAM), the importance of cache size and cache replacement policy at this node, and the significance of TV program popularity and of the size of the set of available on-demand programs.

The scheduled TV channels are distributed with multicast and all programs then also become available for time-shifted viewing. 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, the caching strategy in use, and what programs the viewers under the cache have chosen to watch (the request pattern). We monitor the cache hit ratio and the load on the link from the server to the local cache and we investigate how these change for different parameter settings.

TABLE I
SIMULATION PARAMETERS.

Number of viewers	1000
Number of TV channels	13
Number of TV programs	11635
Programs available time-shifted	21 days
Simulated time	28 days
Scheduled TV/Time-shifted TV	50/50
TV stream bit rate	2 Mbps
Total program volume (21 days)	5064 GB

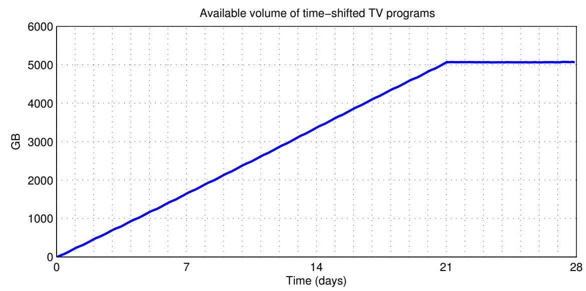

Fig. 4. Available volume of time-shifted programs over 28 simulated days.

B. TV viewers

In the simulator we follow the graph from the viewer statistics in Figure 1 closely and in each time step adjust the fraction of the viewers that are active and watch TV. A viewer that is activated 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 chosen randomly following the empirical probability distribution for the popularity of the currently available programs. Table I shows the parameter settings for the simulations we present in this paper. 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.

C. TV programs

We step through, minute by minute, the TV schedule with 13 channels over 28 days. In each time step we update the set of programs. We also re-calculate the relative popularity of each program using the number of viewers of each scheduled program that we got from the input data set. The latter sets the probability that a simulated viewer will choose to watch a particular TV program.

The set of time-shifted programs is also 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 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 that we investigate in Section IV-B. As a default value we let the programs be available on-demand for three weeks. With this parameter setting we have a steady state after 21 days when programs are deleted from the set of time-shifted programs in the same pace as new ones are scheduled and introduced. Figure 4 shows the available volume of time-shifted programs over 28 simulated days. When evaluating the effect of caching in Section IV we only consider the last simulated week. There are then on average 8736 time-shifted programs to choose between and a program volume of 5064 GB.

Each time-shifted program has a value of popularity which determines the probability that a viewer will choose to watch this particular program. This is initially set to the same value as the program had when scheduled (which is the number of viewers the program had in the input data set). But the popularity of a time-shifted program declines with time. Inspired by the work by Avramova et al [6] we investigated different functions for how TV program popularity changes. We

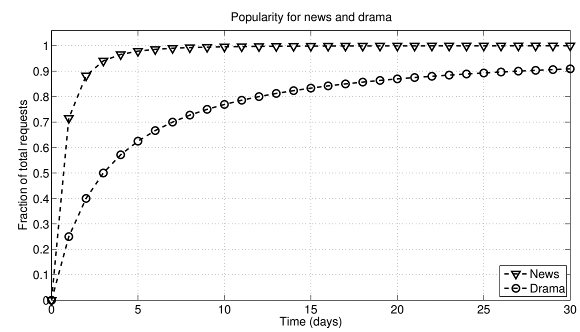

Fig. 5. Popularity over time for programs categorized as news or drama.

categorized each program in our data set as either *News*-like or *Drama*. Programs such as news, business and weather reports, sports, game shows, and morning shows focusing on current events, were all classified into the News category of programs that quickly decrease in popularity. Other programs including movies, TV-series, and documentaries were classified into the Drama category with more slowly declining popularity. The mix of programs differs a lot between different channels. Some offers a lot of news programs, others have only movies and documentaries. With the 13 channels in our dataset 68% of the programs were sorted into the drama category of programs. In the simulations we let the news programs quickly lose their popularity when available on-demand, while the drama programs retain their interest over a longer time following the functions plotted in Figure 5.

D. Cache replacement policy

When the cache is full, and a new program part arrives, a cache replacement policy is needed to decide what should stay in the cache and what to delete. In this work we simulate and compare three classic policies: Least Recently Used (LRU), Least Frequently Used (LFU) and Clairvoyant.

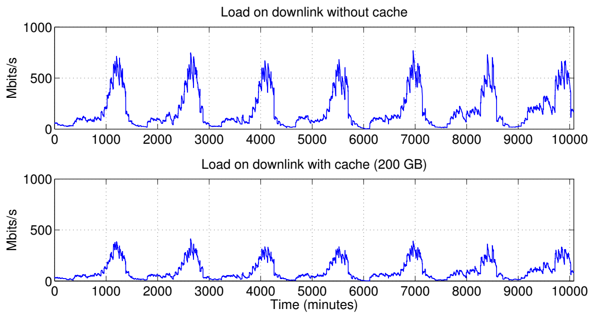

Fig. 6. Load on downlink during 7 days for the cases with no cache (top) and with a 200GB cache (bottom).

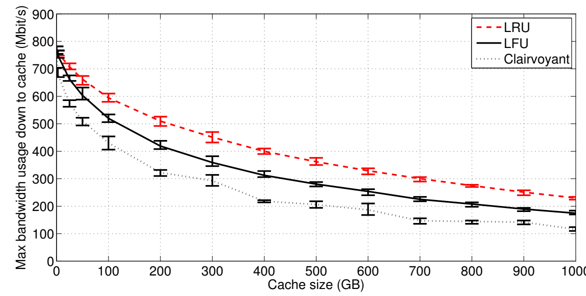

Fig. 7. Comparing maximum bandwidth usage on the downlink for different cache sizes and cache replacement policies. The bars show minimum and maximum values from 5 simulation runs.

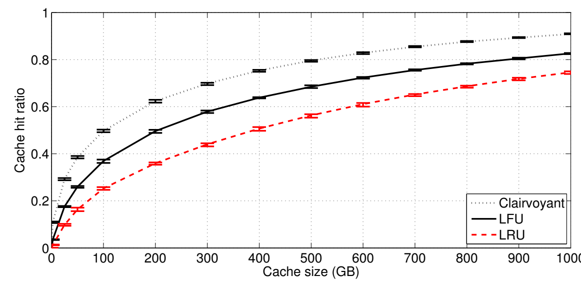

Fig. 8. Cache hit ratio for different cache sizes and cache replacement policies. The total volume of available on-demand programs are on average 5064 GB. A cache size of 100 GB corresponds to 2%, and 500 GB is close to 10% of the available program volume.

With the LRU strategy we delete from the cache the program that has not been requested for the longest time. With 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 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 upper limit to which we can compare the LRU and LFU strategies.

IV. Simulation results

We simulate TV distribution with time-shift in a simple scenario as described in Section III. In Section IV-A we study the effect of introducing a local cache, and how the resulting cache hit ratio and link load depend on cache size and cache replacement policy. Here we use the default parameter settings described in Section III where all programs from 13 channels are available on-demand for three weeks, and all programs are categorized as either news-like or drama and have a popularity that decreases over time correspondingly. In IV-B we vary the time that the programs are available on-demand and study how this impacts on the caching efficiency. In IV-C we investigate what significance program popularity have for caching.

A. Impact of cache size and cache replacement policy

Figure 6 shows the link load on the down link during the last 7 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 bottom graph shows the link load when a 200 GB local cache is introduced using the LFU cache replacement policy. The volume of on-demand TV programs to choose from is on average 5064 GB during the last week of the simulation that we study here. A 200 GB cache can hold 4% of the available program volume and on average it reduces the link load by 45.7%. The highest peak during primetime is reduced by 49.1% from 770 Mbps to 392 Mbps. The graphs in Figures 7 and 8 show the maximum link loads and cache hit ratios (CHR) for different cache sizes. Here we also see a comparison between different cache replacement policies.

We see in Figure 8 that LFU performs better than LRU. Caching 4% (200 GB) of the available programs with LFU gives a 50% hit rate. We can also see that there is still a lot of room for improvement up to the optimal clairvoyant replacement policy. With optimal caching (and a cache size of 200 GB) the hit ratio is increased from 50% to 62% compared to LFU.

B. Impact of on-demand time and program set size

If we change the time period that programs are available on-demand then we also change the size of the set of programs that are available for on-demand requests at a given time. Figure 9 shows the results for the cases when programs are available 1 day, 1 week, and 3 weeks. We can see that even though the program popularity declines with time, the time period that we let the programs be available on-demand has a big impact on the cache hit ratio. In these simulations the LFU cache replacement policy was used. For the case when the programs only are available 1 day there are on the average 417 on-demand programs to choose between in our simulation scenario with an inflow from 13 channels.

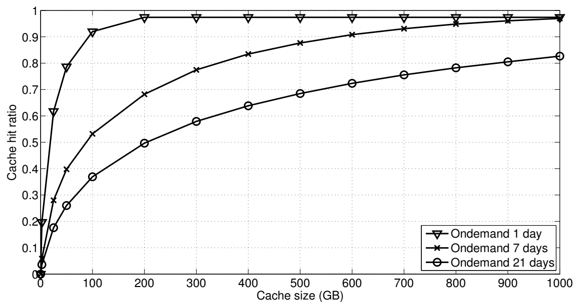

Fig. 9. Comparing cache hit ratio for the cases when programs from 13 channels are available on-demand 1 day, 7 days and 21 days.

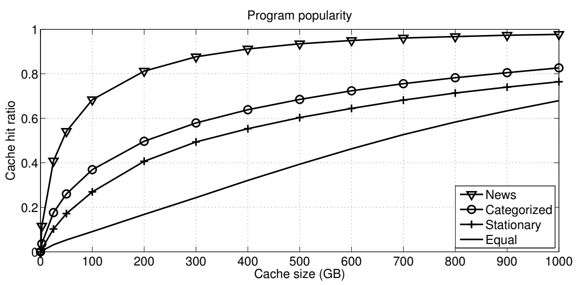

Fig. 10. Influence of program popularity on cache hit ratio. The results for the set of categorized programs with decreasing popularity is compared to stationary program popularity, equal popularity, and with the case where all programs (as news programs) quickly decrease in popularity.

C. Impact of program popularity

The popularity distribution of TV-programs is important for the usefulness of caching. This is also one aspect where TV with time-shift differs from traditional Video-on-Demand in that many popular TV programs (such as news and weather forecasts) have a very short life span.

In order to get an idea of what impact different aspects of program popularity have on the results we compare the cache hit ratio that we get with our work load model with the cases: stationary, equal, and news-like popularity. Figure 10 shows the results. In these simulations the LFU cache replacement policy was used.

In the *equal* case the popularity of all programs are set to the same constant value. In the *stationary* case all programs keep their original values of popularity (which are the number of viewers the programs had when they were first aired) during the time they are available for on-demand requests. In our workload model we also categorize all programs and take into account that program popularity changes over time. A comparison between the stationary and categorized cases in the graph shows what impact this have on the results. In the *news* case we see what the cache hit ratio would be if all programs were news-like and quickly decreased in popularity. For example, with a 200 GB cache (that can hold 4% of the program volume) we get a cache hit ratio of 17% for the *equal* case, we get a 50% hit ratio for our categorized workload model, and for the *news* case 81% of the requests can be handled by the local cache. This shows that the program type, or the type of TV channel that offers on-demand services, has a big influence on the caching results.

V. Related Work

The recent growth and popularity of IPTV services have led to an increasing interest from researchers to measure and model IPTV viewing behaviour. Cha et al. [8] study viewing behaviour including channel popularity and channel switching in an operational IPTV network. Qiu et al. model TV channel popularity [9] and user activities [10] in a large IPTV system and present the SimulWatch workload generator. These studies are similar to ours in that they model IPTV viewer behaviour – 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.

Gopalakrishnan et al. [11] measure and model in detail the interactive user behaviour in an IPTV environment, including how users fast-forward, pause and rewind to control their viewing. There are also many interesting studies of video popularity. Griwodz et al. [12] model long-term popularity of videos on the time scale of days based on VHS rental statistics. Tang et al. [13] analyse and model many aspects of media server access and implement a workload generator. Their model include both static and temporal file popularity and they distinguish between files with regular and news-like lifespan. Kang et al. [14] analyse workload on the Yahoo video sharing site. Gill et al. [15] and Cha et al. [16] present extensive studies of YouTube video sharing. Borghol et al. [17] study the popularity dynamics of Youtube videos. Yu et al. [18] study content access patterns in a large Video-on-Demand system. Lou et al. [19] study the popularity evolution of video files from a Chinese television station and use trace-driven simulation to evaluate caching in a p2p Video-on-Demand system. Dan and Carlsson [20] measure and analyse BitTorrent content popularity. Avramova et al. [6] study and model the popularity evolution of TV-on-Demand and video traces. Szabo and Huberman [21] predict the long-term popularity of online content at Digg and Youtube based on early measurements of user accesses.

Related work on caching include the work by Borst et al. [22] that study caching algorithms for content distribution networks. Wauters et al. [23], Vleeschauwer et al. [24], and Vleeschauwer and Laevens [25] use analytical models and simulations to study the performance of caching strategies in IPTV on-demand systems. These studies have a more theoretical approach and is in this sense complementary to our work. They do not use real TV schedules or TV statistics to run the simulations. Krogfoss et al. [26] also investigate several aspects of caching and optimization strategies for IPTV networks including network dimensioning and cache placement.

Much research has also focused on peer-to-peer techniques for TV and VoD including [27], [28], [29], [30].

VI. Discussion

Our simulation results show that a comparatively small local cache can be used to significantly reduce the peak link loads during prime time. Caching 4% of the on-demand program volume gives a 50% hit rate with the LFU cache replacement policy. The simulation results also show that the program type, or the type of TV channel that offers on-demand services, has a big influence on the caching results. It matters whether we have news programs that quickly become outdated or movies that keep their popularity over a longer time.

For future work more complex IPTV scenarios and caching strategies may be considered. There are also several possible refinements of the simulation model such as separating the popularity of different segments of a time-shifted program and introducing more complex viewer behaviour. The large, predictable, daily variations in user demand described in Figure 1 means that it is important to have the right programs in the cache during prime time. It also makes pre-caching during low traffic an interesting area for further study. Furthermore, the monetary cost of introducing memory into the network versus providing the bandwidth needed is important for operators.

Acknowledgment

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

References

[1] J. Wang, "A Survey of Web Caching Schemes for the Internet," *ACM SIGCOMM Computer Communication Review*, vol. 29, pp. 36–46, 1999.

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

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

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

[5] On–line: http://www.uitzendinggemist.nl/.

[6] Z. Avramova, S. Wittevrongel, H. Bruneel, and D. Vleeschauwer, "Analysis and Modeling of Video Popularity Evolution in Various Online Video Content Systems: Power Law versus Exponential Decay," in *Proceedings of International Conference on Evolving Internet*, 2009.

[7] H. Abrahamsson and M. Björkman, "Simulation of IPTV caching strategies," in *Proceedings of International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS10)*, Ottawa, Canada, July 2010.

[8] 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.

[9] 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.

[10] 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*, Chicago, USA, November 2009, pp. 430–441.

[11] V. Gopalakrishnan, R. Jana, K. Ramakrishnan, D. Swayne, and V. Vaishampayan, "Understanding Couch Potatoes: Measurement and Modeling of Interactive Usage of IPTV at large scale," in *Proceedings of Internet Measurement Conference (IMC'11)*, 2011.

[12] C. Griwodz, M. Bär, and L. Wolf, "Long-term Movie Popularity Models in Video-on-Demand Systems or The Life of an on-Demand Movie," in *Proceedings of the fifth ACM international conference on Multimedia*, Seattle, USA, 1997, pp. 349 – 357.

[13] W. Tang, Y. Fu, L. Cherkasova, and A. Vahdat, "Modeling and Generating Realistic Streaming Media Server Workloads," *Computer Networks*, vol. 51, pp. 336–356, 2007.

[14] X. Kang, H. Zhang, G. Jiang, H. Chen, X. Meng, and K. Yoshihira, "Measurement, Modeling, and Analysis of Internet Video Sharing Site Workload: A Case Study," in *Proceedings of IEEE Conference on Web Services*, 2008.

[15] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "YouTube Traffic Characterization: A view From the Edge," in *Proceedings of ACM SIGCOMM Internet Measurement Conference*, San Diego, USA, October 2007.

[16] M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon, "Analyzing the Video Popularity Characteristics of Large-Scale User Generated Content Systems," *IEEE Transactions on Networking*, vol. 17, pp. 1357–1370, October 2009.

[17] Y. Borghol, S. Mitra, S. Ardon, N. Carlsson, D. Eager, and A. Mahanti, "Characterizing and Modeling Popularity of User-generated Videos," in *Proceedings of IFIP International Symposium on Computer Performance, Modeling, Measurements and Evaluation (PERFORMANCE)*, 2011.

[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] J. Lou, Y. Tang, M. Zhang, and S. Yang, "Characterizing User Behavior Model to Evaluate Hard Cache in Peer-to-Peer Based Video-on-demand Service," in *Proceedings of MMM'07*, 2007, pp. 125–134.

[20] G. Dan and N. Carlsson, "Power-law Revisited: A Large Scale Measurement Study of P2P Content Popularity," in *Proceedings of the 9th international conference on Peer-to-peer systems (IPTPS'10)*, San Jose, USA, April 2010.

[21] G. Szabo and B. Huberman, "Predicting the Popularity of Online Content," *Communications of the ACM*, vol. 53, August 2010.

[22] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proceedings of INFOCOM'10*, San Diego, USA, 2010.

[23] 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.

[24] 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.

[25] D. D. Vleeschauwer and K. Laevens, "Performance of caching algorithms for IPTV on-demand services," *IEEE Transactions on broadcasting*, vol. 55, pp. 491 – 501, 2009.

[26] 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.

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

[28] 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.

[29] 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.

[30] T. Silverston, O. Fourmaux, A. Botta, A. Dainotti, A. Pescapé, G. Ventre, and K. Salamatian, "Traffic analysis of peer-to-peer IPTV communities," *Computer Networks*, vol. 53, no. 4, pp. 470–484, March 2009.