Improving IPTV Services by H.264/SVC Adaptation and Traffic Control

Ingo Kofler, Robert Kuschnig and Hermann Hellwagner
Institute of Information Technology (ITEC)
Faculty of Technical Sciences, Klagenfurt University
Klagenfurt, Austria
Email: firstname.lastname@itec.uni-klu.ac.at

***Abstract*—This paper presents a novel approach that combines both in-network, application-layer adaptation and network-layer traffic control of scalable video streams based on the H.264/SVC standard. In the IPTV/VoD scenario considered, an intercepting RTSP/RTP proxy performs admission control of the requested video, based on the signaled scalability information, and decides whether the content can be streamed without changes or in an adapted version. The proxy configures the network layer appropriately in order to separate the video stream from best-effort traffic on the same link. Rather than performing fixed bandwidth allocation, our proxy approach uses the Hierarchical Token Bucket (HTB) queuing discipline to allow for borrowing bandwidth between traffic classes. In that setting, two different allocation policies are introduced. The Hard Reservation Policy (HRP) performs admission control and adaptation on the video streams and does not modify video bandwidth allocation after admission. In contrast, the Flexible Borrowing Policy (FBP) restricts the admission control to the base layer of the SVC stream. The packets carrying MGS enhancement layer data are marked with priorities by the proxy and are handled at the network layer by a priority-based queuing mechanism. Both a qualitative comparison and an experimental evaluation of the two policies are given.**

I. Introduction

In recent years, substantial research and standardization effort was spent in the context of scalable video coding. This led to the development and standardization of the scalable extension of the H.264/AVC video coding standard, subsequently denoted by H.264/SVC or simply SVC [1]. In contrast to scalable video coding approaches that emerged in the past, the design goal for this standard was to maintain the good compression efficiency of H.264/AVC while introducing scalability mechanisms. Scalability is achieved by introducing a layered encoding of the bitstream. The adaptation of the video stream in the spatial, temporal or/and SNR domains can then be performed by simply truncating layers.

Obviously this kind of video adaptation can be implemented in a very efficient way, especially when compared to adaptation by transcoding. In our previous work, we showed that adaptation of SVC video streams can be even performed on a low-end network device [2]. The adaptation is performed by a proxy application that runs on the network device and adapts the video stream on-the-fly as it is streamed over the network. An important aspect of in-network adaptation is that the adaptation component (e.g., the proxy) has to be signaling and session aware and that the adaptation has to be done on a stateful basis. In the context of IETF standards, such a component is often called a media-aware network element (MANE) [3]. A number of proxy-based approaches for multimedia and in particular video adaptation can be found in the literature [4], [5]. They all have in common that they benefit from the signaling at the application layer and their session awareness. However, they often lack information about the underlying network conditions which can be subject to changes during a video session. Other approaches on the network layer make use of packet priorities which reflect the importance of the payload w.r.t. the video stream [6]. These priorities can then be used for packet scheduling and queuing mechanisms in the network. But most of these approaches are not session aware and rely on priorities that have to be assigned by a dedicated component in the system.

In this paper, we present our approach of a proxy that combines the adaptation of H.264/SVC with traffic control at the network layer. The proxy performs admission control for RTSP requests for video streams and dynamically allocates bandwidth at the network layer. We demonstrate its applicability in an IPTV deployment scenario and the benefits that arise for the subscribers of such an IPTV service. The paper is organized as follows. Section II introduces the deployment of our proxy-based approach in a typical IPTV scenario. Our proxy-based approach and two admission control and bandwidth allocation policies that can be enforced by the proxy are discussed in Section III. An evaluation of the approach in an experimental test-bed is given in Section IV. Section V concludes the paper.

II. IPTV Use Case

The ability to perform cheap in-network adaptation of video streams allows a more flexible provisioning of services in the context of IPTV. The main advantage of IPTV solutions in contrast to traditional terrestrial or satellite-based distribution channels is the existence of a dedicated link to each subscriber. Nowadays, Digital Subscriber Line (DSL) technologies are commonly deployed to serve both IPTV and broadband Internet services [7]. Since the DSL link is not a shared medium but dedicated to a single subscriber [8], it is possible to offer Video-on-Demand (VoD) services with trick-play capabilities.

As a consequence that both the VoD and the broadband Internet service are offered over a single DSL link, it is necessary

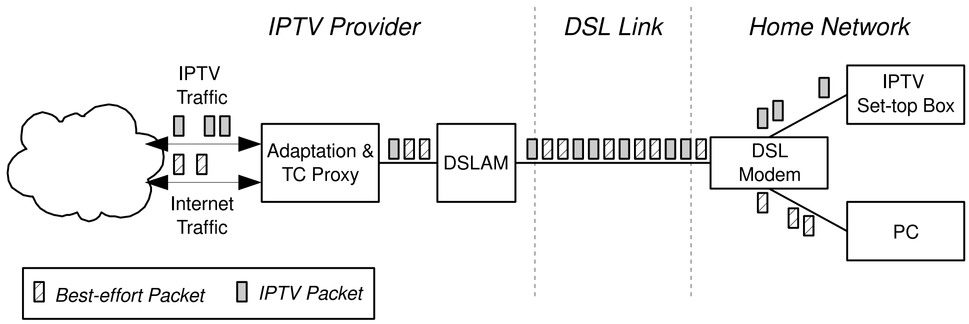

Fig. 1. IPTV deployment

to employ traffic control and QoS mechanisms to separate both services. As we have learned from an experimental analysis of a commercial IPTV solution offered in Austria (aonTV), this separation is often realized as fixed allocation of the bandwidth to both IPTV and Internet services. This allocation prevents the service from maximizing the utilization of the DSL link and from increasing the subscriber's benefit of the services.

In order to overcome this shortcoming of existing solutions, we propose a proxy-based approach that makes use of H.264/SVC adaptation and traffic control (TC). The deployment of the proxy in the IPTV scenario at the IPTV provider's premises is depicted in Figure 1. The advantage of this approach is that the proxy can influence the traffic control and QoS mechanisms in a dynamic way. The key technology utilized for the traffic control is the Hierarchical Token Bucket (HTB) queuing discipline [9]. Additionally, the proxy performs admission control for the Video-on-Demand sessions issued by the IPTV subscriber and can optionally adapt the H.264/SVC-based video streams at the application layer. In the next section, the cooperation of the proxy with the traffic control at the network layer and its implications on the IPTV service is discussed in detail.

III. Adaptation and Traffic Control Proxy

The Hierarchical Token Bucket (HTB) queueing discipline allows to distribute the capacity of a link to different classes of network traffic, e.g., to separate best-effort traffic from real-time traffic. Each traffic class is associated with two numbers - the assured rate (ar) and the ceil rate (cr). The assured rate is the amount of bandwidth that is allocated and assured for that specific traffic class while the ceil rate is the maximum amount of bandwidth that can be assigned to that class by the packet scheduler. If one traffic class does not make use of its assured rate, the remaining (and otherwise unused) bandwidth can be borrowed to other classes. The amount of bandwidth that one traffic class can borrow from other ones is represented by the difference between the ceil rate and the assured rate. The assignment of traffic flows to each of the classes is realized by filters. Filters can be configured to match packets based on some criteria (e.g., transport protocol, IP addresses, port numbers, TOS/DSCP field) and to delegate the matched packets into a certain traffic class. As the name of the HTB queueing discipline implies, the traffic classes can be organized in a hierarchical manner, e.g., a traffic class could be further separated into traffic classes. Furthermore, each class can make use of its own queuing discipline, e.g., FIFO, SFQ, RED.

An example of bandwidth borrowing is illustrated in Figure 2. The root class of the tree represents the actual physical network device which has a bandwidth of 6000 kbps. The bandwidth of the device is distributed among three traffic classes A, B and C with assured rates of 2000, 3000 and 1000 kbps, respectively. Let's assume that class A serves a constant bit rate traffic of 500 kbps, which is only one fourth of its assured rate. The unused bandwidth of 1500 kbps can therefore be borrowed to other traffic classes if they have demand for it. In the example, class B serves a TCP download which could make use of additional bandwidth. Since the ceil rate of class B is higher than the assured rate, it can borrow up to 1000 kbps from other traffic classes. So the packet scheduler will borrow 1000 kbps unused bandwidth of class A to class B. However, class C would never borrow bandwidth since its assured and ceil rate values are identical.

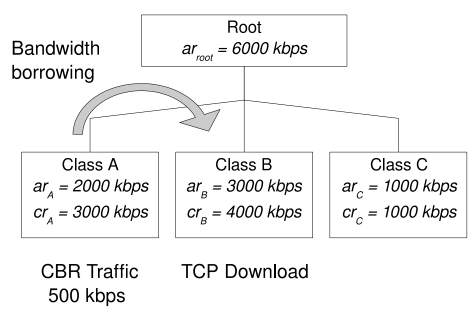

Fig. 2. HTB borrowing

In our approach the adaptation proxy is both responsible to handle the RTSP/RTP connections at the application layer and to manage the traffic classes on the network layer. The adaptation proxy intercepts each RTSP request automatically and can therefore be categorized as intercepting proxy. Since the proxy is deployed at the ISP's end of the DSL link, it is assured that each request of the IPTV subscriber will pass through the adaptation proxy. The decision, which actions the proxy performs on the RTSP sessions (e.g., adaptation, admission control), depends on the policy that is actually engaged. As the proxy is involved in the RTSP-based signalling of the multimedia session, it is also aware of the session description which is typically offered by the streaming server as a response to an RTSP DESCRIBE request by the client. The predominant standard for describing multimedia sessions is the Session Description Protocol (SDP). In our approach we utilize the *sprop-scalability-info* parameter in the SDP which is used to encapsulate a Scalability SEI message. By parsing this SEI message, the proxy is able to figure out which scalability dimensions are offered by the stream. Additionally, the SEI message contains a mapping between the adaptation parameters (TID, DID, QID) [10] and the resulting bit rate of the video stream. The proxy makes use of this information for both performing admission control at the application layer and for allocating the corresponding bandwidth at the network layer.

In the following, two policies for the proxy are introduced and discussed in detail. Both have in common that the available bandwidth of the DSL link is in principle divided into video traffic and best-effort traffic. In the following these bandwidths are denoted as BW_v and BW_{be}. The allocation of bandwidth to these two classes of traffic could be delegated up to a certain degree to the subscriber of the DSL. E.g., one subscriber might allocate 4 Mbps to video traffic and 2 Mbps to best-effort traffic, depending on his/her usage preferences. Of course, the system should not allow the subscriber to define allocations that might not make sense (e.g., assigning only 0.1 Mbps to video traffic). This basic allocation can be interpreted in the HTB terminology as assigning assured rates to both types of traffic.

A. Hard Reservation Policy

The Hard Reservation Policy (HRP) can be categorized as a conservative one. As already discussed, the capacity of the link is divided into video and best-effort traffic. According to the HTB approach, two traffic classes with corresponding assured rates are created at the network layer. The management of the network layer which includes the creation, deletion, and modification of traffic classes is performed by the proxy. Each attempt of initiating a VoD session via RTSP is intercepted by the proxy. The proxy examines the session description of the scalable video content in order to analyze the possible adaptations and their impact on the bit rate. Based on that scalability information, the knowledge about the bandwidth allocated for video traffic and the bandwidth requirement of the sessions already served by the proxy, the proxy performs admission control (AC).

The admission control algorithm works as follows. If the bandwidth requirement for the requested session is less than the remaining video bandwidth then the requested video stream can be delivered to the client without any kind of adaptation. The remaining video bandwidth is calculated as the bandwidth assigned to the video traffic BW_v minus the bandwidth allocated to each session currently served by the proxy. The proxy allocates the required bandwidth for the session and serves the requested video. At the network layer, an HTB traffic class with the corresponding assured rate and appropriate filters are created by the proxy. If the video cannot be served in its highest quality version due to insufficient bandwidth, then the proxy checks if an adapted version of the video can be served. By examining the scalability information, the proxy searches for adaptation parameters that cause the adapted video to meet the bandwidth requirements. If such a parameter combination exists, the proxy initiates the on-the-fly adaptation of the video and allocates a corresponding traffic class including filters at the network layer. If it turns out that no such parameter combination exists, the proxy refrains from serving the video and returns the client the appropriate RTSP error message. In that case, no interaction with the traffic control at the network layer takes place. Once the proxy is serving a video stream, the bandwidth allocated for that stream is assured and policing takes place if the stream does not comply to its traffic specification. The traffic classes and filters are maintained as long as the video is streamed and are removed in case of an explicit teardown or in case of a session timeout.

The advantage of this policy is that once the proxy has admitted a session and allocated the bandwidth, the video will be streamed with the guaranteed bit rate until the end of the session. Requesting a new session does not affect the sessions that are already served by the proxy. By using this policy, uni-directional bandwidth borrowing from video traffic to best-effort traffic can take place. This means that the subscriber of the DSL link could make use of the full capacity of the link for downloading purposes as long as there are no VoD sessions.

The following example illustrates the operation of this policy. The capacity of the DSL link is assumed to be 6 Mbps, of which the subscriber assigns 4 Mbps for video traffic and 2 Mbps for best-effort traffic. Additionally, we assume that a client within the DSL subscriber's home network requests video content which is encoded with one base and four enhancement layers leading to bandwidth requirements of 1200, 1400, 1700, 2100 and 2500 kbps. As illustrated in Figure 3 the bandwidth allocation is enforced by the proxy during its initialization phase at t_0 by setting up two traffic classes with appropriate assured and ceil rates. The rate assignment allows for borrowing unused video bandwidth to the best-effort class but not vice versa. Let's assume that there is a request for a VoD session at time t_1. According to the admission control algorithm the proxy would decide to serve the video stream without any adaptation since 4 Mbps of bandwidth are available to serve the full quality (2.5 Mbps). Since the unadapted version of the video stream (base layer incl. all enhancement layers) can be served with the allocated bandwidth of 4 Mbps, the proxy creates an appropriate traffic class (Video Stream 1). At a later point in time, t_2, a second request for a VoD session is issued by another set-top box of the DSL subscriber. For the sake of simplicity we assume the same layer configuration and bandwidth requirements for this second stream. The proxy performs admission control again and analyzes the scalability information exchanged in the session description. Since the remaining video bandwidth is 1500 kbps, the second session cannot be served in its initial version. Adaption becomes necessary. The proxy determines the adaptation parameters that are necessary to achieve a video stream that can be served with the available bandwidth. In our example, it decides to serve the base layer in combination with one enhancement layer and the setup of the appropriate traffic class at the network layer takes place. A request for a third VoD session would be rejected by the proxy, since it is serving two sessions with a total of 3900 kbps so the remaining 100 kbps do not allow to serve even the base layer of a third stream.

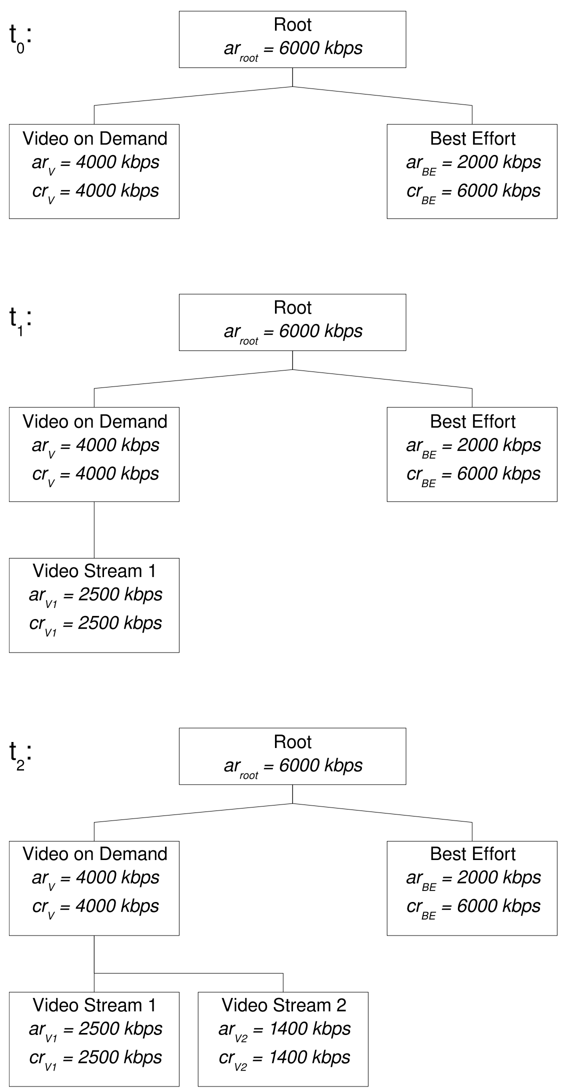

Fig. 3. Example - Hard Reservation Policy

B. Flexible Borrowing Policy

A second policy, which we call Flexible Borrowing Policy (FBP), that can be enforced by the proxy is based on the Medium Grained Scalability (MGS) feature [10] of H.264/SVC. As with the first policy, the subscriber of the DSL link has the possibility to allocate bandwidth for both video and best-effort traffic (BW_v and BW_{be}) according to his/her preferences. The main difference is that AC is now performed only for the base layer. Since the base layer cannot be further adapted to a achieve a lower bit rate of the video, the AC outcome can only be to either admit the base layer or not. If the base layer cannot be served with the remaining video bandwidth, the proxy denies the session setup. Otherwise, the base layer is admitted and a corresponding HTB traffic class and the appropriate filters are installed at the network layer.

Packets containing the enhancement layers are handled in a different way. The proxy maintains a single HTB traffic class (shared by all VoD sessions) that serves the enhancement layer packets. This traffic class employs a priority-based scheduling discipline with a configurable number of priority classes. For our first evaluations we used a fixed configuration of three priorities. The packets carrying the enhancement layer are marked by the proxy according to their importance. The proxy makes use of the scalability information contained in the session description in order to derive a mapping between the quality layer id (QID) and the resulting priority class on the network layer. The mapping is linear and considers the bit rate information of each layer. The marking of the packets by the proxy is realized by using the TOS/DSCP field in the IP header.

The assured rate and ceil rate of the traffic class for the enhancement layers are configured by the proxy in a very dynamic way. The values are updated each time a new VoD session is admitted or in case of a teardown of a session. The assured rate ar_{el} is calculated by subtracting the bandwidths allocated for the base layers of all VoD sessions from the amount of bandwidth allocated to video traffic (BW_v). The ceil rate cr_{el} of the enhancement layer traffic class is the sum of the assured rate and the bandwidth assigned to the best-effort traffic: $cr_{el} = ar_{el} + BW_{be}$. This means that unused best-effort bandwidth can be borrowed to the traffic class that serves the enhancement layers. In the absence of best-effort traffic, the unused bandwidth can be used to transmit more enhancement layer packets which results in a higher quality of all VoD sessions. If there is best-effort traffic, there is no bandwidth borrowing and the enhancement layers of all VoD sessions compete for the bandwidth of the enhancement layer class (ar_{el}). If the bandwidth requirement of all enhancement layer packets exceeds that bandwidth (ar_{el}), the priority-based queuing discipline will begin to drop packets. According to the priorities assigned by the proxy, the less important packets will be dropped first, resulting in a degradation of all video streams.

The behaviour of FBP applied to the example introduced above is illustrated in Figure 4. As with HRP, admission control is performed each time a VoD session is requested.

Therefore, at t_1 AC is performed for the base layer only, which requires 1200 kbps as mentioned above. The remaining video bandwidth of 2800 kbps is assigned to the traffic class for the enhancement layers (ELs). Additionally, the traffic class allows bandwidth borrowing of additional 2000 kbps, which is the bandwidth allocated to best-effort traffic. Following the layer configuration and the three configured priority classes, the priority mapping for the EL packets is derived as follows. Since the scalable video stream can be adapted between 1200 and 2500 kbps, this range is divided into three intervals $I_1 = [1200 \ldots 1633]$, $I_2 = [1634 \ldots 2066]$, and $I_3 = [2067 \ldots 2500]$. Based on the bit rate information of each layer, the ELs are mapped to these intervals and corresponding priorities. In this example, the first EL will get the highest priority, the second and third EL are getting both medium priority and the fourth EL will be handled with the lowest priority. During the second request at t_2 the same AC and priority assignment takes place. Additionally, the assured and ceil rates of the EL traffic class are updated according to the remaining video bandwidth. As one can see from the figures, a request for a third VoD session would also be possible, since the remaining bandwidth would still allow to admit a third base layer.

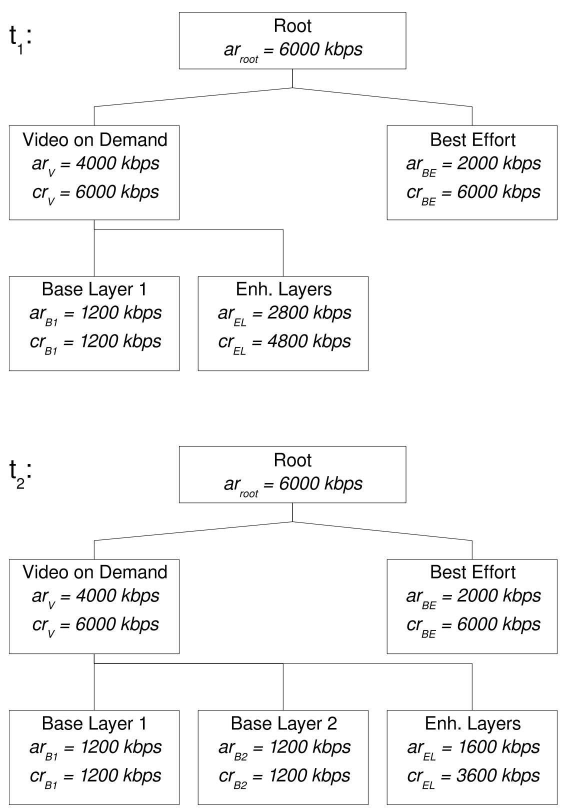

Fig. 4. Example - Flexible Borrowing Policy

C. Policy Comparison

When comparing both policies, it cannot be simply stated that one policy is better than the other. Rather, their usage depends on the preferences of the IPTV provider. HRP ensures that once a VoD session is admitted, it is served in the same quality until the end of the session. On the other hand, this prohibits making use of unused bandwidth, so that the quality cannot increase during the course of the session. AC based on the base layer – as it is applied in the FBP – also leads to a higher number of concurrent sessions as compared to HRP. This is due to the fact that HRP always tries to allocate as much bandwidth as possible for the requested VoD session. As a result the bandwidth assigned to each session by the HRP policy depends on the time of request. Such an effect cannot be observed when using the FBP since admission control is solely based on the base layer and the enhancement layers of all VoD sessions compete within a single traffic class.

Layer	harbour	football
Base layer	551 kbps	665 kbps
QID=1	1170 kbps	1010 kbps
QID=2	1670 kbps	1575 kbps
QID=3	2219 kbps	1879 kbps
QID=4	2814 kbps	2754 kbps

TABLE I
LAYER CONFIGURATION

Another difference between both policies is the direction of bandwidth borrowing. When applying HRP, the borrowing is uni-directional. Unused bandwidth that is allocated for video streams can be borrowed for serving best-effort traffic, but not vice versa. In contrast to that, the FBP allows for bi-directional bandwidth borrowing. Here, unused best-effort bandwidth can also be borrowed to the traffic class that serves the enhancement layer packets of the VoD sessions. This can be considered as a minor advantage of FBP in terms of efficient usage of the DSL link capacity.

IV. Evaluation

In order to validate our approach, we performed an experimental evaluation of the proxy and its implications on the bandwidth utilization. The experimental setup that represents the IPTV deployment as discussed in Section II was as follows. The scalable video content was streamed by an enhanced version of Apple's Darwin Streaming Server (DSS)[1] which was extended for support of H.264/SVC. The VoD session was initiated by the usual RTSP signaling. The media data itself was transmitted using RTP according to [3]. The video streams are requested by the client by using the openRTSP tool, which is an RTSP/RTP client that ships with the Live555 multimedia library[2]. It supports to dump the received video stream into a file which can then be used during an off-line decoding and analysis process. The client and the server are communicating through the proxy. Client, server and proxy software were deployed on dedicated Linux-based PCs and were connected by 100 Mbps Fast Ethernet links. The emulation of the DSL link between the proxy and the client was accomplished by

[1]http://developer.apple.com/opensource/server/streaming
[2]http://www.live555.com/liveMedia

Scenario	harbour		football	
	stream 1 [dB]	stream 2 [dB]	stream 1 [dB]	stream 2 [dB]
HRP	37.71	34.92	34.49	29.86
FBP w/o BE	37.71	37.71	34.49	34.49
FBP w BE	36.36	36.31	33.46	33.53

TABLE II
EVALUATION RESULTS

limiting the bandwidth of the link to 6 Mbps by using the network emulation plugin [11]. For the generation of best-effort traffic over the DSL link we used the iperf tool which uses a TCP connection to measure the available bandwidth on a link. Since the TCP connection continuously tries to maximize the utilization of the link we consider this as the worst-case situation. For the evaluation we encoded two different video clips (harbour and football) with the JSVM reference software (version 9.12.2). The clips were encoded at a resolution of 352x288 (CIF) and four MGS enhancement layers. The encoding parameters were selected to allow a bit rate adaptation of the video content in the range of about 500 to 2800 kbps. An overview of the layer configuration is given in Table I.

Both policies were evaluated using the same procedure. At the beginning of the experiment, a VoD session was requested by the client (stream 1). After the first stream was streamed for one second, another session was requested by the same client (stream 2). Both video streams were streamed for 100 seconds and were dumped into a file at the client. For both policies the allocated bandwidth for the video and best-effort traffic was 4 and 2 Mbps, respectively. The video streams were decoded off-line and the average resulting PSNR value was calculated on a per-frame basis. While HRP was evaluated in the absence of best-effort traffic, the Flexible Borrowing Policy was evaluated both with and without best-effort traffic. The reason for these settings is that the FBP can make use of unused best-effort bandwidth which results in a higher quality. The resulting PSNR values can be found in Table II. As can be seen from the performance figures of HRP, the quality of the first and the second stream significantly differs. While the first session is served at the highest quality (no adaptation), the second stream is adapted by the proxy to meet the 4 Mbps restriction imposed by the proxy. In the case of the FBP the quality of both sessions is quite the same. If there is no best-effort traffic, the unused bandwidth of 2 Mbps is borrowed to the video streams which allows to serve both sessions without any adaptation. In the case of best-effort traffic no borrowing is performed which results in insufficient bandwidth to deliver all enhancement layers of both sessions. Therefore, the priority-based queuing discipline that handles the enhancement layer packets will drop the least important packets. As a consequence the quality of both sessions decreases.

V. CONCLUSIONS

In this paper we presented a novel approach that combines both in-network adaptation and traffic control for scalable video streams based on the H.264/SVC standard. An intercepting RTSP/RTP proxy performs admission control and decides if the requested video content can be streamed in its original form or if it has to be adapted to meet the bandwidth restrictions of the network link. Depending on the outcome of the admission control process the proxy configures the network layer appropriately in order to separate the video streams from best-effort traffic. In the IPTV scenario which we considered, a subscriber is given the flexibility to decide how the capacity of the DSL link is assigned to both video and best-effort traffic. In contrast to fixed bandwidth allocation policies, our approach makes use of the Hierarchical Token Bucket (HTB) queuing discipline which allows for borrowing bandwidth between traffic classes. This maximizes the utilization of the network link. Two different policies that can be enforced by the proxy are introduced and discussed in the paper. The Hard Reservation Policy (HRP) performs admission control on the video streams and does not influence video streams once they are admitted. In contrast to that, the Flexible Borrowing Policy (FBP) restricts the admission control to the base layer of the scalable video stream. The packets carrying MGS enhancement layers of the SVC video are marked with priorities by the proxy and are handled at the network layer by a priority-based queuing mechanism. The advantages and disadvantages of both policies are subject to both a qualitative discussion and a quantitative evaluation of our prototype in an experimental test-bed.

ACKNOWLEDGEMENT

This work was supported by the Austrian Science Fund (FWF) under project "Adaptive Streaming of Secure Scalable Wavelet-based Video (P19159)" and by the EC in the context of the P2P-Next project (FP7-ICT-216217).

REFERENCES

[1] T. Wiegand, G. Sullivan, H. Schwarz, and M. Wien, eds., *ISO/IEC 14496-10:2005/Amd3: Scalable Video Coding*. International Standardization Organization, 2007.

[2] I. Kofler, M. Prangl, R. Kuschnig, and H. Hellwagner, "An H.264/SVC-based adaptation proxy on a WiFi router," in *Proc. NOSSDAV 2008*, pp. 63–68, May 2008.

[3] S. Wenger, Y.-K. Wang, T. Schierl, and A. Eleftheriadis, "RTP Payload Format for SVC Video." Internet Draft, Nov. 2008.

[4] M. Wien et. al., "Real-Time System for Adaptive Video Streaming Based on SVC," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 17, pp. 1227–1237, September 2007.

[5] J. Brandt and L. Wolf, "Adaptive Video Streaming for Mobile Clients," in *Proc. NOSSDAV 2008*, pp. 113–114, May 2008.

[6] Y.-K. Wang, M. M. Hannuksela, S. Pateux, A. Eleftheriadis, and S. Wenger, "System and Transport Interface of SVC," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 17, pp. 1149–1163, September 2007.

[7] S. Han et. al., "IPTV Transport Architecture: Alternatives and Economic Considerations," *IEEE Communications Magazine*, vol. 46, pp. 70–77, February 2008.

[8] M. Dischinger, A. Haeberlen, K. Gummadi, and S. Saroiu, "Characterizing Residential Broadband Networks," in *Proc. of IMC07*, (New York, NY, USA), ACM Press, 2007.

[9] M. Devera, "HTB Homepage." URL: http://luxik.cdi.cz/~devik/qos/htb/.

[10] H. Schwarz et. al., "Overview of the Scalable Video Coding Extension of the H.264/AVC Standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 17, pp. 1103–1120, September 2007.

[11] S. Hemminger, "Network Emulation with NetEm," in *Proceedings of the linux.conf.au 2005*, April 2005.