

A Quality Adaptation Scheme for Internet Video Streams

Panagiotis Papadimitriou and Vassilis Tsaoussidis

Demokritos University, Electrical & Computer Engineering Department,
12 Vas. Sofias Street, Xanthi, 67100, Greece
{ppapadim, vtsaousi}@ee.duth.gr

Abstract. We propose a layered quality adaptation scheme for video streams to smooth the short-term oscillations induced by *Additive Increase Multiplicative Decrease (AIMD)* mechanisms, and eventually refine the perceptual video quality. The layered scheme utilizes receiver buffering, adapting the video quality along with long-term variations in the available bandwidth. The allocation of a new layer is based on explicit criteria that consider the available bandwidth, as well as the amount of buffering at the receiver. Consequently, the adaptation mechanism prevents wasteful layer changes that have an adverse effect on user-perceived quality. In the sequel, we concentrate on the interactions of the layered approach with *Scalable Streaming Video Protocol (SSVP)*. Exploiting performance measures related to the perceived quality of rate-adaptive video streams, we quantify the combination of SSVP rate control and receiver-buffered layered adaptation.

1 Introduction

An increasing demand for multimedia data delivery coupled with reliance on best-effort networks, such as the Internet, has spurred interest in rate-adaptive multimedia streams. Video streaming, in particular, is comparatively intolerant to delay and variations of throughput and delay. Unlike bulk-data transfers, video delivery requires a minimum and continuous bandwidth guarantee. Rate adaptive video streams offer the clients the benefit of being resilient to changing network conditions and allow for a large number of streams to concurrently share network resources. Video streams can be adaptive, since user-perceived *Quality of Service (QoS)* is often satisfactory over a range of stream compression levels. Although this adaptivity is limited (i.e. multimedia streams have minimum subscription levels, below which service quality is unacceptable), they have the capability of adjusting their subscription levels in response to congestion, much as elastic flows do.

Today's Internet is governed by the rules of *Additive Increase Multiplicative Decrease (AIMD)* [2], which effectively contribute to its stability. Essentially, the goal of such algorithms is to prevent applications from either overloading or under-utilizing the available network resources. Although *Transmission Control Protocol (TCP)* provides reliable and efficient services for bulk-data transfers, several design issues render the protocol a less attractive solution for multimedia applications. More precisely, the process of probing for bandwidth and reacting to the observed congestion causes oscillations to the achievable transmission rate. Furthermore, TCP occasionally introduces arbitrary delays, since it enforces reliability and in-order delivery. In response to

standard TCP's limitations, several TCP protocol extensions [1, 4, 12] have emerged, providing more effective bandwidth utilization and sophisticated mechanisms for congestion control. *TCP-friendly* protocols [4, 11, 12] achieve smoother window adjustments, while they manage to compete fairly with TCP flows. In order to achieve smoothness, they use gentle backward adjustments upon congestion. However, they compromise responsiveness through moderated upward adjustments [9].

Considering TCP's limitations and the impending threat of unresponsive UDP traffic, rate-based congestion control has become an attractive alternative [4, 6, 8]. Avoiding the burstiness occasionally induced by the window-based mechanisms, rate-based protocols generate a smooth data flow by spreading the data transmission across a time interval. Therefore, rate-based mechanisms compose plausible candidates for media-streaming applications.

In addition to congestion control, a streaming video server should be able to control the quality of the video stream depending on the prevailing network conditions. *Simulcast* and *layered adaptation* compose two widely-used quality adaptation techniques. Simulcast uses multiple versions of the stream, encoded at different bitrates. The server transmits all the alternate streams and the client switches to the stream version that best matches its capacity. Layered adaptation has been proposed as a solution to the bandwidth redundancy introduced by simulcast. This approach is based on information decomposition. That is, the video stream is encoded at a *base layer* and one or more *enhancement layers*, which can be combined to render the stream at high quality. Layered adaptation is performed by adding or dropping enhancement layers depending on current conditions. In general, a large number of video layers results in greater bandwidth resolution. However, depending on the coding method, there is a trade-off between the number of layers and the video coding efficiency.

We are particularly interested in the interactions between layered adaptation and AIMD-based congestion control. In this context, we propose a receiver-buffered layered scheme to adapt video quality along with long-term variations in the available bandwidth. Receiver buffering reduces jitter, and depending on the amount of buffered data, the receiver is enabled to sustain temporary drops in the sending rate. In order to prevent wasteful layer changes which impair the perceived video quality, the proposed mechanism allocates additional layers based on certain criteria that consider the available bandwidth, as well the amount of buffering at the receiver. The layered scheme is designed to effectively interact with AIMD mechanisms: layered encoding allows video quality adjustments over long periods of time, whereas AIMD congestion control adjusts the transmission rate rapidly over short-time intervals. In particular, we study the interactions of the layered approach with *Scalable Streaming Video Protocol (SSVP)* [6], adjusting the quality of congestion-controlled video on-the-fly. SSVP is an AIMD-oriented rate control scheme optimized for video streaming applications. Quantifying the interactions of SSVP with the specific adaptation scheme, we identify that layered adaptation in conjunction with receiver buffering smoothes the short-term oscillations induced by AIMD mechanisms, and eventually refines the perceptual video quality.

The remainder of the paper is organized as follows. The following section summarizes related work. In Section 3, we analyze the parameters of the proposed layered scheme and study its interactions with SSVP. Section 4 includes conclusive performance studies based on simulations. Finally, Section 5 concludes the paper.

2 Related Work

Numerous studies for adaptive video delivery appear in [3, 5, 7, 10]. An overview of existing solutions for video adaptation is presented in [5]. Authors in [3] analyze the impact of selected congestion control algorithms on the performance of streaming video delivery. They concentrate on binomial congestion control [1] and especially on SQRT, which responds to packet drops by reducing the congestion window size proportional to the square root of its value instead of halving it. [7] proposes a layered mechanism to adapt the quality of congestion-controlled video. The mechanism is able to control the level of smoothing in order to improve the quality of the delivered video stream. Authors in [10] provide a system-level analysis of performance and design issues surrounding rate adaptive networks.

The literature includes numerous studies and proposals towards efficient rate/congestion control for multimedia applications in the Internet. *Rate Adaptation Protocol (RAP)* [8] is a rate-based protocol which employs an AIMD algorithm for the transmission of real-time streams. The sending rate is continuously adjusted by RAP in a TCP-friendly fashion, using feedback from the receiver. However, since RAP employs TCP's congestion control parameters (i.e. 1, 0.5), it causes short-term rate oscillations, primarily due to the multiplicative decrease. Furthermore, RAP occasionally does not result in inter-protocol fairness.

TFRC [4] is a representative equation-based protocol, which adjusts its transmission rate in response to the level of congestion, as estimated based on the calculated loss rate. Multiple packet drops in the same *Round Trip Time (RTT)* are considered as a single loss event by TFRC and hence, the protocol follows a more gentle congestion control strategy. More precisely, the TFRC sender uses the following response function:

$$T(p, RTT, RTO) = \frac{1}{RTT \sqrt{\frac{2p}{3}} + RTO (3 \sqrt{\frac{3p}{8}}) p (1 + 32p^2)} \quad (1)$$

where p is the steady-state loss event rate and RTO is the retransmission timeout value. Equation (1) enforces an upper bound on the sending rate T . However, the throughput model is quite sensitive to parameters (e.g. p , RTT), which are often difficult to measure efficiently and to predict accurately. Also, the long-term TCP throughput equation does not capture the transit and short-lived TCP behaviours, and it is less responsive to short-term network and session dynamics. TFRC eventually achieves the smoothing of the transmission gaps and is suitable for applications requiring a smooth sending rate. However, this smoothness has a negative impact, as the protocol becomes less responsive to bandwidth availability [9].

GAIMD [12] is a TCP-friendly protocol that generalizes AIMD congestion control by parameterizing the additive increase rate α and multiplicative decrease ratio β . For the family of AIMD protocols, authors in [12] derive a simple relationship between α and β in order to be friendly to standard TCP (i.e. $\alpha = 4(1 - \beta^2) / 3$). Based on experiments, they propose an adjustment of $\beta = 0.875$ as an appropriate smooth decrease ratio, and a moderated increase value $\alpha = 0.31$ to achieve TCP friendliness.

3 Quality Adaptation

The rationale of a quality adaptation scheme mainly rests on the assumption that a user's perception is sensitive to the changes in video quality, as well as to potential interruptions in the stream playback. Despite the degradation in visual quality, we consider smooth video of reduced bitrate more preferable than inconsistent and jerky video at highest quality.

3.1 Receiver-Buffered Layered Adaptation

We propose a quality adaptation mechanism in order to sustain smooth video delivery in a wide range of network dynamics. We specifically adopt the layered approach, where the streaming server coarsely adjusts video quality on-the-fly, without the need to implement transcoding. The server encodes raw video into n cumulative layers using a layered coder: layer 1 is the base layer and layer n is the least important enhancement layer. The layer rates are given by r_i , $i = 1, 2, \dots, n$ and l_j denotes the cumulative layer up to layer j , i.e. $l_j = \sum_{i=1}^j r_i$, $j = 1, 2, \dots, n$. The receiver obtains a certain number of layers depending on bandwidth availability. All active layers k ($1 \leq k \leq n$) are typically multiplexed in a single AIMD flow with rate $R = \sum_{i=1}^k r_i = k\bar{r}$, where \bar{r} denotes the average rate among the k layers.

Quality adaptation should not be simply supported by initial buffering at the receiver, because long-lived mismatch between the available bandwidth and the playback quality results in either buffer overflow or underflow. Alternatively, we rely on a receiver-buffered layered scheme, where each layer supplies a corresponding buffer, as depicted in Fig. 1. We assume that each buffer is drained with a constant rate c_i , $i = 1, 2, \dots, k$. Consequently, the total consumption rate at the receiver buffers is $C = \sum_{i=1}^k c_i = k\bar{c}$, where \bar{c} denotes the average consumption rate

of the buffers that correspond to the k active layers. The efficiency of a layered mechanism is significantly affected by the frequency of layer changes. Since we do not know in advance how long we will be able to sustain a specific layer, minimal rate variations should not directly trigger video quality adjustments. Generally, switching layers unduly usually induces perceptible video quality variations with frustrating consequences to the end-user. A simplified approach for the allocation of additional layers could be based exclusively on the instantaneous available bandwidth. Subsequently, the server would send video at the highest possible layer at any given time. In this case, the frequency and magnitude of oscillations in the congestion control algorithm would eventually govern the corresponding frequency and magnitude of layer changes. This would inevitably distract the end-user, considering that the Internet operates in the transient than in the stationary regime with occasional and rapid changes in the available bandwidth.

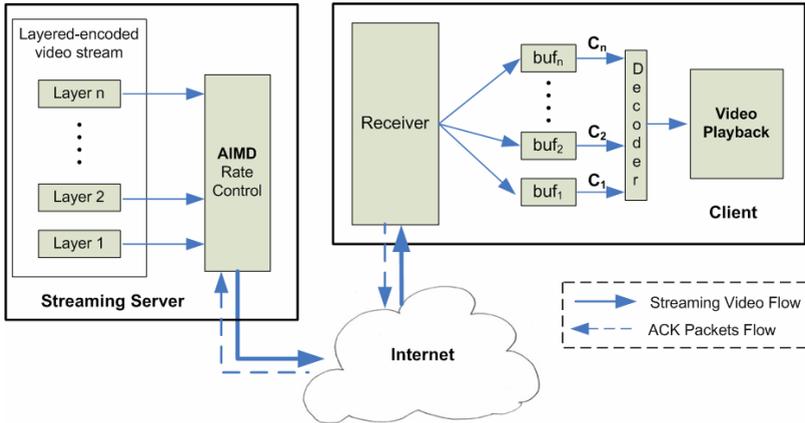


Fig. 1. Receiver-buffered layered adaptation

Following [7], certain conditions in the instantaneous available bandwidth and the amount of buffering at the receiver can be applied in order to reduce the number of layer changes. Initially, we concentrate on defining a priori whether a new layer should be allocated under the properties of AIMD rate control. Instead of delivering the maximum number of layers that can be accommodated by the network resources, we add a new layer, as soon as the current transmission rate R exceeds the total consumption rate of all currently active layers plus a new one:

$$R > (k + 1) \bar{c} \tag{2}$$

In this case, the new layer can be decoded immediately. However, relying on such rule does not eventually prevent layer changes; oscillations in the congestion control algorithm may still result in adding or dropping layers. Only sufficient buffering at the receiver can smooth out the variations in the available bandwidth and sustain a relatively constant number of active layers throughout the connection.

Along these lines, we attempt to derive a second rule associated with the amount of buffering required at the receiver. Fig. 2 depicts the behavior of an adaptive video

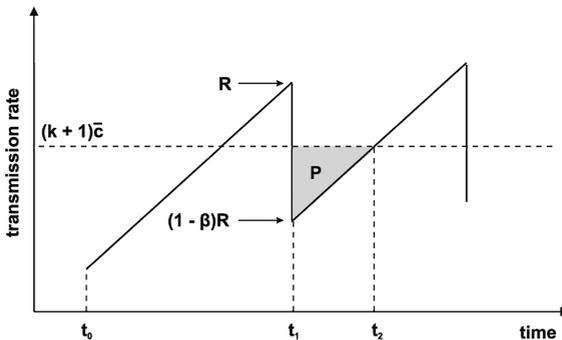


Fig. 2. Layered adaptation with AIMD rate control

flow under generalized AIMD (α , β) congestion control. After one backoff at time t_1 , we observe that at time t_2 , condition (2) holds; in addition, the buffering requirements are satisfied when:

$$\sum_{i=1}^k \text{buf}_i \geq P \quad (3)$$

where buf_i is the amount of video-data buffered for the i^{th} layer and P is the area of the shaded portion in Fig. 2. P represents the area of a triangle and can be derived by:

$$P = \frac{1}{2}(t_2 - t_1)[(k+1)\bar{c} - (1-\beta)R] \quad (4)$$

Assuming that the AIMD flow is in the congestion avoidance phase, the transmission rate R follows a line with slope α (Fig. 2). Subsequently, we obtain:

$$(k+1)\bar{c} = (1-\beta)R + \alpha(t_2 - t_1) \Rightarrow t_2 - t_1 = \frac{(k+1)\bar{c} - (1-\beta)R}{\alpha} \quad (5)$$

Combining equations (4) and (5), P is given by:

$$P = \frac{[(k+1)\bar{c} - (1-\beta)R]^2}{2\alpha} \quad (6)$$

Finally, we derive the second necessary and sufficient condition in order to allocate a new layer:

$$\sum_{i=1}^k \text{buf}_i \geq \frac{[(k+1)\bar{c} - (1-\beta)R]^2}{2\alpha} \quad (7)$$

Employing both rules for the allocation of additional layers enables the adaptation mechanism to trade short-term improvements for long-term quality smoothing, preventing buffer overflows and eventually rapid fluctuations in quality that frustrate the end-user.

3.2 Receiver-Buffered Layered Adaptation with SSVP Rate Control

We exploit the proposed receiver-buffered scheme to abolish the negative impact of short-term oscillations induced by AIMD mechanisms on video delivery. We hereby concentrate on the interactions between the layered approach and SSVP [6] rate control. SSVP, in a complementary role, operates on top of UDP, relying on sender and receiver interaction. The recipient uses control packets to send feedback of reception statistics to the sender. In accordance with the relaxed packet loss requirements of streaming video and considering the delays induced by retransmitted packets, SSVP does not integrate reliability into UDP datagrams. Hence, control packets do not trigger retransmissions. However, they are effectively used in order to determine bandwidth and RTT estimates, and properly adjust the rate of the transmitted video stream. The recipient uses packet drops or re-ordering as congestion indicator. SSVP employs an AIMD-oriented congestion control mechanism with $\alpha = 0.2$ and $\beta = 0.875$. The transmission rate is controlled by properly adjusting the inter-packet-gap (*IPG*). If no congestion is sensed, *IPG* is reduced additively; otherwise, it is increased multiplicatively. Further details in the operation of SSVP can be found in [6].

Considering SSVP's rate control parameters ($\alpha = 0.2$, $\beta = 0.875$), we rewrite equation (7), as:

$$\sum_{i=1}^k \text{buf}_i \geq \frac{5}{2} \left[(k+1) \bar{c} - \frac{R}{8} \right]^2 \quad (8)$$

SSVP enables a closed-loop control between server and client, where the receiver detects the state of congestion, determines the proper transmission rate and eventually opts for the number of layers that should be delivered. As a receiver-centric protocol, SSVP has first knowledge on the amount of buffering at the receiver and can augment the layered scheme to determine the optimal number of layers according to the prevailing conditions. Delegating on the client the selection of the number of layers to send composes a more tractable approach, since the sender may not be always able to accurately predict the amount of buffered data at the receiver. The allocation of a new layer is enabled only when both conditions (2) and (8) are satisfied. Feedback is provided to the server via the SSVP control packets.

3.3 QoS Assessment for Rate-Adaptive Video Streams

Consider that changes in stream subscription level have an adverse effect on user-perceived QoS, due to the distraction caused by changing video resolution. The term *subscription level* corresponds to the average rate of a stream for all possible encoding and compression levels for the given stream. In the case of a server utilizing layered encoding, a subscription level S infers that the client receives the base layer plus $S-1$ enhancement layers. We are interested in the assessment of dynamic rate adaptation, where a client changes its subscription level throughout stream playback in response to bandwidth availability along its route. Following [10], we use two QoS measures for rate-adaptive video streams: (i) the normalized average subscription level \bar{S} within $(0, 1]$, and (ii) the frequency of rate adaptation observed by an end-user, defined as:

$$F = \frac{1}{d} \sum |S(t^+) - S(t^-)|, \quad 0 \leq t \leq d \quad (9)$$

where $S(t)$ denotes the instantaneous subscription level at each time t . For example, $F = 1$ corresponds to one change in the stream subscription level per second on average. A combination of both QoS measures provides a useful insight to the user-perceived performance for rate-adaptive video streams.

4 Performance Evaluation

4.1 Experimental Environment

The evaluation plan was implemented on the *NS-2* network simulator. Simulations were conducted on a single-bottleneck *dumbbell* topology (Fig. 3) with a round-trip link delay of 64 ms. The bottleneck link is shared by competing MPEG and FTP connections and its capacity is configured depending on the experiment. The capacity of all access links to the source and sink nodes is set to 1 Mbps. The routers are drop-tail with buffer size adjusted in accordance with the *bandwidth-delay* product. We set the

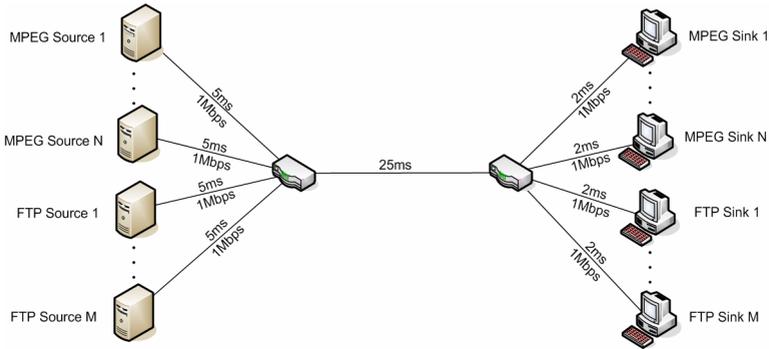


Fig. 3. Simulation topology

packet size to 1000 bytes for all system flows and the maximum congestion window to 64 KB for all TCP connections. The duration of each simulation is 60 sec.

In order to simulate MPEG traffic, we developed an *MPEG-4 Traffic Generator*. The traffic generated closely matches the statistical characteristics of an original MPEG-4 video trace. The model developed is based on *Transform Expand Sample (TES)*. We used three separate *TES* models for modeling I, P and B frames, respectively. The resulting MPEG-4 stream is generated by interleaving data obtained by the three models.

We hereby refer to the performance metrics supported by our simulation model. *Goodput* is used to measure the efficiency in bandwidth utilization. Following the metric in [11], we use *Coefficient of Variation (CoV)* in order to gauge the throughput smoothness experienced by flow i :

$$CoV_i = \frac{\sqrt{E_t\{\text{throughput}_i^2(t)\} - E_t\{\text{throughput}_i(t)\}^2}}{E_t\{\text{throughput}_i(t)\}}$$

where $E_t\{\}$ denotes the computation of the mean along time. Throughput rates are sampled every 150 ms. Long-term fairness is measured by the *Fairness Index*, derived from the formula given in [2], and defined as $(\sum_{i=1}^n \text{Throughput}_i)^2 / (n \sum_{i=1}^n \text{Throughput}_i^2)$, where Throughput_i is the throughput of the i^{th} flow and n is the total number of flows.

In order to quantify the performance on video delivery, we monitor packet inter-arrival times and eventually distinguish the packets that can be effectively used by the client application from delayed packets according to a configurable packet inter-arrival threshold. The proportion of delayed packets is denoted as *Delayed Packets Rate*. In accordance with video streaming requirements, we adjusted the packet inter-arrival threshold at 75 ms. Since MPEG traffic is sensitive to packet loss, we additionally define *Packet Loss Rate*, as the ratio of the number of lost packets over the number of packets sent by the application.

4.2 Results and Discussion

In the sequel, we demonstrate conclusive performance studies based on selected simulation results. Initially, we simulated one MPEG flow over SSVP (i) without adaptation, (ii) with *instantaneous* adaptation (SSVP-IA), and (iii) with the proposed receiver-buffered layered adaptation scheme (SSVP-LA). In the case of instantaneous adaptation, the server adapts the number of layers based on the instantaneous available bandwidth. The simulations were conducted on the dumbbell topology with bottleneck capacity of 1 Mbps. The system also includes two FTP flows over TCP Reno in order to enforce contention with the interfering MPEG flow. Fig. 4 illustrates the sending rate from the MPEG transfer in each case. On the occurrence of layered adaptation (i.e. Figs. 4b, 4c), the corresponding layer changes are also demonstrated. Table 1 includes useful statistics to assess the QoS perceived by the end-user.

According to Fig. 4c and the frequency of rate adaptation (i.e. $F = 0.11667$), the receiver-buffered scheme results in a minimal number of layer changes, adapting video quality along with long-term variations in the available bandwidth. Receiver buffering is able to sustain temporary drops in the sending rate and prevents wasteful dropping of layers. A new layer is allocated based on the available bandwidth, as well the amount of buffering at the receiver. The normalized average subscription level for the MPEG stream in the case of SSVP-LA is 0.75758. On the contrary, instantaneous adaptation causes frequent layer changes in response to the variations in the available bandwidth. The frequency of rate adaptation (i.e. $F = 1.05$) is significantly higher in comparison with the receiver-buffered approach, and has an adverse effect on user-perceiver quality. Despite a slightly higher $\bar{S} = 0.84413$ for the instantaneously adapted MPEG stream, the frequency of oscillations in the congestion control algorithm eventually governs the corresponding frequency of layer switching. Subsequently, the end-user is often distracted by the changes in video quality.

We are also interested in the interactions between the receiver-buffered adaptation scheme and the underlying AIMD protocol. We exploit CoV (Table 1) to quantify the magnitude of variation in the transmission rate. Note that a lower CoV expresses a lower variation in the sending rate, and consequently higher smoothness. Apparently, the receiver-buffered scheme reduces SSVP's short-term oscillations and the protocol delivers a smoothed video flow, which remains relatively immunized from the disturbances caused by the interfering FTP connections. The gains in smoothness (compared to SSVP) are less perceptible in the case of instantaneous adaptation; the measured CoV is still lower than SSVP alone, but notably higher than SSVP-LA. We point out that SSVP is a protocol that anticipates smoothness and delivers a smoothed flow, even without the support of a quality adaptation mechanism. As a result, the protocol does not leave much room for improvement in terms of smoothness. Therefore, we expect more significant gains, when the receiver-buffered mechanism augments other AIMD protocols that exhibit lower levels of smoothness.

We also carried out a series of simulations in order to assess the performance of receiver-buffered adaptation with SSVP-LA and multiple flows. In this context, we simulated a wide range of MPEG flows (1-50) over (i) SSVP-LA, (ii) TFRC, and (iii) GAIMD, competing with 5 FTP connections over TCP Reno, successively. GAIMD

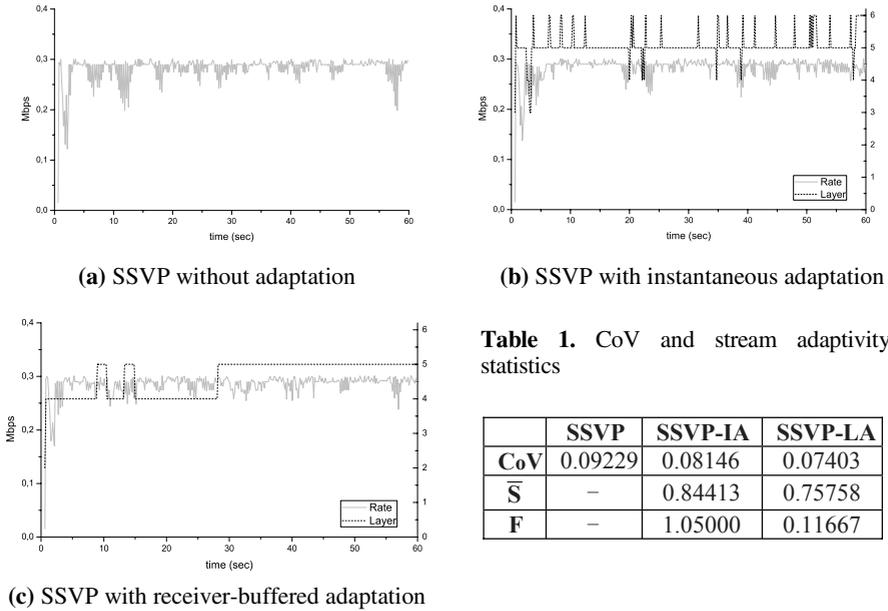


Table 1. CoV and stream adaptivity statistics

	SSVP	SSVP-IA	SSVP-LA
CoV	0.09229	0.08146	0.07403
\bar{S}	–	0.84413	0.75758
F	–	1.05000	0.11667

Fig. 4. Transmission rate and layer variation

and especially TFRC are implied to achieve remarkable efficiency on media delivery over a wide range of network and session dynamics. When the MPEG-4 server transmits over TFRC or GAIMD, the video is streamed at the optimal quality (i.e. layered adaptation is disabled). As a result, we evaluate the transmission rate control performed by each protocol at the transport layer. The experiments were conducted on the dumb-bell topology with a bottleneck capacity of 10 Mbps. Since today’s multimedia applications are expected to run in physically heterogeneous environments composed of both wired and wireless components, we included *NS-2* error models into the access links to the MPEG sink nodes. We used the Bernoulli model in order to simulate the link errors with packet error rate adjusted at 0.01. We measured *Goodput*, *Fairness Index*, and we additionally demonstrate statistics from delayed and lost packets, since both compose influencing factors for perceived video quality (Figs. 5-7).

According to Fig. 5, SSVP-LA flows utilize a high fraction of the available bandwidth, despite the delivery of lower bitrate video (the difference in rate/quality is subject to the prevailing network conditions) in comparison with TFRC and GAIMD. We observe that with high link-multiplexing, SSVP-LA outperforms the rest of the protocols in terms of bandwidth utilization, since the interaction of layered encoding with AIMD congestion control enables the effective adaptation of the video rate to the current network dynamics. GAIMD, in particular, yields inadequate bandwidth utilization, since it composes a *blind* window mechanism that relies on specific events triggered by violated thresholds.

Fig. 6 illustrates that SSVP-LA and GAIMD achieve high levels of fairness. The AIMD-based responses during congestion enforce competing flows to converge to the fairness point for both protocols. Furthermore, the receiver-buffered adaptation

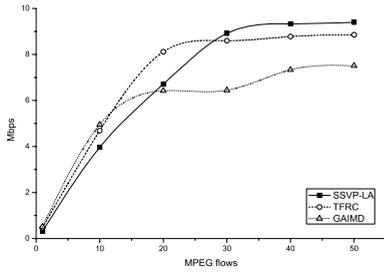


Fig. 5. Goodput of MPEG flows

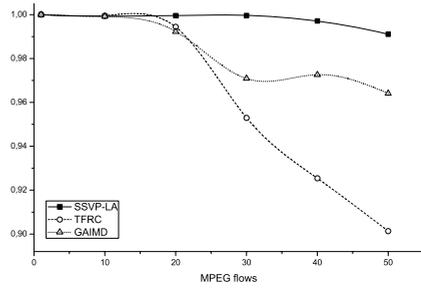
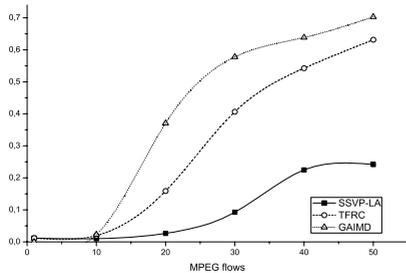
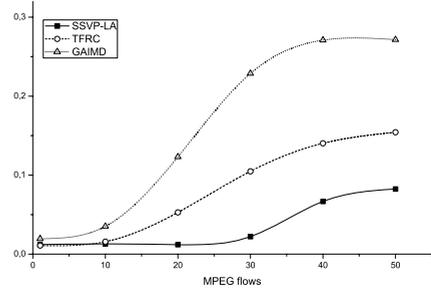


Fig. 6. Fairness Index



(a) Packet Loss Rate



(b) Delayed Packets Rate

Fig. 7. Performance on video delivery

mechanism confines packet loss (Fig. 7a) and therefore reduces the number of congestion cycles (i.e. the period between two consecutive loss indications), which results in fairness gains for SSVP-LA. On the other hand, we observe that the *Fairness Index* for TFRC degrades abruptly, reflecting a throughput imbalance between the connections. Apparently, TFRC's equation-based responses to packet loss undermine long-term fairness, along with contention increase.

According to Fig. 7, SSVP-LA achieves the timely delivery of most packets maintaining an uninterrupted and smooth video flow that is slightly affected by contention. The layered mechanism eventually bridges the gap between the short-term adjustments in the transmission rate caused by congestion control with the need for stable quality in streaming applications. Therefore, the flexibility of the receiver-buffered layered approach contributes in the delivery of smooth video in a wide range of network dynamics. SSVP-LA eventually enforces an upper bound to the magnitude of delay variation (Fig. 7b), providing a possible guarantee for streaming applications that can efficiently operate within this QoS provision. On the other hand, TFRC and GAIMD induce considerable variations in the receiving rate, with the effect of jitter becoming evident to the end-user. Furthermore, Fig. 7a depicts considerable packet loss for TFRC, which inevitably deteriorates the perceived video quality. The protocol occasionally fails to obtain accurate estimates of the loss event rate, invoking an inappropriate equation-based recovery, since TFRC's throughput model is sensitive to packet loss. We also observe a high packet loss rate for GAIMD, due to its gentle window decrease ratio at periods of congestion.

5 Conclusions

We have proposed a quality adaptation scheme that takes advantage of receiver buffering to alleviate most of the impairments caused by the oscillatory behavior of AIMD congestion control algorithms, improving the perceptual video quality. Our performance studies have validated the efficiency of the specific mechanism, uncovering significant performance gains, especially in the presence of limited bandwidth. Essentially, the proposed adaptation scheme provides a strong incentive for streaming applications to rely on AIMD rate control for efficient and smooth video delivery.

References

1. D. Bansal and H. Balakrishnan, Binomial Congestion Control Algorithms, in Proc. IEEE INFOCOM 2001, Anchorage, Alaska, USA, (2001)
2. D. Chiu and R. Jain, Analysis of the increase/decrease algorithms for congestion avoidance in computer networks, *Journal of Computer Networks and ISDN*, Vol. 17, (1989) 1-14
3. N. Feamster, D. Bansal, and H. Balakrishnan, On the Interactions Between Layered Quality Adaptation and Congestion Control for Streaming Video, in Proc. 11th Int/nal Packet Video Workshop, Korea, (2001)
4. S. Floyd, M. Handley, J. Padhye, and J. Widmer, Equation-Based Congestion Control for Unicast Applications, in Proc. ACM SIGCOMM 2000, Stockholm, Sweden, (2000)
5. J. Liu, B. Li and Y. Zhang, Adaptive Video Multicast over the Internet, *IEEE Multimedia*, Vol. 10, (2003) 22-33
6. P. Papadimitriou and V. Tsaoussidis, End-to-end Congestion Management for Real-Time Streaming Video over the Internet, in Proc. 49th IEEE GLOBECOM, San Francisco, USA, (2006)
7. R. Rejaie, M. Handley, and D. Estrin, "Layered Quality Adaptation for Internet Video Streaming", *IEEE Journal on Selected Areas in Communications (JSAC)*, Vol. 18, (2000) 2530-2544
8. R. Rejaie, M. Handley, and D. Estrin, RAP: An End-to-end Rate-based Congestion Control Mechanism for Realtime Streams in the Internet, in Proc. IEEE INFOCOM 1999, New York, USA, (1999)
9. V. Tsaoussidis and C. Zhang, The Dynamics of Responsiveness and Smoothness in Heterogeneous networks, *IEEE Journal on Selected Areas in Communications (JSAC)*, Vol. 23, (2005) 1178-1189
10. S. Weber and G. Veciana, Network Design for Rate Adaptive Media Streams, in Proc. IEEE INFOCOM 2003, San Francisco, USA, (2003)
11. Y. R. Yang, M. S. Kim, and S. S. Lam, Transient Behaviors of TCP-friendly Congestion Control Protocols, in Proc. IEEE INFOCOM 2001, Anchorage, Alaska, USA, (2001)
12. Y. R. Yang and S. S. Lam, General AIMD Congestion Control, in Proc. 8th IEEE Int/nal Conference on Network Protocols (ICNP), Osaka, Japan, (2000)