Smoothing Transcoded MPEG-1 Video Streams for Internet Transmission*

E. Jammeh, M. Fleury, M. Ghanbari
University of Essex
Electronic Systems Engineering Department
Audio and Video Networking Laboratory
Colchester, Essex, CO4 3SQ, United Kingdom
Tel.: +44 1206 872817
Fax.: +44 1206 872900
{eajamm,fleum,ghan}@essex.ac.uk

Abstract

This paper presents an efficient smoothing scheme for the real-time transmission of MPEG-1 transcoded video over 'best-effort' IP networks. The scheme uses intelligent partitioning and multiplexing of the packetised bitstream. Bit-rate smoothing is achieved by partitioning packets according to their picture type (I, P or B). Subsequently, the partitioned packets are multiplexed in such a way that each packet from an anchor (I and P) picture is followed by two packets from B-pictures. The proposed scheme smoothes the bit rate of the encoded video, making it more suitable for adaptive video streaming applications. In such applications, the transmission bit rate is varied so as to adapt to the available network bandwidth. The scheme re-organises the transmission order of the packets, spreading the less important packets from the B-pictures into the more important packets from the anchor pictures. This has the effect of reducing the likelihood

*A preliminary version of this paper was presented at IEE International Conference on Visual Information Engineering, VIE 2003

of losing the more important anchor packets, thereby improving the quality of transmitted video.

A variety of simulations results are presented to demonstrate these points.

Keywords: Video Streaming, Bit-rate smoothing, Transcoding, Quality Adaptation, and Video over IP.

MANUSCRIPT FOR CONSIDERATION FOR VIE2003 SPECIAL SECTION

1 Introduction

The growth of the Internet has encouraged the deployment of new multimedia applications, such as streaming of stored video. Most of these new multimedia applications require Quality of Service (QoS) guarantees from the underlying network. This requirement is largely due to assumptions made during the video encoding process, the availability of: sufficient channel bandwidth; low-loss delivery; and low delay, and variation of delay (jitter). At the network layer, the IP (Internet Protocol) is 'best-effort', meaning that packets may be lost or duplicated. The well-known TCP (Transmission Control Protocol), intended to impose reliability, has a 'bursty' transmission behaviour, in part caused by its Additive Increase Multiplicative Decrease (AIMD) congestion control mechanism. In addition, the mechanisms employed by TCP to achieve reliability such as Go-Back-n retransmission make TCP unsuitable for encoded video transport. Consequently, UDP (User Datagram Protocol) is used, which provides a rudimentary and unreliable service, requiring transcoder applications to re-invent control algorithms for improved video quality delivery. There are two varieties of control required: 1) congestion control and 2) smoothing and error reduction. Previous work by the authors [1; 2] has considered congestion control, while the subject of this paper is smoothing and error reduction of the encoded video stream. However, the two forms of control are related.

In previous work, a congestion control algorithm was developed that measured the level of network congestion, and computed a sending rate that adapted to the congestion level. A low-delay DCT-domain transcoder [3] adapted the encoded video variable bit rate to that computed by the congestion controller. However, the bit rate of transcoded video is inherently very bursty (refer to Section 2), inevitably making adaptive congestion control schemes inefficient [4].

For MPEG-1 transcoded video, the bit rate is picture dependent, with the anchor pictures (I and P) having a higher bit rate compared to the B pictures [5]. A B picture will consume approximately 20 Kb, if an I picture consume about 200 Kb [6]. Because of their higher bit rates, packets from the anchor pictures are more likely to be lost during periods of network congestion, compared to packets from the lower-rate B-pictures. Furthermore, lost packets from the anchor pictures will result in error propagation through a Group of Pictures (GOP), which is only reset by I pictures. A lost packet from an anchor picture will affect all subsequent pictures in the GOP. However, a lost packet from a B picture will result in errors confined within a single picture, as B pictures are not used for prediction. This leads to the situation where the more important packets from anchor pictures are more likely to be lost during periods of network congestion [7]. Adaptive video streaming congestion control schemes should therefore be complemented with bit-rate smoothing algorithms. Smoothing will reduce the transcoded video bit rate burstiness, enabling efficient rate adaptation with improved quality delivery. When smoothing is achieved with a reduction in the peak-to-mean bandwidth ratio, there is an added reduction in bit-rate burstiness.

This paper proposes smoothing the bit-rate through packet partitioning and multiplexing such that each packet from anchor pictures (I/P) is always followed by two packets from a B picture. The smoothing has the effect of a) reducing the peak-to-mean bit rate of the transcoded video, and b) reduce the depth and frequency of the bit rate variations of a typical transcoded bit stream. Besides, multiplexing packets from B-pictures with packets from the anchor pictures reduces the likelihood of dropping anchor packets, making the less important B packets readily available for dropping during periods of network congestion, as, if needed, an additional form of congestion control.

The remainder of this paper is organised as follows. Section 2 establishes some background knowledge. Section 3 explains the proposed smoothing scheme, and its performance is demonstrated in Section 4. Finally, following an examination of other work on smoothing schemes, Section 5, some conclusions are drawn in Section 6.

2 Background

The MPEG-1 standard [8] was designed to be compatible with a video cassette recorder (VCR) in terms of random access and high coding efficiency. Three main picture or frame types with different encoding methods were identified to accommodate a high compression efficiency with random access: 1) I (intra-coded), 2) P (predicted), and 3) B (bi-directional). The standard defines a Group of Pictures (GOP), which is a series of pictures that help to randomly access the video sequence. Each I picture in a GOP is followed by a combination of P and B pictures, Figure 1a), which is typical GOP structure. The transmission order is somewhat different, Figure 1b), to allow the reference pictures for the B pictures to first be decoded. The relative bit rates of pictures in an MPEG bit stream can be shown [5] to be 6:3:2, for respectively I, P, and B-pictures. Figure 2 shows the inherent 'burstiness' across time of a typical MPEG-1 bit rate. Not only is the bit rate 'bursty', but the peak-to-mean bit-rate ratio is high. The changing scene complexity adds to the fluctuations, and in scenes with considerable motion, the P- and B-picture bit-rate will increase, by as much as a factor of three.

Transcoding, in the context of this work, is the bit rate reduction of encoded video for encoded video bit rate adaptation. A convenient way to do so is to change the quantization level within the transcoder, though other options exist such as reducing the sensitivity to inter-frame motion. The bit-rate is adjusted on a picture-by-picture basis. Internet applications that can tolerate variable video quality include teleconferencing, and the low-cost distribution of information such as news or sport. Because the video frame structure, syntax, temporal and spatial resolution are all unchanged, the bit-rate burstiness of the encoded video is still inherent in transcoded video. Rate smoothing techniques enhance the performance of congestion control schemes [9], improving the QoS of the delivered video.

The purpose of congestion control is to adapt the transmission rate to the available network bandwidth. The desired transmission rate is a control on the transcoder's output rate. A fuzzy logic congestion controller was employed in the experiments in Section 4.1. Fuzzy logic emulates the control process, *as if* a human expert were regulating the transmission rate. The fuzzy membership functions model the uncertainty in that expert's perception of the feedback, whereas an output rate decision is made precise by the process of defuzzification. The feedback to the controller is the congestion level,

C_L and its rate of change δC_L, with C_L defined by:

$$C_L = 1 - \frac{G^s_{av}}{G^c_{av}}, \tag{1}$$

where G^c_{av} and G^s_{av} are moving averages of packet transfer times at the client and the server (transcoder) respectively, which have been normalized for packet size.

3 Smoothing

This section describes the proposed smoothing scheme. An ideal smoothing method is described, and its drawbacks highlighted. A more practical, improved smoothing scheme is described.

3.1 Ideal Transcoded Video Bit-Rate Smoothing

Consider an MPEG-1 transcoded video with pictures P_1, P_2, $\ldots P_N$, with a repetitive GOP structure, as shown in Figure 1. The aim of the smoothing scheme is to smooth the transmission bit rate of the bitstream such that each picture in a GOP transmits at an average ideal smoothed rate I_r. Where

$$I_r = \frac{\sum_{i=1}^{N} R_i}{N} = \frac{\sum_{i=1}^{N} S_i}{NT}. \tag{2}$$

R_i is the rate of the i_{th} picture in the GOP, S_i is the size of the i_{th} picture in the GOP, N is the number of pictures in a GOP, and T is the transmission duration of each picture. However, pictures have different sizes (number of bits) and are required to have the same transmission duration. This is the root cause of the bursty transmission bit rate shown in Figure 2. The difference in picture sizes is mainly caused by having different picture types, and having each picture type encoded differently to the other picture types. To enable pictures with different sizes to transmit at the same ideal smoothed rate, I_r, requires having pictures with unequal transmission durations, D_i — calculated from picture sizes and I_r, as given by (3).

$$D_i = \frac{S_i}{I_r}, \tag{3}$$

where D_i is the duration of the I_{th} picture and S_i is the size of the I_{th} picture. Assuming that each picture is made up of P_n slices[1], and that each slice is packed into a single packet for transmission, then the inter-packet gap (IPG) of transmitted packets in the i_{th} picture will be given by (4).

$$IPG_i = \frac{D_i}{P_n} = \frac{S_i}{P_n \times I_r} \tag{4}$$

The ideal smoothing of (4), has two main disadvantages. Firstly, the streamed video's temporal resolution might be different from that of the coded video, because the IPG is varied to achieve the smoothing effect. Secondly, varying the IPG will introduce packet inter-departure jitter, which translates into packet inter-arrival jitter at the destination, regardless of the underlying available network bandwidth.

3.2 Proposed Smoothing Scheme.

The idea of the proposed scheme is to average the bit rates of I-pictures with the rates of two B-pictures, to obtain an $I - B_r$ smoothed rate, which is the sending rate for the I and the two B pictures. A smoothed $P - B_r$ bit rate is also computed from each P-picture with its subsequent two B-pictures as described formally in (5) and (6).

$$I - B_r = \frac{I_r + B_r^1 + B_r^2}{T_n} \tag{5}$$

$$P - B_r = \frac{P_r + B_r^1 + B_r^2}{T_n}, \tag{6}$$

where I_r is the I-picture bit rate; Pr is the P-picture bit rate, and B_r^1 and B_r^2 are the bit rates of the first and second B-pictures following the anchor pictures. T_n is the number of picture types; this work assumes a value of 3 for T_n, i.e. $\{$I, P and B pictures$\}$.

It is assumed that: each picture is made up of eighteen slices, and that a single slice per packet (SSPP) packetisation scheme is used for transmission. Packets from the transcoder are partitioned and buffered according to their picture type. Three buffers, corresponding to the three picture types, are used to buffer the packets which are subsequently multiplexed before transmission. Two multiplexing schemes are used; the I-B and P-B methods. These generate the $I - B_r$, and the $P - B_r$ smoothed bit

[1] A slice is a group of frame macroblocks [5].

rates, for respectively the I-B and P-B multiplexing methods. The multiplexing method used depends on the current picture type being transcoded.

The I-B scheme multiplexes packets such that each packet from an I picture is always followed by two packets from a B picture, whilst two packets from a B picture always follow each packet from a P picture for the P-B scheme. This arrangement is illustrated in Figure 4.

The transmission order of Figure 15 requires a small adjustment to the basic scheme, due to the presence of only I and P packets in the initial part of the GOP transmission cycle. Rather than increase latency, waiting for generation of the B pictures, initially an I followed by two P packets are multiplexed onto the output channel. This arrangement continues until all of the P picture packets are exhausted, whereupon the I-B pattern in Figure 4 continues, and so on.

The proposed smoothing scheme has two main advantages apart from enhancing the performance of congestion control algorithms. Firstly, the probability of losing anchor packets (I/P) is reduced by a factor of three, as packets in anchor pictures are spread within three pictures. Secondly, the scheme reduces the peak-to-mean rate ratio, which will help to reduce the variation in network congestion level, resulting in a more stable delivered video QoS. Furthermore, the scheme will reduce the probability of loosing the more important anchor packets, and enhance the quality of the streamed video during periods of network bandwidth constraint. B-picture packets are also, if required, readily available for dropping during periods of network congestion, providing an additional form of congestion control.

4 Performance Evaluation

A video streaming application using the proposed scheme was simulated using the ns-2 network simulator [10; 11]. Two 1.75 Mbit/s encoded video traces were used for this experiment. The first trace was obtained from a raw and un-smoothed transcoded video, and the second trace was obtained from a transcoded video stream that implemented the smoothing scheme by partitioning and multiplexing the packets.

Two sets of experiments were carried out on simulated network, Figure 5, which includes a variable bandwidth link that can be made to form a bottleneck. In the first experiment, the smoothed video trace was sent from the server to the client through a bottleneck link. Using ns-2's built-in traffic generators, packets with a known inter-packet departure time probability density function (pdf) were injected from the source. A Pareto pdf is used [12] to model (interactive) Internet traffic flows. As the traffic is 'bursty', transmission only takes place during 'on' periods. This heavy-tailed pdf is represented by (7).

$$f(x) = ax^{(1-a)}, \tag{7}$$

where a =shape parameter, $a > 0$, $1 \leq x \leq \infty$. a near to one gives rise to self-similar traffic, whereas a near to two, has similar fractal properties to exponential traffic. In general terms, a heavy-tailed distribution can give rise to very large file size sizes, with a non-negligible probability, and such files, when packetized, might delay or cause loss of co-existent traffic. The Pareto configuration for the simulation is summarized in Table 6, with the maximum file size as 25 KB.[2]

The bottleneck link bandwidth was varied from 100 Kbit/s to 1 Mbit/s, in steps of 100 Kbit/s. The following performance metrics were used as a measure of performance for each streaming session: network congestion level, cumulative packet loss, and packet inter-arrival jitter. The experiment was repeated with a raw video bitstream instead of the smoothed bitstream. The results obtained are presented in the Sections 4.1–4.4.

4.1 Rate Burstiness Reduction

The main objectives of the smoothing scheme are: to reduce the peak-to-minimum rate ratio, reduce the depth of variation, and reduce the frequency of rate variations of the transcoded video bitstream. Figure 6 shows the resulting bit rate of a 1.75 Mbit/s transcoded video stream for which the smoothing scheme was enabled. It can be observed that the smoothing scheme achieved a reduction in the peak-to-minimum rate ratio from 9 to 2.5. This is evident when the raw unsmoothed transcoded bit stream, Figure 2, is compared to the smoothed version, Figure 6. Furthermore, the depth and frequency of rate variation has been reduced significantly, as demonstrated by Figure 6. The network performance

[2]Note that surveys, e.g. [13], conclude that, due to TCP protocol behavior and Ethernet frame size, the Internet packet size is trimodal at around 40, 560, and 1500 bytes. Thus, a setting of around 500 bytes tests a typical traffic scenario.

tool, `tcpflw` [11], was used to evaluate the frequency distribution of the packet throughput of both the raw and smoothed video stream, which are controlled by means of a Fuzzy Logic Congestion controller [2]. This reduction in rate variation is clearly shown when the frequency distribution of the rates of the smoothed and unsmoothed bit streams are compared for each bottleneck bandwidth. For example, the reduction in bit-rate burstiness is evident, when matching graphs of congestion controlled traffic, with and without smoothing are compared (Figures 7 and 8, Figures 9 and 10, for respectively a bottleneck rate of 400 Kbit/s and 800 Kbit/s)). In Figures 7–10, the comparison metric is the Instantaneous Packet Throughput (IPT), which is the Packet Size (PS) divided by the packet Inter-Arrival Time (IAT) between the packet and the following arriving packet, as defined in (8). The packet size includes the payload and header.

$$IPT_{n-1} = \frac{PS_{n-1}}{t_n - t_{n-1}}, \tag{8}$$

where t_n is the arrival time of the n^{th} packet.

4.2 Network Congestion Level

It has already been shown in Section 4.1 that the bit rate of the smoothed transcoded video has a reduced peak-to-minimum rate fluctuations. This will result in a reduced level of network congestion, due to the rate overwhelming the network capacity. Furthermore, the reduced depth and frequency of such rate variations will result in a more stable network congestion level. For example, Figures 11 and 12 show the performance of the smoothing scheme in terms of network congestion level for bottleneck bandwidths of respectively 400Kbits/s, and 800 Kbit/s. From the graphs, it is apparent that the network congestion level from the smoothed bit stream is inherently more stable in terms of reduced congestion, and a reduced frequency of variation.

4.3 Loss

It was shown in 4.2 that the smoothing scheme introduces a smoother and reduced network congestion level. This will result in a reduced packet loss rate. The smoothing scheme therefore, has a reduced packet loss rate when compared to an un-smoothed transcoded bit stream. Figures 13shows the performance of the smoothing scheme, before congestion control has been applied to either raw or smoothed streams, in terms of cumulative packet loss for a bottleneck bandwidth of 1200 Kbit/s. The improvement from using the smoothing scheme is about 6%. Combining smoothing with fuzzy congestion control shows a significant improvement in loss rates for the smoothed scheme, at tighter bottleneck constrictions, for example at 400 Kbit/s in Figure 14. Use of congestion control on the raw stream without smoothing, brings some reduction in the cumulative total, but it is the combination of smoothing and congestion that is most effective. It is evident from these graphs that combined with congestion control, the smoothing scheme gives an improved quality of delivered video, due to a reduced packet loss rate.

The frequency distribution of packet loss was also compared between raw and smoothed video streams, with Figure 15 being an example over a 5 s period. It was found that the loss distribution was more evenly distributed, with a tendency for packet losses to be bunched into clusters in the raw stream case. This implies that there is a greater probability of only anchor packets being lost during a cluster, whereas B picture packets are also vulnerable for a more even distribution.

4.4 Delay

Buffering packets at the source introduces some delay. The maximum delay introduced by the proposed smoothing scheme has an upper limit of three frames. A maximum delay of 120 ms is, therefore, introduced for encoded video with a temporal resolution of 25 frame/s. This is the delay incurred by packets from generation to just before transmission, and not the latency after sending the packets through the network.

5 Related work

This section is a brief look at some other work on on-line smoothing of encoded video. Unlike the transmission of stored video [14], on-line source-based smoothing assumes no knowledge of future pictures beyond the current GOP, as is appropriate to video-casting applications, and to a greater extent for interactive applications such as teleconferencing. In [15], it is pointed out that most stored video algorithms assume a constant end-to-end delay through a packet-switched network, whereas end-to-end delay jitter is normal. The work in [15], adopts the dynamic programming algorithm of [16] to these circumstances. A review [17] of earlier algorithms for the smoothing of real-time video especially that of [6], concludes that the number of rate changes in such algorithms is more than is necessary. In [18], an on-line smoothing scheme employs a piece-wise algorithm at a smoothing server, in which buffering is concentrated at the server, potentially reducing the client playback delay. The proposal herein is largely appropriate to real-time flows, being based on short-term variations at the GOP-scale. However, being packet-based, it is largely independent of the mathematically-based optimisation algorithms that have been employed elsewhere.

6 Conclusions

This paper presented the need for transcoded video smoothing to complement video congestion control. A smoothing scheme that partitions and multiplexes a packetised transcoded bit stream according to packet type, at the transport layer, was proposed. The scheme was shown to reduce the depth and frequency of rate variations of an MPEG-1 transcoded video bit stream, and reduced the peak-to-minimum rate ratio from 9 to 2.5. It was also shown that the reduced bit stream rate burstiness led to a reduction in the level of introduced network congestion. The scheme has the effect of improving the quality of transcoded video, as evidenced by a huge reduction in packet loss when combined with congestion control. Besides, the proposed scheme reduces the probability of losing the important anchor packets, thus, further improving the quality of the received video.

References

[1] E. Jammeh and M. Ghanbari. Transmission of pre-encoded video over best-effort IP networks. In *13$^{\underline{th}}$ International PacketVideo workshop (PV2003)*, 2003.

[2] E. Jammeh, M. Fleury, and M. Ghanbari. Fuzzy-logic control for transcoded video streams. *IEEE Signal Processing Letters*, 2003. Submitted in 2003.

[3] T. Shanableh and M. Ghanbari. Transcoding architectures for DCT-domain heterogenous video transcoding. In *International Conference on Image Processing, ICIP 2001*, pages 433 – 436, 2001.

[4] S. Wang, H. Zheng, and J. A. Copeland. Video multiplexing with QoS constraints. In *IEEE SPIE Conference on Internet Routing and Quality of Service*, pages 82–91, 1998. SPIE vol. 3529.

[5] M. Ghanbari. *Standard Codecs: Image Compression to Advanced Video Coding*. IEE, London, UK, 2003.

[6] S. S. Lam, S. Chow, and D. K. Y. Yau. An algorithm for lossless smoothing of MPEG video. In *ACM SIGCOMM'94*, pages 281–293, 1994.

[7] D. Wilson. *Efficient Coding and Transmission of Two-Layer Video*. PhD thesis, University of Essex, 1997.

[8] MPEG-1: Coding of moving pictures and associated audio for digital storage media at up to about 1.5 Mbit/s, 1991. ISO/IEC 1117-2: video.

[9] T. J. Shanableh. *Heteregenous Video Transcoding for Matching Network-Bandwidth and End-Systems Constraints*. PhD thesis, University of Essex, 2001.

[10] L. Breaslu, K. Estrin, D. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, X. Ya, and H. Yu. Advances in network simulation. *IEEE Computer*, 33(5):59–67, 2000.

[11] G. F. Lucio, M. Paredes-Farrera, E. Jammeh, M. Fleury, and M. J. Reed. OPNET modeler and Ns-2: Comparing the accuracy of network simulators for packet-level analysis using a network testbed. In *3rd WEAS International Conference on Simulation, Modelling and Optimization (ICOSMO 2003)*, 2003.

[12] K. Park, G. Kim, and M. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *4th International Conference on Network Protocols*, pages 171–180, 1996.

[13] K. Claffy, G. Miller, and K. Thompson. The nature of the beast: Recent traffic measurements from an Internet backbone. In *Internet Society INET'98*, 1998.

[14] W.-C. Feng and J. Rexford. Performance evaluation of smoothing algorithms for transmitting prerecorded variable-bit-rate video. *IEEE Transactions on Multimedia*, 1(3):302–313, 1999.

[15] C. Tryfonas. Practical considerations for smoothing multimedia traffic transported over packet-switched networks. In *12$^{\underline{th}}$ International PacketVideo workshop (PV2000)*, 2002.

[16] Z. Jiang and L. Kleinrock. A general optimal smoothing algorithm. In *IEEE INFOCOM '98*, volume 1, pages 676–684, 1998.

[17] J. Zhang and J. Hui. Optimal smoothness results and approximation techniques for real-time VBR video traffic smoothing. In *18$^{\underline{th}}$ IEEE Real-Time Systems Symposium¡ RTSS'97*, pages 253–263, 1997.

[18] S. Sen, J. Rexford, K. Dey, J. F. Kurose, and D. Towsley. Online smoothing of variable-bit-rate streaming video. *IEEE Transactions on Multimedia*, 2(1):37–48, 2000.

Packet size	500 bytes
Burst time	200 ms
Idle time	300 ms
Rate	1 Mbit/s
Shape	1.5

Table 1: Pareto traffic configuration

List of Figures

Figure 1: a) Typical GOP picture structure b) Picture transmission order.

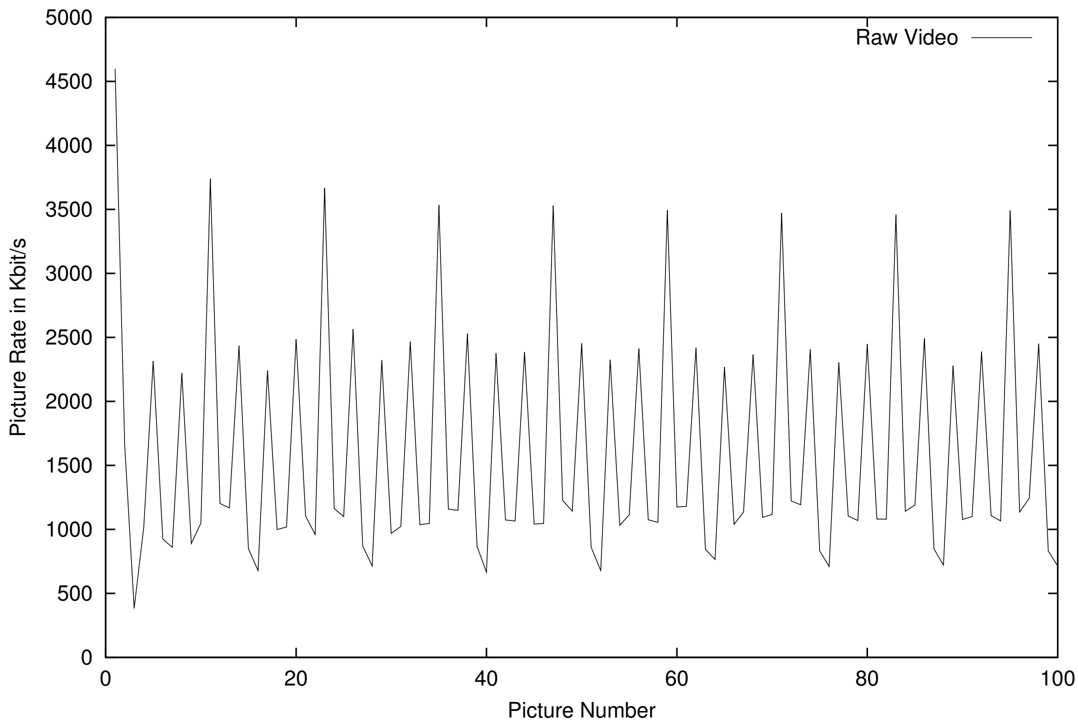

Figure 2: Bit rate of an MPEG-1 video coded at 1.75 Mbit/s, showing the bursty bit rate at picture level.

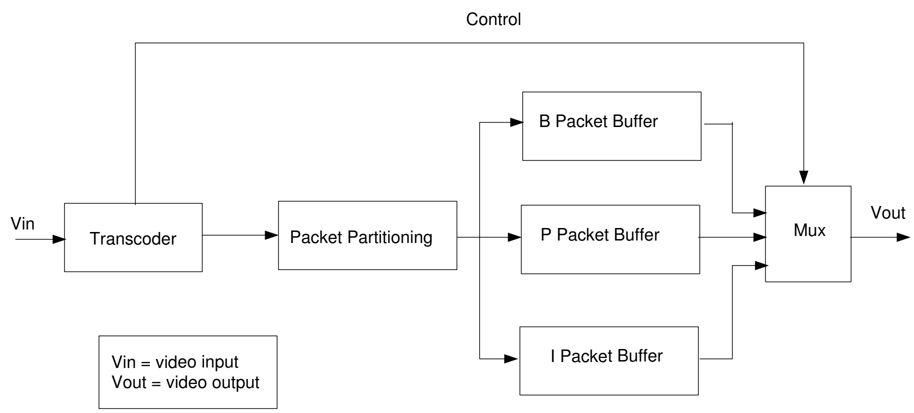

Figure 3: Smoothing through multiplexing architecture.

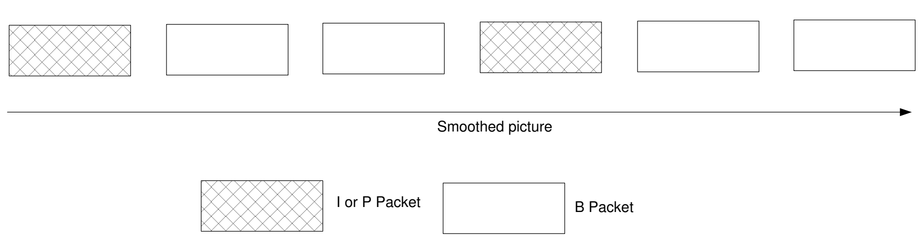

Figure 4: Packeting structure of smoothed I/P-B pictures.

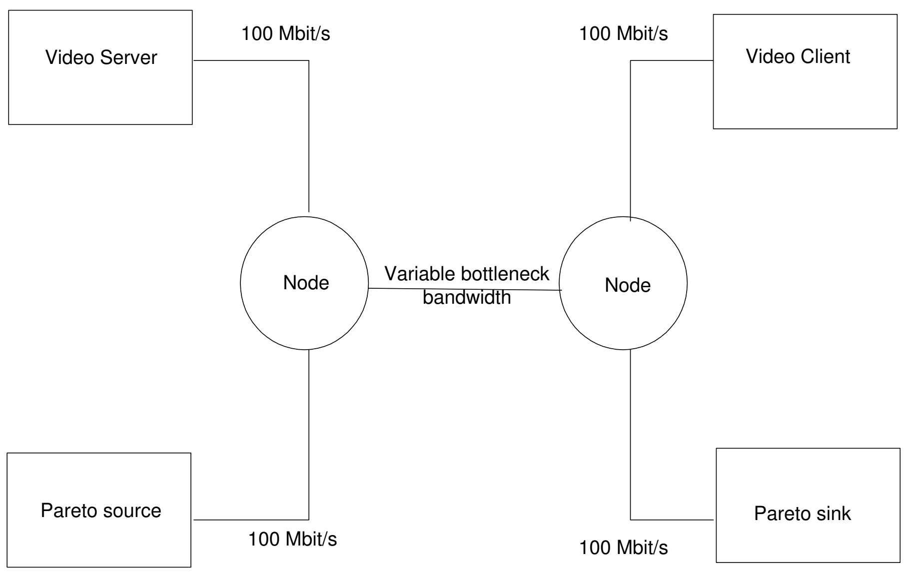

Figure 5: Simple simulation network used to determine the performance enhancement of the smoothing scheme, showing link bandwidths.

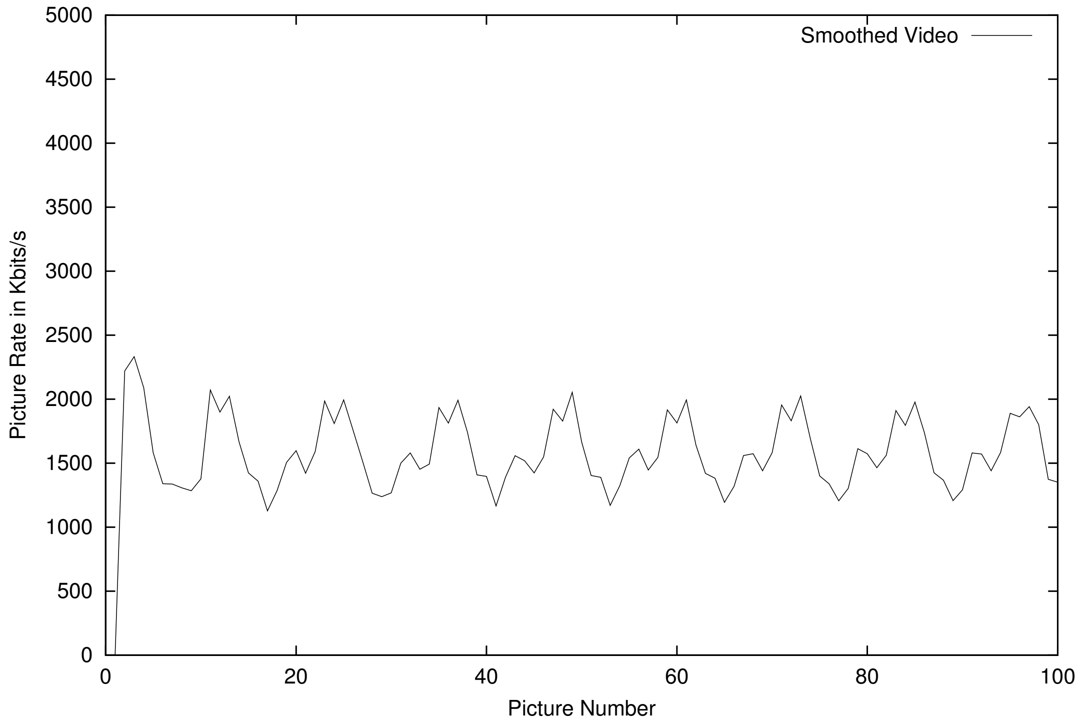

Figure 6: Bit rate of a smoothed MPEG-1 transcoded video originally encoded at 1.75 Mbit/s, showing a reduced rate burstiness.

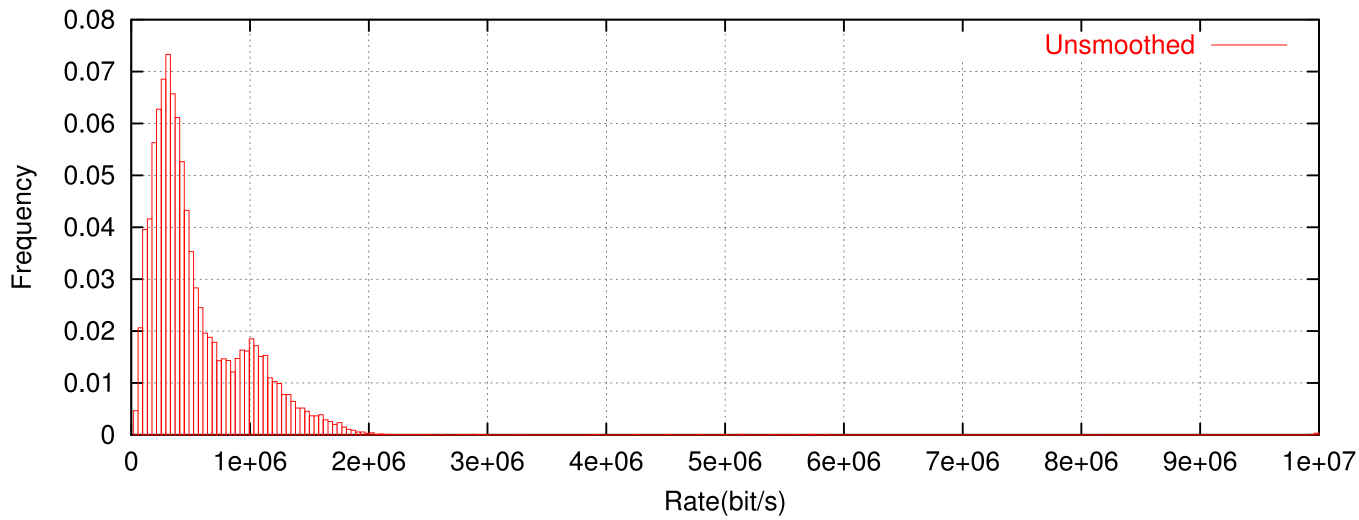

Figure 7: Frequency distribution of the packet throughput of unsmoothed bit stream for a 400 Kbit/s bottleneck bandwidth.

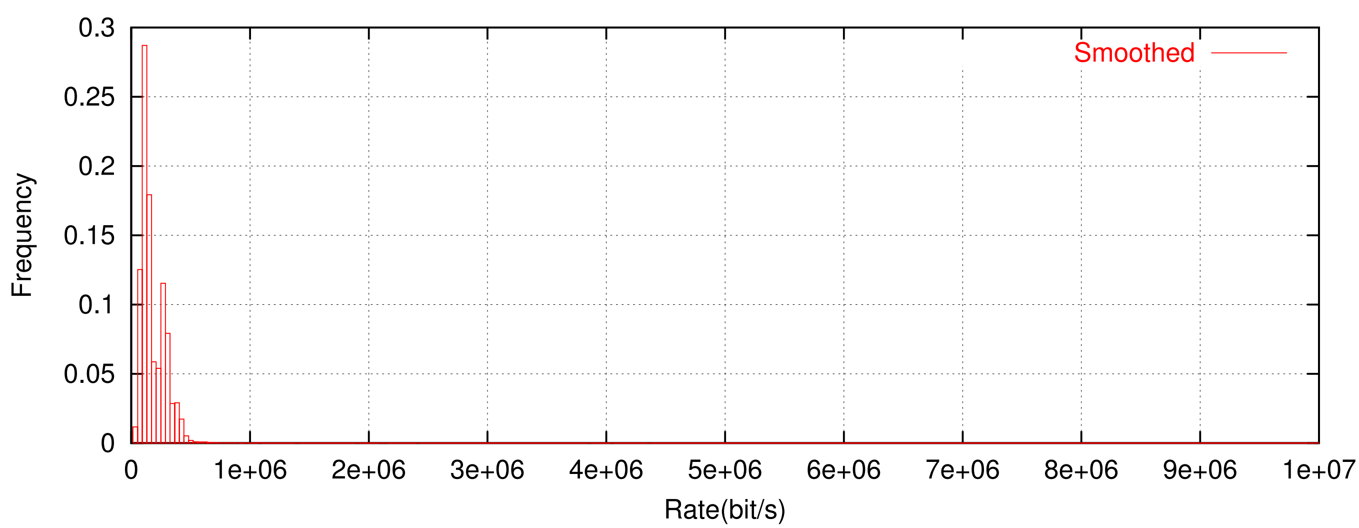

Figure 8: Frequency distribution of the packet throughput of smoothed bit stream for a 400 Kbit/s bottleneck bandwidth.

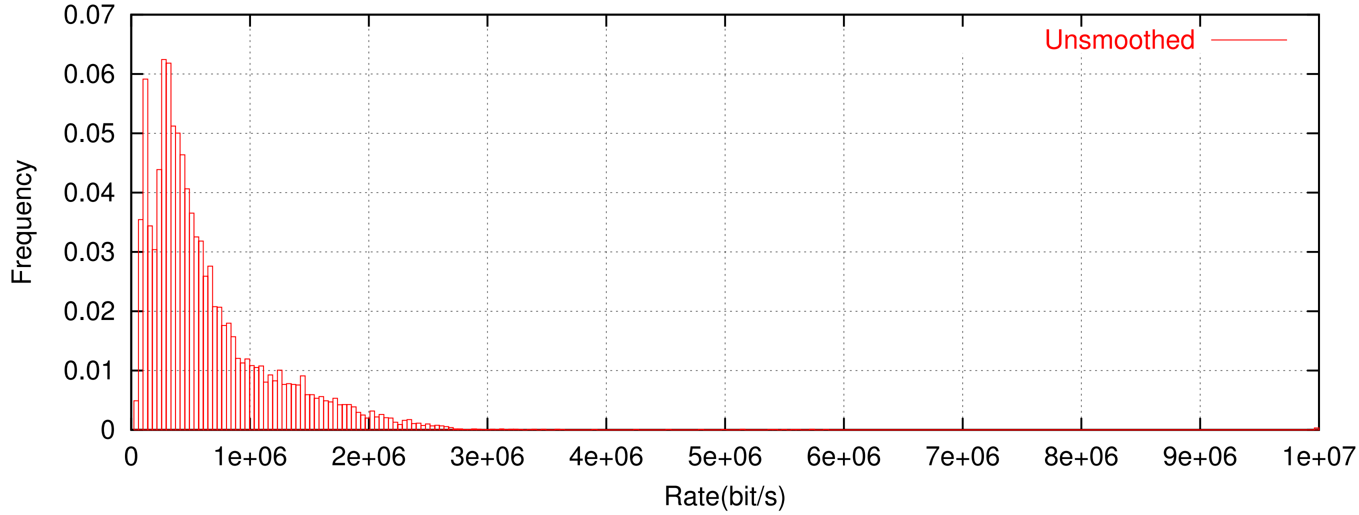

Figure 9: Frequency distribution of the packet throughput of unsmoothed bit stream for a 800 Kbit/s bottleneck bandwidth.

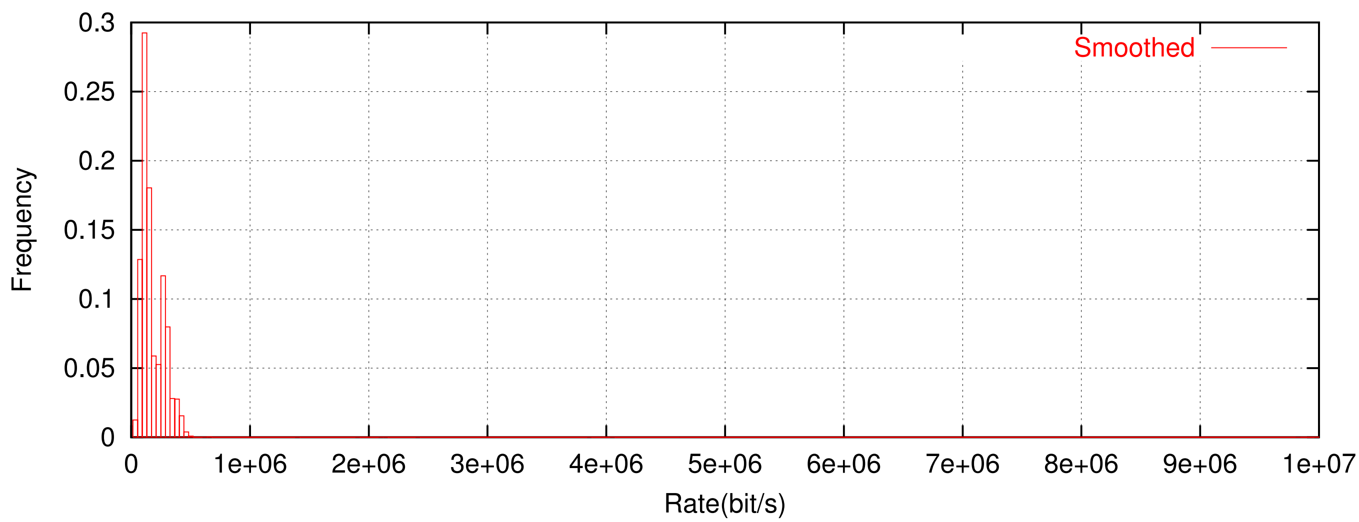

Figure 10: Frequency distribution of the packet throughput of smoothed bit stream for a 800 Kbit/s bottleneck bandwidth.

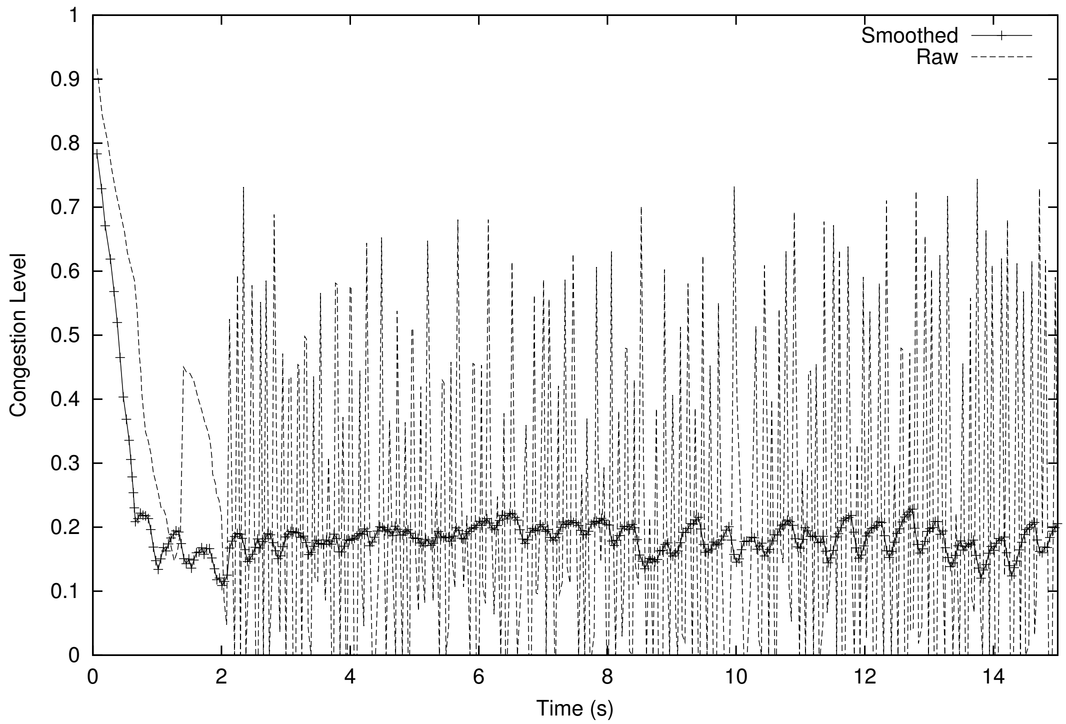

Figure 11: Graphs showing congestion level performance comparison for a 400 Kbit/s bottleneck bandwidth.

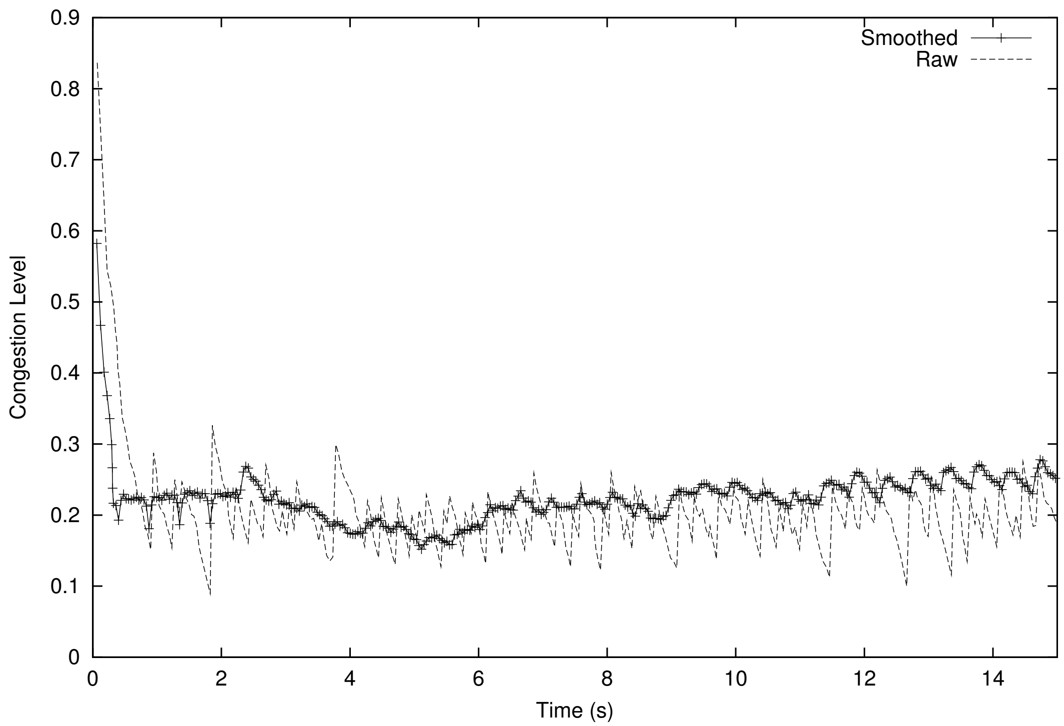

Figure 12: Graphs showing congestion level performance comparison for a 800 Kbit/s bottleneck bandwidth.

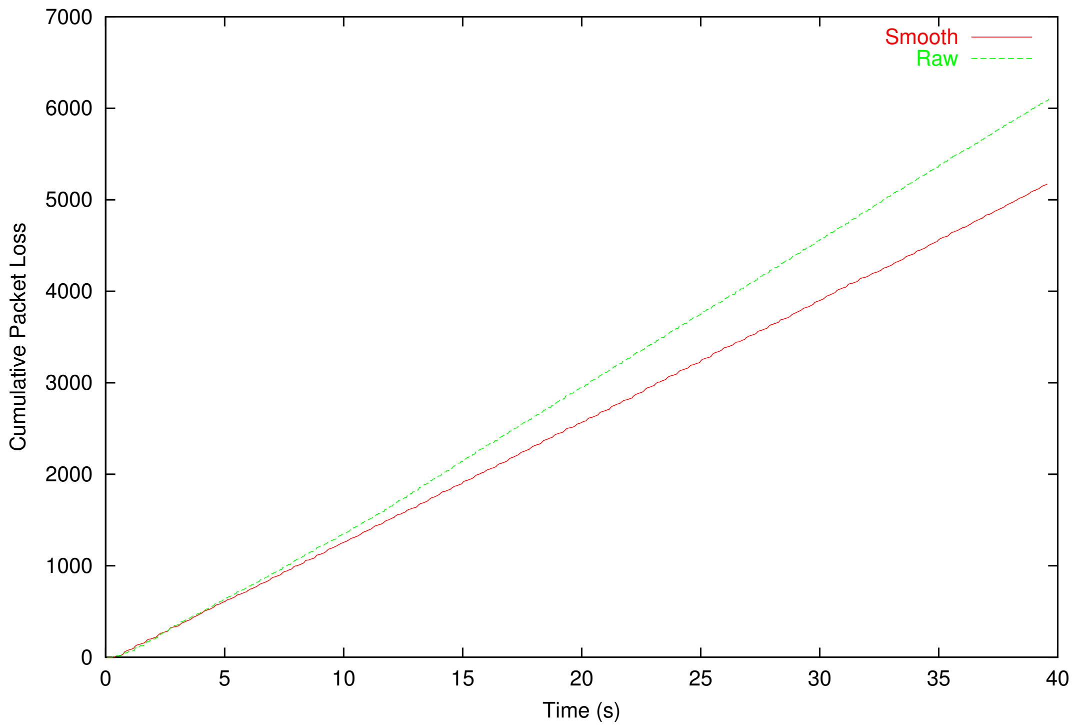

Figure 13: Graph showing packet loss performance comparison for a 1200 Kbit/s bottleneck bandwidth with no congestion control.

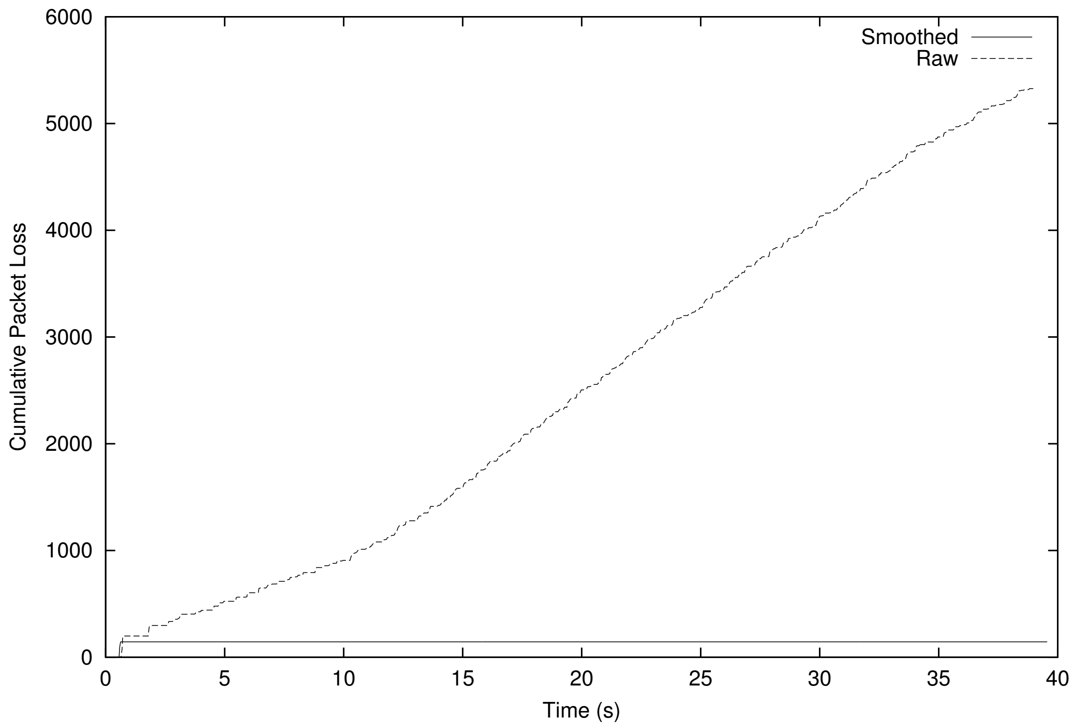

Figure 14: Graph showing packet loss performance comparison for a 400 Kbit/s bottleneck bandwidth with congestion control applied.

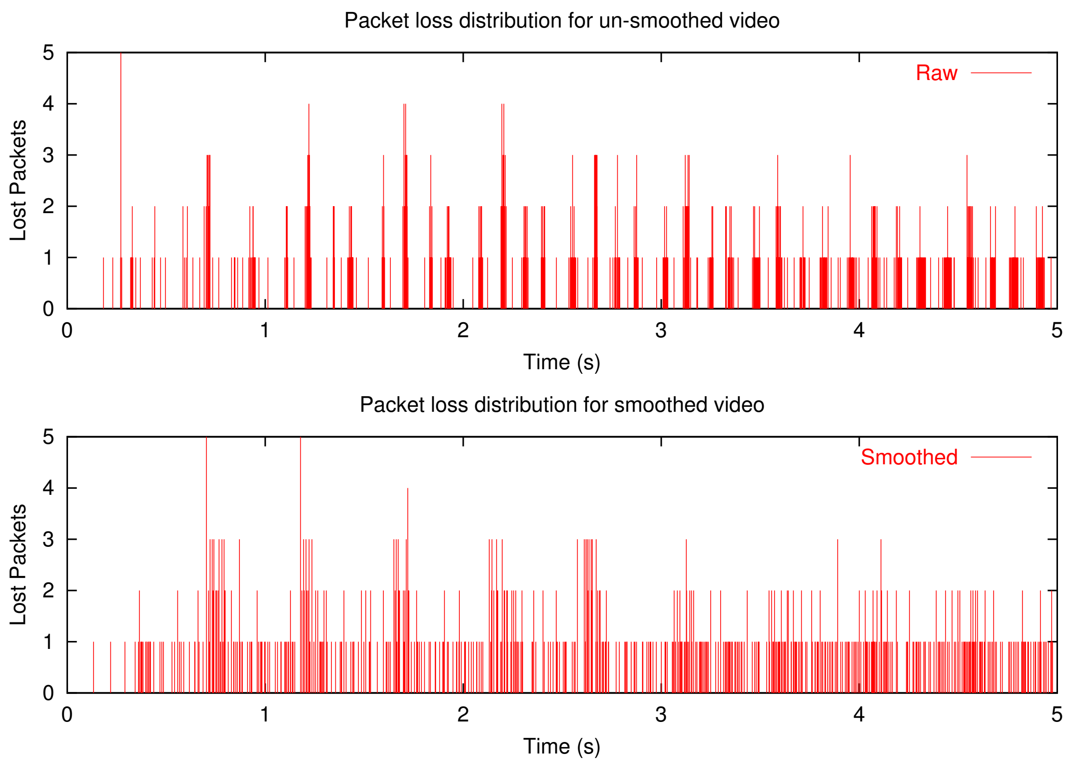

Figure 15: Graphs showing the frequency distribution of packet loss for raw and smoothed video.