12

Impact of MPEG Video Traffic on an ATM Multiplexer

Oliver Rose [1] and *Michael R. Frater* [2]

[1] *Institute of Computer Science,*
University of Würzburg, Am Hubland, 97074 Würzburg, Germany
Phone: +49-931-888-5507, Fax: +49-931-888-4601
Email: rose@informatik.uni-wuerzburg.de

[2] *Dept. Electrical Engineering, University College,*
University of New South Wales,
Australian Defence Force Academy, Canberra ACT 2600, Australia
Email: mrf@truk.ee.adfa.oz.au

Abstract

Variable Bit Rate video is expected to be a major source of traffic for the Broadband Integrated Services Digital Network. While many methods for coding and compressing this video have been proposed, it is likely that the MPEG standard will be commonly used for the delivery of new video services over high speed networks, such as the Broadband ISDN. In this paper, we explore the effect of the MPEG Group of Pictures (GOP) structure on the performance of a multiplexer carrying this video traffic from a number of independent sources. We conclude that the alignment between the GOP structures of the different sources will have a major impact on network performance and lead to difficulties concerning the Connection Admission Control and the Usage Parameter Control of MPEG video sources. We also suggest means by which this effect might be reduced.

Keywords

ATM, Broadband ISDN, Video Communications

R. Puigjaner (ed.), *High Performance Networking*

1 Introduction

The Broadband ISDN provides a very flexible digital communications network, over which many different types of services can be transmitted. This flexibility comes about largely due to the fact that all data is transmitted in the network in cells, rather than having channel bandwidth permanently allocated. As a result, services whose requirements for channel capacity change with time can be more easily accommodated. Services that might use this network include conventional voice-telephony circuits, digital data, and a number of services incorporating digital video.

In this paper, we describe the results of a number of experiments on the transmission of video sequences compressed according to the MPEG standard through an ATM multiplexer. In particular, we explore how the temporal properties of this video traffic, and in particular the alignments between different sources, affect network performance. The purpose of this paper is not to propose new models; rather, it is to examine what features of real traffic should be captured by a model in order to predict cell loss statistics.

One of the main conclusions is that the deterministic periodic properties of the traffic play a very important role, and that when many traffic streams are multiplexed onto a single channel, the effect of the deterministic periodic properties is more important than stochastic properties such as autocorrelation and long range dependence. We also conclude that there is significant potential for some form of rate shaping to increase the number of video channels that can be supported by a multiplexer of fixed capacity.

In Section 2, the properties of MPEG-compressed video are summarized. The results, which are based on switch simulations using real data as input, are presented in Section 3. The implications of these results are discussed in Section 4.

2 Properties of video sources

A video sequence consists of a series of frames, each containing a two-dimensional array of pixels. For each pixel, both luminance and chrominance information is stored. In its raw state, digital video would require a huge channel capacity for transmission (approximately 160 Mbit per second for normal television pictures). As a result, it is invariably compressed before transmission. Compression technology can be classified into two groups:

1. Constant bit rate (CBR): The output bit rate of the encoder is held constant by means of a feedback loop control. As soon as the output buffer exceeds a given limit, the coding quality is reduced to decrease the number of bytes per frame. If the number of bytes per frame is too small stuffing bits are used to increase the amount of data.

2. Variable bit rate (VBR): The output bit rate is variable, but the quality of the video is held approximately constant.

From the point of view of the network provider, CBR video has several advantages. Because of the known cell rate the connection admission control (CAC) is no problem.

During the holding time of the connection only this cell rate has to be controlled, i.e. only peak cell rate (PCR) monitoring takes place. Therefore the usage parameter control (UPC) of such a CBR source is simple, too.

For VBR video there are some problems. The definition of an effective bandwidth of a VBR video stream which is needed for CAC is difficult, because the statistical properties of video streams can be very different depending on the coding scheme and the content of the video sequence. Thus it will be hard to find a small set of parameters to calculate the effective bandwidth of this type of video streams. In a close relationship to this problem is the UPC problem. The selection of parameters of a VBR video stream to be controlled, and techniques for implementing this strategy are open questions. In [11] it is shown that the control of the sustainable cell rate (SCR) using the generic cell rate algorithm (GCRA), a method which is suggested by the ITU and the ATM Forum for the UPC of data streams, will not work, because useful SCR parameters cannot be defined for a VBR video stream.

Digital video has a number of properties that lead to quality of service requirements that differ from other services. Among these quality of service requirements are:

Cell loss. The compression of digital video removes a large amount of the redundancy present in the video images. By doing this, it increases the impact of cell loss on the quality of service. At present, it is not clear what cell loss probability will be tolerable. This will depend on a number of factors, including the sensitivity of the human visual system to different types of degradation.

Cell delay. Because of the coding and decoding there is always some delay even when there is no media access, buffering and transmission delay. The delay requirements depend on the video service. For interactive services like video conferenceing and video telephony, there should be as little delay as possible. For distribution services like video on demand and TV broadcasting the delay is normally no problem, because the user is not able to notice it. The consequence of these requirements is that no traffic shaping can be done for interactive services, because traffic shaping always produces delay due to buffering cells. Again, the level of delay tolerable in interactive services is not precisely known. One problem will be to give a reliable statement about the maximum delay introduced by the network.

It is easier to guarantee characteristics of cell loss and delay for CBR services. However, for a given network capacity, it is possible to achieve higher quality in the decoded video using VBR compression than with CBR. Hence, there is good reason to search for techniques for managing networks carrying VBR traffic. Possible algorithms may contain "constrained variability", either by sophisticated loop-back controls within the coder or by signalling schemes with the network, cells with different priorities, or multi-layer coding.

In the following we will use the MPEG coded streams [5] as an example for VBR video. The MPEG compression technique is well established now and will be used for coding videos for transmission on an ATM network in the next few years. In the following we give an outline of the statistical properties of MPEG video stream which are important for the following sections. In MPEG coded streams, there are three types of frames, each using a slightly different coding scheme:

I-frames use only intra-frame coding, based on the discrete cosine transform and entropy coding;

P-frames use a similar coding algorithm to I-frames, but with the addition of motion compensation with respect to the previous I- or P-frame;

B-frames are similar to P-frames, except that the motion compensation can be with respect to the previous I- or P-frame, the next I- or P-frame, or an interpolation between them.

Typically, I-frames require more bits than P-frames. B-frames have the lowest bandwidth requirement.

The different ways in coding frames result in different traffic characteristics for the different frame types. After coding, the frames are arranged in a deterministic periodic sequence, e.g. "IBBPBB" or "IBBPBBPBBPBB", which is called *Group of Pictures* (GOP). In addition to the distributions of the frame sizes, there are the following correlation properties of MPEG coded video streams.

- Dependences introduced by the coding algorithm due to the use of a certain GOP (short-term correlations)
- Long-term correlations within the frame process of a single stream due to the content of the film

The GOP plays the most important role concerning autocorrelation effects of an MPEG video stream coded with different frame types, because it fixes the periodic nature of the stream. This unique property of MPEG coded videos prevents us from using video models which are based on statistical data from video sequences which have only one frame type or ignore the GOP structure, like [1], [2], [4], [6], [7], [8], [9], and [12]. Thus, there is a need to develop a new model, which describes the number of bytes per frame of the coder output.

In this paper, the following simple model is used. For each frame type, (i. e. I, P, B,) we assume that the output rate of the coder is always equal to the average rate associated with that type of frame. In other words, we describe the coder output process by walking cyclically through the mean frame sizes of the GOP owned by a video sequence [10]. For the *Star Wars* sequence we will obtain therefore a sequence of 12 mean values.

This way of modeling includes the simplifying assumption, that the correlations among the frames of one GOP which are caused by the MPEG coding scheme can be neglected for our study. The long-term dependences among frames of consecutive GOPs, e.g. the correlations introduced by similar pictures of one movie scene, seem also to be less important in our case.

3 Multiplexing of MPEG VBR video streams

In this section, cell loss results for a multiplexer buffer are presented. Cells waiting in the multiplexer buffer are served in a FIFO manner. There are 16 input streams, each of which is constructed from real data obtained by coding and compressing the *Star Wars* movie sequence [3]. The sequence consists of 174126 frames, which corresponds to about 2 hours of movie. The GOP of this video is "IBBPBBPBBPBB". The size of the decoded video frames is 504 x 480 Pixels.

Frame type	Number	Mean	Min	Max	CoV
		in cells			
all	174126	41.12	2	483	1.15
I	14511	157.74	31	483	0.33
P	43531	60.58	6	454	0.63
B	116084	19.25	2	169	0.65

Table 1: *Statistical data of the Star Wars sequence.*

Table 1 shows some statistical data of the video stream and Figure 1 shows the histograms of the number of cells per frame of the different frame types. In both diagrams the frame size is measured in ATM cells, where a payload of 48 bytes per cell is assumed.

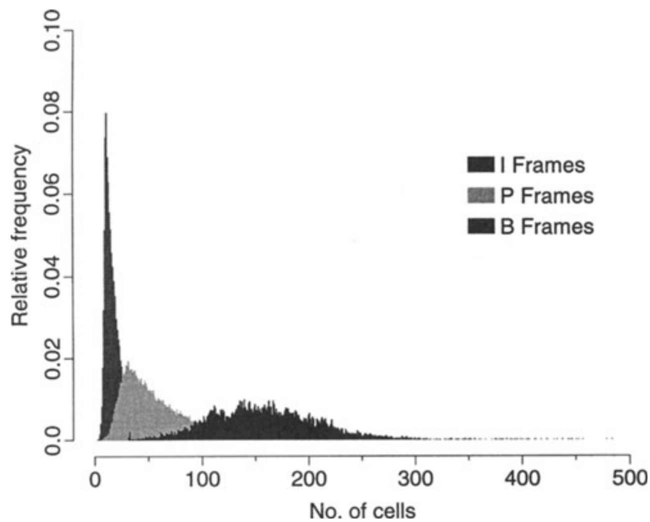

Figure 1: *Distributions of the I-, B-, and P-frame sizes of the Star Wars sequence*

To show the influence of the GOP pattern on the cell losses we chose four ways of multiplexing. The first three alternatives are frame based, the fourth one is GOP based. The time interval to transmit one GOP is denoted by T_{GOP}, and the total amount of data transmitted during T_{GOP} is called V_{GOP}. It is assumed for the frame based alternatives that the video data cells are equally distributed over the whole frame duration and therefore the cell rate is constant during the transmission of a frame.

- **High source alignment**: All sources start the GOP at the same time, i.e. all sources transmit their I-, P-, and B-frames during the same time intervals.
- **Random source alignment**: The starting time of the GOPs is chosen randomly.
- **Low source alignment**: The starting times of the GOPs of the N multiplexed streams are shifted by a time interval of the length T_{GOP}/N, i.e. the minimum overlapping of the I-frames is achieved.
- **Rate averaging**: The starting time of the GOP is chosen randomly and the data of each GOP is transmitted at the rate of V_{GOP}/T_{GOP}, i.e. the rate is constant during the transmission of one GOP.

Figures 2 to 4 show the cell loss results for the first three alignment strategies for multiplexer buffer sizes of 10, 100, and 1000 cells.

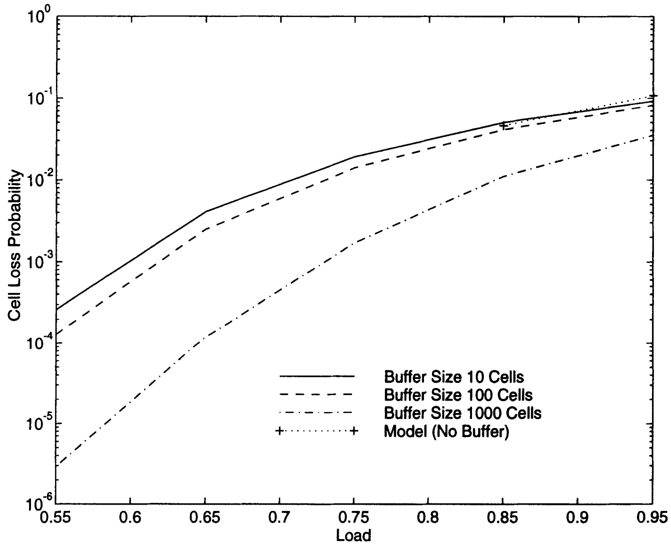

Figure 2: *Cell loss probability for Low Source Alignment*

For the Low Source Alignment case (Figure 2) we get the smallest cell loss rates. The Random Source Alignment (Figure 3) leads to higher cell losses and the High Source Alignment (Figure 4) shows the worst performance with cell losses of more than 10 % for multiplexer loads of more than 50 %. Both the best and the worst case alignment scenario may appear in real multiplexers, but not frequently. In this study, we use these scenarios to show which properties of MPEG coded video stream have a major impact on the multiplexer performance.

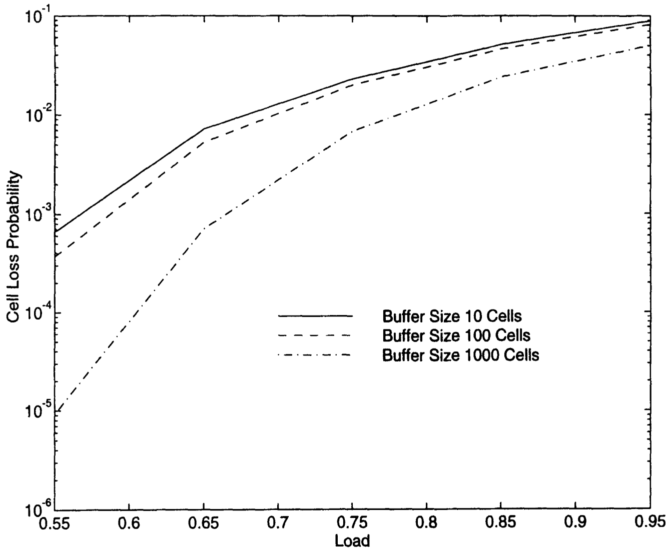

Figure 3: *Cell loss probability for a Random Source Alignment*

In Figures 2 and 4 the cell loss results using the simple model presented in Section 2 as input for a multiplexer with no buffer are also shown. These curves show good conformance with the other results. The good agreement between the predictions of the model and the simulation results indicates that the cell losses are occurring due to the alignment of the GOP structures, rather than the statistical variations within these structures.

Figure 5 is a further indication of the importance of the GOP structure compared to the actual distribution of the frame sizes. The diagram shows the cell loss curves for randomly aligned sources, where the frame sizes are stepwise replaced by their average values. The solid line shows the losses for the original data set. If we replace the frame sizes of the B frames by their average value, only for a load of less than 0.7 the losses are significantly underestimated. If the P frames are replaced, the situation is almost the same. Even if we replace all frames by their corresponding average values, for loads larger than 0.8 the curves are in good agreement. We conclude from this behavior that in case of medium to high loads, the cell loss probabilities depend mainly on the periodic structure of the GOP and the average frame sizes, but not on the other sources of rate variation. This result is even more convincing, if we regard the curves of the mean cell losses +/- standard deviation of the original data set. All curves obtained by averaging can be found within these bounds. This means that the variations of the mean cell losses caused by different alignments of the video streams have a larger effect on the cell losses than replacing the actual frame sizes by their average values.

In Figure 6 the frame based strategies are compared to the Rate Averaging alternative

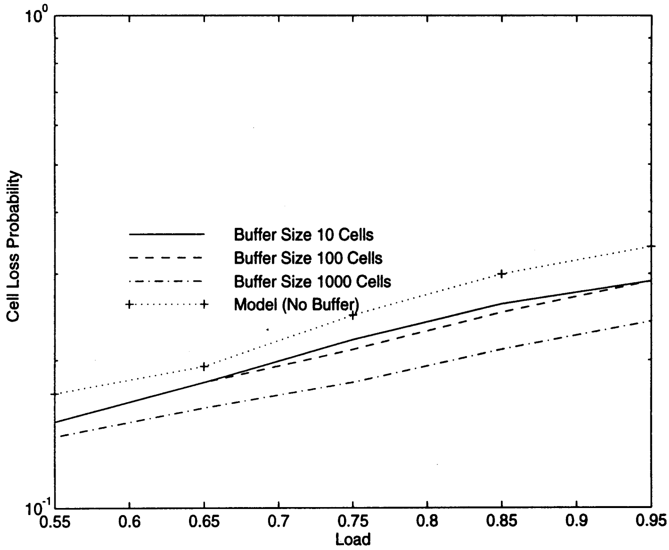

Figure 4: *Cell loss probability for High Source Alignment*

for a buffer size of 100 cells. The cell losses using Rate Averaging are very low, for lower loads the difference is orders of magnitude.

In a further experiment, the multiplexer capacity was fixed at 155 Mbps. Simulations were performed to find out how many VBR video sources could be multiplexed for a specified cell loss probability. These results are shown in Table 2.

Strategy	10^{-3}		10^{-5}	
	N_{max}	Load	N_{max}	Load
Rate Averaging	40	0.85	39	0.82
Low Alignment	34	0.73	33	0.71
Random Alignment	30	0.64	29	0.62

Table 2: *Maximum number of Sources N_{max} and corresponding multiplexer load*

In these results, it is clear that the results obtained using the sequence Rate Averaging over the length of each GOP are superior to all other schemes, amounting a performance increase of approximately 16 % over the best case without averaging (Low Alignment), and approximately 30 % over the randomly chosen frame alignment.

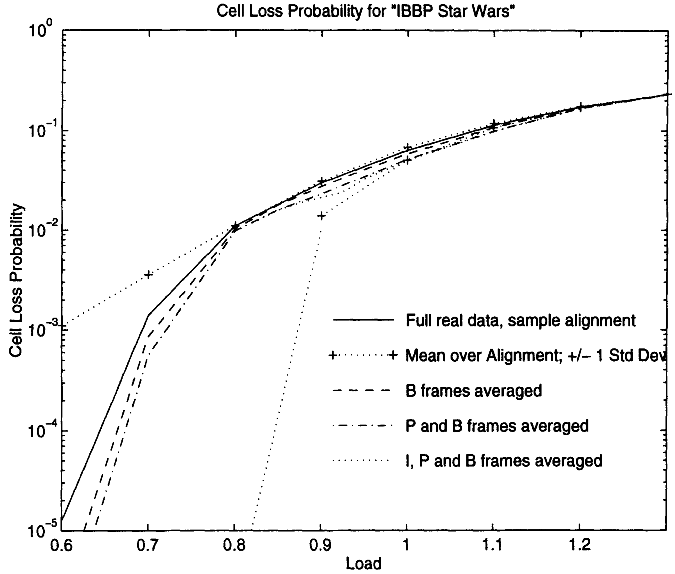

Figure 5: *Effect of exchanging variable frame sizes by their mean values*

4 Discussion

From the results presented above, there are a number of obvious conclusions. Firstly, when a number of MPEG video sources are fed into a single multiplexer, the cell loss probability is very sensitive to the alignment between the GOP structures of these sources. This means that it will be very difficult to guarantee the cell loss probability for a multiplexer (or alternatively to perform the CAC and UPC functions of the network), without knowing in advance how the GOP structures of the video sources are aligned.

This suggests that it would be useful for the network to be able to control this alignment. However, this is not feasible in practice. Hence, there can be no guarantee that performance close to the best case in our plots will be achieved. In fact, it is certain that sometimes the performance will approach the worst case, with cell loss probabilities exceeding 10 %.

The fact that an quite simple model, which does not take into account any statistical variation in the rates, can accurately predict the cell loss probability in some cases, indicates that this loss is occurring due to the average behavior of these sources. In other words, it is the deterministic behavior of the sources that causes the arrival rate at the multiplexer to exceed its service rate. This is not at all surprising for the high source alignment, where all of the sources transmit their I-frames simultaneously. However, the fact that it occurs for low source alignment for loads above 0.85 suggests that only much

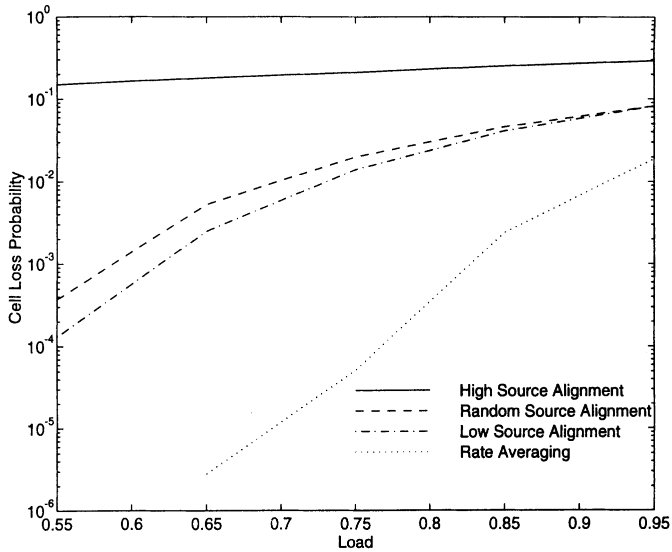

Figure 6: *Cell loss probability for a buffer size of 100 cells*

smaller offered loads will be acceptable. Even with a relatively large buffer size of 1000 cells, a load of not much more than 50 % will be required in order to keep the cell loss probability to even 10^{-5}.

We have observed that a considerable increase in the cell-loss performance can be obtained by having sources average their rate over the period of a GOP. For the case of 16 sources feeding a switch, we have found that the loss probability is reduced by at least a factor of 10^{-2} for reasonable network loads. We have also found that for a 155 Mbit/s switch, the total offered load that can be accommodated for a cell loss probability of 10^{-5} is increase by approximately 30 % by rate averaging.

Rate Averaging, as described here, will create large delays for the video transmission, because a whole GOP has to be stored until the transmission rate for this GOP is known. For example, if a twelve frame GOP is used for coding, this strategy will cause an extra delay of about half a second. In addition to that larger buffers are needed at the encoder. Because of these properties, the Rate Averaging cannot be used for interactive video services like video conferencing and video telephony. For TV and movie distribution, which are not delay sensitive, it is applicable. However, we note that current CBR coders perform a similar averaging process without this large delay. We suggest that a good compromise might be obtained by using a modified form of Rate Averaging, which smoothes out the variation due to the different frame types, but does not impose absolutely the constraint that the rate remain constant for the whole length of a GOP.

It is unlikely, that Low Alignment of the frames of several (at least tens) video sources can be forced in practice. Therefore, unless some form of rate averaging is used, the CAC must be based on a statistical analysis of both the properties of the individual sources, and their interactions. However, one should keep in mind that High Alignment case is not ruled out by this approach. Another fact that will worsen the results is that in practice MPEG video will not fill a whole 155 Mbps channel and thus the multiplexing gain will be smaller than suggested by the results of this study.

We have noted that the alignment between the GOP structures of sources feeding a multiplexer will have a significant effect on the cell loss performance of this multiplexer. One of the consequences of this is that the quality of service given to the different video traffic streams will be very different. Specifically, those sources that are unfortunate enough to have their I-frames closely aligned will experience very much more loss than other sources.

5 Conclusion

In our study, we examined the effect of the GOP structure on the performance of a multiplexer carrying MPEG video traffic. We showed, that the cell loss probability is very sensitive to the alignment between the GOP structures of the multiplexed video sources. This leads difficulties defining CAC and UPC procedures for video sources, because depending on the amount of alignment a group of video sources can show very different impacts on a network. In addition, the cell losses perceived by the video sources can be very different. We also presented some results for averaging the cell over a GOP. With this method the cell loss are reduced and the cell losses are equally distributed among the multiplexed video sources. This is achieved however at the cost of buffering delays.

A subject of further research will be to have a closer look on the tradeoff between delay and cell losses, i.e. averaging strategies with large delays vs. frame by frame transmission with large cell losses.

In our study, we used only transmission strategies which did not affect the way of coding of the video sequence. Another way of smoothing the video streams or improving their ATM suitability is to change the coding scheme. It is still possible to use MPEG coding, but adapted to the current network needs. As mentioned in Section 2 this can either be done by controlling the encoder output rate in a preventive manner within the encoder or in a reactive way via signalling from the ATM network.

References

[1] C. Blondia and O. Casals. Statistical multiplexing of VBR sources: A matrix-analytic approach. *Performance Evaluation*, (16):5–20, 1992.

[2] M. R. Frater, J. F. Arnold, and P. Tan. A new statistical model for traffic generated by VBR CODECs for television on the broadband ISDN. *IEEE Transactions on Circuits & Systems for Video Technology*, 4(6):521–526, Dec. 1994.

[3] M. W. Garrett. *Contributions toward real-time services on packet switched networks.* PhD thesis, Columbia University, 1993.

[4] D. P. Heyman, A. Tabatabai, and T. V. Lakshman. Statistical analysis and simulation study of video teleconference traffic in ATM networks. *IEEE Transactions on Circuits and Systems for Video Technology*, 2(1):49–59, Mar. 1992.

[5] D. Le Gall. MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):46–58, Apr. 1991.

[6] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J. D. Robbins. Performance models of statistical multiplexing in packet video communications. *IEEE Transactions on Communications*, 36(7):834–844, July 1988.

[7] B. Melamed and B. Sengupta. TES modeling of video traffic. *IEICE Transactions on Communications*, (12):1292–1300, Dec. 1992.

[8] P. Pancha and M. E. Zarki. Bandwidth requirements of variable bit rate MPEG sources in ATM networks. In *Proceedings of the Conference on Modelling and Performance Evaluation of ATM Technology, Martinique*, pages 5.2.1–25, Jan. 1993.

[9] G. Ramamurthy and B. Sengupta. Modelling and analysis of a variable bit rate video multiplexer. In *Proceedings of the Infocom '92*, pages 6C.1.1–11, 1992.

[10] O. Rose. Approximate analysis of an ATM multiplexer with MPEG video input. Institute of Computer Science Research Report Series 79, University of Würzburg, Am Hubland, 97074 Würzburg, Germany, Jan. 1994.

[11] O. Rose and M. Ritter. MPEG-video sources in ATM-systems – a new approach for the dimensioning of policing functions. Institute of computer science research report series, University of Würzburg, Am Hubland, 97074 Würzburg, Germany, Aug. 1994. Submitted to the International Conference on Local and Metropolitan Communication Systems: LAN & MAN, December 1994, Kyoto, Japan.

[12] P. Sen, B. Maglaris, N.-E. Rikli, and D. Anastassiou. Models for packet switching of variable-bit-rate video sources. *IEEE Journal on Selected Areas in Communications*, 7(5):865–869, June 1989.