4

Correlation properties of MPEG VBR video sources: Relative importance of short and medium term correlation on ATM multiplexer performance

Ragnar Ø Andreassen
Telenor R&D
Mail: P.O. Box 83 Email: ragnar.andreassen@nta.no
N-2007 Kjeller Telephone: +47 63 84 88 51
Norway Fax: +47 63 81 00 76

Abstract

An investigation of correlation properties of MPEG VBR video based on three different video sequences is presented. It is found that the short range properties of these sources resemble those of uncorrelated sources. To investigate the relative importance of short and medium term correlation properties versus the longer term correlation properties of MPEG sources, analytical source models taking into account the periodic properties of the sources are used. Results from these models are compared with trace driven simulations in which long term properties are inherent. One conclusion of this investigation is that for the ATM cell loss probability, interesting regions in terms of buffer capacity and number of sources exist, where long term correlation properties will not influence multiplexer performance. For a number of 12 multiplexed sources, this region extended to buffer lengths of about 10 times the mean video frame size in high loss regions. In low loss regions, no effect of longer term correlation properties was seen at buffer lengths up to 15 times the mean frame lengths, with the same number of sources.

Keywords

ATM multiplexing, MPEG VBR video, correlation properties

1 INTRODUCTION

With the advent of multimedia services like teleconferencing and video distribution, it is expected that video can become a major traffic source in ATM networks. MPEG is the international ISO/ITU standard for coding/decoding of digital video signals, and it is likely that a major part of the video information transmitted will be MPEG signals. In the process of engineering networks for the transmission of such signals, there will be a demand for robust, tractable and precise models for the MPEG VBR video sources. It is by now well established that many video sources display self-similar properties, and it may be seen that the MPEG sources are also autocorrelated over large time scales. The reduction of temporal

L. Mason et al. (eds.), *Broadband Communications*

redundancy as performed by the MPEG algorithm will additionally induce periodic components in the produced information streams, invoking strong periodic correlations at the frame level. In the process of developing models for multiplexer loss behaviour of such sources, it is not evident which properties should be most precisely modelled: short, medium or long term correlations. This question may not be given a simple answer, as the parameters of buffer length and number of sources, and the considered loss probability region, will determine which properties of the source will dominate.

The aim of this paper is to shed some light on the questions outlined above. To this end we perform an investigation of correlation properties of MPEG sources using three different approaches: autocorrelation functions, time-variance plots and leaky bucket loss contour diagrams. To determine which correlation properties dominate loss behaviour of multiplexed MPEG sources, we compare trace driven simulations, inherently containing long term correlation properties, with Markov models of MPEG sources. By using Markov models in the investigation, it will be possible to utilize certain source properties in a controlled manner. To aid our purpose, a Markov model is presented which makes possible the modelling of the periodic properties of multiplexed MPEG sources.

Inspirations for the current paper are the results published in (Garret, 1993), (Garret, 1994) and (Beran, 1995) demonstrating the self-similar properties of video sequences, and the work documented in (Reininger, 1994) and (Huang, 1995) developing simulator models for MPEG sources taking into account source correlation properties. The Markov model presented in a later section is related to the ones found in (Hübner, 1994), and (Landry, 1994). In the treatment of the periodic property of MPEG sources, a combinatorial approach is presented which may be employed both in the context of simulations and analytical models. Such an approach is, to the knowledge of the author, not documented by others.

2 GENERATION OF VBR MPEG VIDEO TRAFFIC

Three frametypes are defined in the MPEG algorithm: intra (I) frames, predicted (P) frames and interpolated (B) frames. A group of pictures (GOP) is composed of a regular pattern of frame types headed with an I frame succeeded by B and P frames. In the coding process, I and P frames may be used as references: I frames are coded with no reference to other frames, P frames are coded with reference to the previous reference frame and B-frames may be coded with references to both the previous and the succeeding reference frame. Hence the B-frames usually contain the least information, P frames somewhat more than B frames, and I frames more than both B and P frames. Typically the sizes of I versus P/B frames will differ by about a factor of 3 to 10, but the relative sizes of the frametypes may vary considerably between sequences. A source may keep its GOP phase during a connection, or it may change phase at scene boundaries in an effort to reduce information transmitted. Thus the phase constellations at a node in the network will vary in an arbitrary manner. The period and structure of the GOP cycle are coding parameters that may be chosen.

The quantization parameter (q), controls the lossy part of the coding algorithm, in effect regulating the coarseness of the decoded picture. Active use of the q-parameter during encoding may be employed as a flow regulation mechanism. In the sequences used here, fixed quality video is investigated, and the q-parameter is set at a constant value for each frame type.

We shall use statistics obtained from three film sequences, each consisting of 40 000

frames, corresponding to about half an hour of realtime video. One is an action film (James Bond), referred to as Bond, another is a 'psycho thriller' (The silence of the Lambs), referred to as Lambs, and the last contains episodes from a cartoon series (Simpsons), referred to as Simpsons. All sequences are coded with a GOP size of 12 frames, following the pattern of **IBBPBBPBBPBB**. Further documentation concerning production of the traces may be found in (Rose, 1995). In table 1 we display maximum (a_{max}) and expected number of cells (E(a)) in frames, where indices *I*, *P* and *B* refer to the frame types

Table 1 Trace statistics

Trace	a_{max}	E(a)	E(a_I)	E(a_P)	E(a_B)
Bond	694	69.1	236.6	117.7	29.9
Simpsons	682	52.8	210.4	61.2	29.9
Lambs	381	20.8	108.0	21.1	9.7

Here and in all subsequent numeric computations, a cell payload of 44 bytes is assumed.

3 CORRELATION PROPERTIES OF VBR MPEG VIDEO

Temporal redundancies in frame contents and temporal correlation of frame sizes in a sequence of frames are related quantities: If 100% removal of temporal redundancy was achieved, only 'new' information should be transmitted in each frame. The number of bits required to code this 'new' information should be unpredictable, otherwise the information would not be 'new'. So, a perfect (in this respect) coding scheme will result in an uncorrelated sequence of frames. On the other hand, a completely redundant stream, such as repeating the same frame, will obviously result in perfect correlation. It then seems reasonable to expect that in reducing temporal redundancy, the temporal correlation properties should also be weakened in a manner reflecting the properties of the redundancy removing mechanism. We shall in the following perform a qualitative investigation of temporal properties of MPEG sequences by three different methods, to see if some compliance with the above informal argument may be observed.

Firstly, we will consider the autocorrelation functions. In Figure 1, autocorrelation at the frame and group levels are displayed for the different traces, restricting the lag range to 10 000 frame periods as to cut off noise effects which may be observed at longer ranges. Lag scale for group level curves is adjusted so as to match the frame level curves. The periodic property of the sources is reflected in the autocorrelation functions, and the correlation amplitudes at different lags are related to the relative mean sizes of frames in the GOP structure. This can be observed by comparing the graphs with the numbers in Table 1. In a deterministic periodic sequence with the above group structure, the amplitude at lags corresponding to distances between B-frames and reference frames (lag values $i+3\cdot j$, i=1,2 , j=0,1,2,...) would be negative. A strong positive relative correlation between neighbouring frames of different types will tend to raise the above correlation values, and may explain the lack of negative amplitude as can be observed in two of the sequences. Group autocorrelations fall off very slowly with increasing lags, indicating the contents of GOPs to be strongly correlated over large time scales. As redundancy reduction is confined to groups, this behaviour is as expected considering the results in (Garret, 1994) and (Beran, 1995).

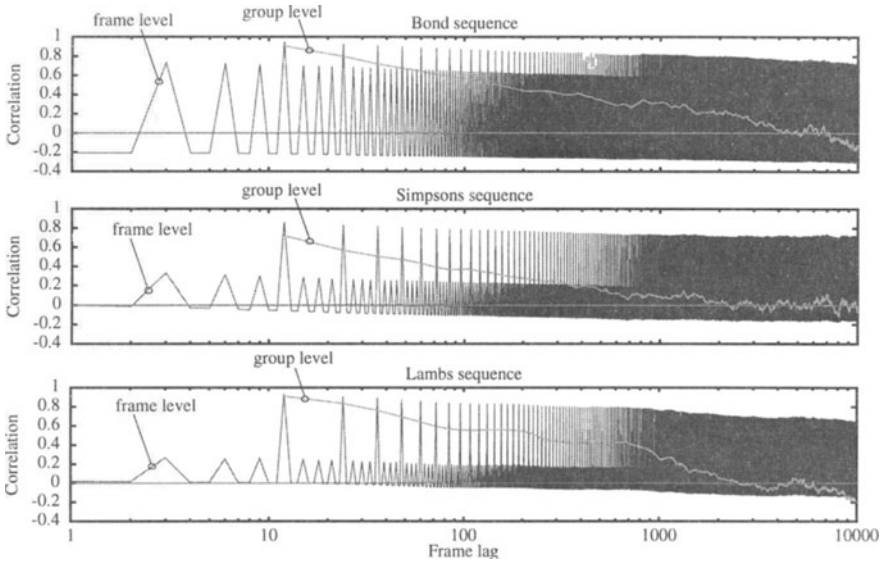

Figure 1 Autocorrelation functions

As another way of investigating temporal properties of MPEG sources, we shall employ the method of time-variance plots. Following (Garret, 1993), sequences of lengths S may be partitioned into a number of blocks $T = \lfloor S/M \rfloor$ for some block size M. Defining the block mean stochastic variable $Y_k = \frac{1}{M}\sum_{i=(k-1)M+1}^{kM} a_i$, where a_i denotes frame sizes, one may form the general expression $s_Y^2 = M^{-\beta} s_a^2$, where s_Y^2, s_a^2 denotes the sample variances. This may be expressed as $ln(s_Y^2)/ln(s_a^2) = 1 - \beta ln(M)/ln(s_a^2)$. Plotting the empirical left hand side of this equation against $ln(M)/ln(s_a^2)$ should result in a linear curve with unity negative slope in the case of an independent time series. i.e. $\beta = 1$. For an exactly self-similar sequence, β should remain constant in the range of [0,1> over all time scales, whereas for an asymptotically self-similar process, β will converge towards a constant value in the same range as the sample sizes increase. In Figure 2 are plotted the time-variance graphs for the sequences (denoted 'Original'). In the same diagrams are plotted the $\beta = 1$ line (denoted 'Independent'), and the time-variance curve for an artificial sequence (denoted 'Group') obtained by subsequently drawing independently the size of the 12 frames of the GOP structure from their respective marginal distributions. The following properties of these diagrams may be noted: The first few points up to $M \approx 10$ follow the independent-curve more or less closely. In this region, the periodic property will also influence the time-variance graph, which can be noted by comparing the 'Group' curve and the 'Original' curve. This effect is strongest in the Bond sequence. As block levels increase, the two curves depart due to the lack of correlations above frame level in the 'Group' sequences. In many sequences

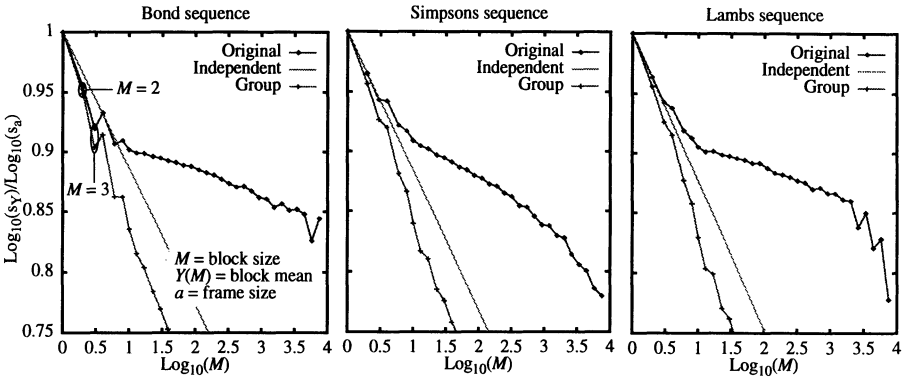

Figure 2 Time variance plots

we have investigated, the rest of the graph will be partitioned in two fairly distinct regions, a middle region with very low β-value, and a region comprising the larger blocks showing a somewhat higher β-value. This behaviour is most clearly seen in the Simpsons sequence.

The last investigation in this section will be devoted to leaky bucket loss contour diagrams. Such diagrams have been proposed in (Lucantoni, 1994) and (Garret, 1993) as an alternative way of characterizing the statistical properties of data sources, and will give an operational picture of source properties. The points of the curves in the diagram were obtained by fixing leaky bucket sizes, then adjusting the load of a trace driven simulator until a specified loss probability was achieved. In each diagram in Figure 2, curves are shown for the real source (denoted 'Original') and an artificial sequence generated by independently drawing frame sizes from the marginal distributions (denoted 'Independent'). In all diagrams can be observed a region in which the two traces show a similar behaviour. As buffer lengths increase, low and medium loss iso-lines of real sequences depart abruptly from the loss iso-lines of the independent sequences, indicating the effect of correlations at the group level.

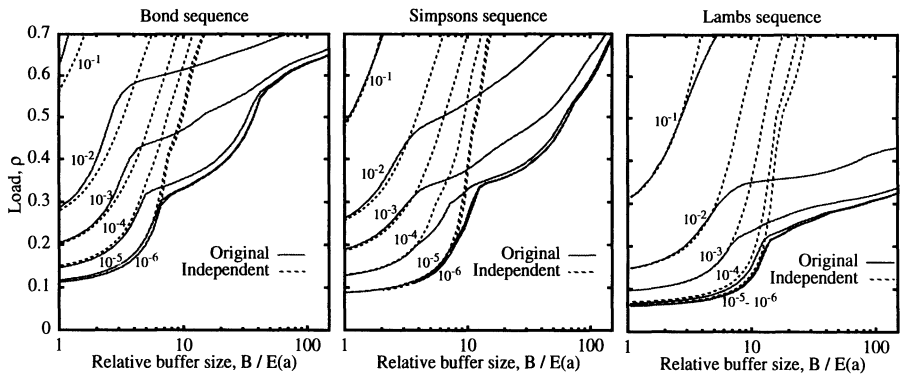

Figure 3 Leaky bucket loss contour diagrams

We draw the following conclusions from these investigations: At time scales smaller than the GOP period, the single MPEG sources investigated display behaviour resembling that of uncorrelated sources. This agrees very well with the operational principles of the MPEG algorithm: removing temporal redundancy within each GOP, but not across GOPs.

4 SOURCE MODELS

4.1 Single source

This system consists of one source generating frames at a fixed rate into a buffer served with constant output. Frames are segmented into cells when fed into the buffer. We shall consider the case where cells are spaced evenly through the inter frame period ('frame stretching'). The number of cells in subsequent frames follows the marginal distribution of cells per frame from a real source. The following definitions will be needed:

- n_k r.v. for the number of cells in the buffer immediately before arrival of the first cell of the kth frame at time t_k
- a_k r.v. for the number of cells in the kth frame.
- d The number of cells that may be serviced in the inter frame period. When cell period is set at unity, d denotes the inter frame period.
- b Buffer limit

It is assumed that the time intervals t_k - t_{k-1} are integral cell periods, which may be justified by observing that the frame period is much longer than the cell period. As frames are stretched, and assuming service before arrival, the system may reach the state of b-1 at the end of the frame period. This situation will happen if the system overflows, or the cell flow into the buffer exactly fills it up during the frame period. At instants t_k, the following recursion relation will then apply:

$$n_{k+1} = Max(Min(n_k + a_k - d, b - 1), 0) \tag{1}$$

It may be noted that this way of modelling the arrival process will correctly handle short term autocorrelation in the cell interarrival times caused by the frame stretching mechanism. An interesting feature of the above relation is a certain scaling property. If, by performing the necessary multiplications with the cell payload, the problem is recast using information content rather than cell numbers, we get the relation:

$$\eta_{k+1} = Max(Min(\eta_k + \alpha_k - \delta, \beta), 0) \tag{2}$$

Where the obvious identifications of symbols for information content rather than cell numbers may be performed. Going to the fluid limit by letting payload size P diminish towards zero, keeping $\eta \equiv n \cdot P$ constant, it may be seen that the resulting problem is completely scalable insofar that for the information loss ratio function Λ, we have that $\Lambda(\delta, \beta) = \Lambda'(\delta', \beta')$ if $\delta' = s\delta$, $\beta' = s\beta$, $\alpha'_k = s\alpha_k$. An important implication of this observation is the invariance of the solution to the relation $\beta/E(\alpha)$, i.e. the ratio between

buffer size and expected frame size. The discrete cell case is not accurately scalable as a scaling of cell size will change the granularity of the involved stochastic variables. The above result will nevertheless give the natural scaling when results from different sources are compared, and as no assumptions about the nature of the sources other than the frame stretching mechanism is made, will have a broad range of application.

We now embed a Markov chain at instants t_k letting n_k denote the states of the Markov model. To satisfy the Markov condition, arrivals a_k must then be assumed to be independent and identically distributed. Omitting sequence indexes, the stationary state probabilities π_j may be expressed as:

$$\pi_j = \begin{cases} 1 - (F * \pi)_{(b-1)+d-1} & , \quad j = b-1 \\ (A * \pi)_{j+d} & , \quad 0 < j < b-1 \\ (F * \pi)_d & , \quad j = 0 \end{cases} \tag{3}$$

Where

$$A_j = Pr(a=j) \qquad F_j = Pr(a \le j) \qquad \pi_j = Pr(n=j) \qquad \sum_j \pi_j = 1$$

And the $*$ operator denotes the discrete convolution operation consistent with the ranges of the operands. If the steady state equation (3) is written out in matrix form as $\pi = \pi \boldsymbol{P}$, the transition probability matrix will be of order b, the buffer length.

Generating function approaches exist both for infinite buffer length (Lin, 1990) and finite buffer lengths (Østerbø, 1995). We have however chosen a numeric solution method iterating on the probability transition matrix. This method turned out to be very fast.

Cell loss probability R can be determined from the ratio between overflowed cells and arrived cells:

$$R(\rho, b) = \frac{1}{E(a)} \sum_{j=1}^{a_{max}-d} j \cdot (A * \pi)_{(b-1)+d+j} \tag{4}$$

The load ρ is given by $\rho = E(a)/d$.

A generalisation of this model to a Markov model consisting of a group of N frames with independent but different framesize distributions is possible. Embedding the Markov chain at the group level rather than at the frame level, the resulting equilibrium equation may be expressed as

$$\pi^{(k+N)} = \pi^{(k)} \boldsymbol{P}_0 \boldsymbol{P}_1 \ldots \boldsymbol{P}_{N-1} \tag{5}$$

Where the $\boldsymbol{P}_k$ matrixes are obtained in a similar manner as previously, but with the use of the structural arrival probabilities. The total loss probability is then given by

$$R(\rho, b) = \frac{\sum_{i=0}^{N-1} R_i E(a_i)}{\sum_{i=0}^{N-1} E(a_i)} \tag{6}$$

where the individual loss ratios may be found by equation (4).

4.2 Multiplexed sources

Two logical multiplexing models will be considered, one in which the periodic GOP property of sources is taken into account, and a simpler one neglecting this property. In the simpler model (denoted 'Frame' in figures), all sources are aligned at frame boundaries and send uncorrelated sequences of frames following the marginal frame size distribution of the original source. The resulting queuing problem is as for the single source, replacing the frame size distribution with the K-fold convolution of the same. The assumption of frame aligned sources greatly simplifies the resulting queuing problem, as the general case implies the treatment of a telescoping recursion relation. In a previous paper (Andreassen, 1995) the consequences of the frame alignment assumption are investigated. For the number of sources and buffer lengths considered here, we have experienced that a model based on frame alignment will upper bound loss probabilities of free running sources by approximately an order of magnitude in the low loss regions ($\sim 10^{-10}$ - $\sim 10^{-6}$), less in high loss regions.

In the other logical multiplexing model that is considered (denoted 'Periodic' in figures) each of the K sources occupies one of N distinct phase values relative to the GOP cycle. Phase constellations change on a single source basis, and the duration of a constellation of phases is assumed to be long enough such that steady state behaviour will dominate over phase change transients. Thus, the treatment of different constellations may be decoupled. A symmetrical phase change model is assumed, such that the probability for a single source to be in a specific state is uniformly distributed over the phases. Each source sends frames according to a periodic pattern of marginal distributions. This model will take into account the periodic correlation property of the sources, but neglecting other correlation properties.

We now consider the distribution of sources in the N discrete phases of the GOP cycle. A specific constellation of sources in the phase space may be identified by a stochastic constellation vector $\boldsymbol{K}$, where the elements K_i denote the number of sources in state i, $i = 0, 1,..., N\text{-}1$. As there are N possible phase values, and $K = \sum_{i=0}^{N-1} K_i$ independent sources, the probability of a specific phase constellation is given by the multinomial distribution:

$$f(\boldsymbol{k}) \equiv Pr(\boldsymbol{K}=\boldsymbol{k}) = \frac{K!}{\prod_{i=0}^{N-1} k_i!} \prod_{i=0}^{N-1} p_i^{k_i} \tag{7}$$

Where $p_i = Pr$(A single source being in phase state i)

With the uniformity assumptions made above, we have that $p_i = 1/N$. As arrival rates do not change between constellations, a mean cell loss probability may in principle be found by taking the expectation of the loss function $R(\rho, b, \boldsymbol{k})$ under the multivariate density $f(\boldsymbol{k})$. In practise, such computations will not be possible, and further simplifications are necessary. To this end, the concept of nearly equivalent phase constellations will be defined.

This is a simplification based on the intuitive idea that the number of in-phase sources in a constellation is more important than the actual phases of the constituent sources in the constellation. Thus, we approximate $R(\rho, b, \boldsymbol{k}) \approx R(\rho, b, perm(\boldsymbol{k}))$, where $perm(\boldsymbol{k})$ denotes an arbitrary permutation of the vector $\boldsymbol{k}$, and note that $f(\boldsymbol{k}) = f(perm(\boldsymbol{k}))$. A set of all vectors that can be permuted into each other form a permutation group A, formally defined as

$A = \{\boldsymbol{k}, \boldsymbol{j} | \boldsymbol{k}, \boldsymbol{j} \in A \Leftrightarrow \boldsymbol{k} = perm(\boldsymbol{j})\}$. All different permutation groups given the GOP length N and the number of sources K make up the set E, which is the set to be traversed using the above simplification. To clarify the above concepts, consider the simple example of $K = 3$ sources with $N = 2$ different phases. Possible constellation vectors are (3,0), (2,1), (1,2) and (0,3), so $A_1 = \{(3,0), (0,3)\}$, $A_2 = \{(2,1), (1,2)\}$, $E = \{A_1, A_2\}$.

To determine the number of members in each permutation group, let the constraint on the vectors be that $\sum_{i=0}^{N-1} k_i = K$, which may be written as $\sum_{i=0}^{K} i \cdot m_i = K$, where an order summation is performed instead of the element summation. The order corresponds to a number of in-phase sources, whereas the factors m_i denote the multiplicity of the ith order in the summation, i.e. how many times a number of in-phase sources are repeated in the constellation vector. An ordered set of m_i-values will uniquely identify a permutation group and vice versa. The m_i factors may take values in the range 0,1,..., K, with the restriction that $\sum_{i=0}^{K} m_i = N$. It may then be seen that the sought number can be expressed by the multinomial weight in the m_i's. To compare with throwing dice, K would be the number of faces of the die, m_i would be the number of dice showing face value i, and N would be the number of dice. The complete expression for the loss probability will be:

$$\hat{R}(\rho, b, K, N) = \sum_{A \in E} \left(\frac{N!}{\prod_{i=0}^{K} m_i!} f(\boldsymbol{k}_A) R(\rho, b, \boldsymbol{k}_A) \right) \tag{8}$$

The composite source used in the above computation may generally be expressed by the N different marginal densities in the MPEG cycle. As we do not want to use any information from the constellation other than whether sources are in phase or not, a natural approach will be to use only two different densities for frame types, one for I-frames and another for P/B-frames. For a given constellation this will give:

$$\hat{A}_i = A_I^{(k_i)} * A_{PB}^{(K-k_i)} \qquad i = 0, 1, \ldots, N-1 \tag{9}$$

Where $A^{(k)}$ denotes the k-fold convolution of the A distribution. The above distributions may then be used in the relations (5) and (6) giving solutions for the periodic group. For each permutation group upper and lower bound on loss probabilities may be computed, assuming that the two cases will occur for the respective maximum concentration and maximum dispersion of the I-frames of the constellations in the permutation group over the GOP period.

To complete the analysis, it will be necessary to find the members of the set E. The problem can be described as the number of ways to form the sum K out of N natural numbers without using more than one member of each permutation group. This problem lends itself to solution by a recursive algorithm, noting that forming subconstellations in a constellation may be done in an invariant manner.

4.3 Simulation model

To compare model results with those from real sources, a trace driven simulator is employed. As the frame alignment assumption is used to get comparable results, the recursion relation

(1) applies in the calculation of buffer occupancy and cell losses. The simulator uses the same algorithm as outlined in the previous section. This will allow estimates of low loss probabilities as compared to a more direct approach. To determine confidence, a loss estimator based on equation (8) is used: Each estimate is based on a weighted mean of simulation results from all permutation groups, where constellations inside a permutation group are chosen at random. To avoid correlation between sources, each source is uniformly offset over the sequence while preserving group constellation.

5 RESULTS AND DISCUSSION

In the following, some results from analytical and simulation models will be compared. All figures display the behaviour of 12 sources, a number which is chosen so as to give a nontrivial situation, tractable computations and simulations, and not to strain the frame alignment assumption too hard. For each source, three scenarios are displayed, each comprising buffer lengths of 5, 10 and 15 times the expected frame lengths of the sources. For the simulation curves, errorbars indicate the 95% confidence interval obtained by performing 100 independent simulations. For the periodic model, an errorbar indicates the upper and lower bound on losses, the line being the mean value of these.

To evaluate results, the following properties of the different models should be kept in mind: The 'Frame' source model is memoryless at the frame level, the 'Periodic' model takes into account the periodic property of the sources, but is otherwise memoryless. The 'Simulation' model contains the naturally occuring correlation properties of the sources, and accounts also for the way different phase constellations are formed.

Results from the Bond, Simpsons and Lambs sources are displayed in Figure 4. In the low loss region, the simulation and model curves fit reasonably well for all buffer lengths displayed. It seems that loss in this region is not related to correlation properties above the frame level of the MPEG sources. Rather the losses are due to the rare situations in which several sources send large frames simultaneously, the scenario which is explicitly modelled both by the constellation simulation model and the Markov models.

The Bond source features a negative-correlation property, easily discernible in the autocorrelation and time variance graphs of Figures 1 and 2. In the latter figure, blocks of sizes two and three frames have a low variance in their sample means as the frames constituting the blocks are negatively correlated. In the loss contour diagram, this property will manifest itself by showing smaller losses for the original source than for the independent source in the upper left regions of that diagram. A similar effect is seen in the case of multiplexed sources. At buffer lengths of 347 and 695 (i.e. five and ten frame lengths), loss probabilities of the real source are less or equals those of the models. Only in the case of buffer length corresponding to 15 frame lengths (b=1043), will the effect of correlations beyond the frame level dominate behaviour in the high loss region, giving higher loss than the Markov models.

The Lambs source is the one most precisely modelled by Markov models in terms of relative buffer lengths. This property is not easily deducible from the autocorrelation and time-variance plots, as these diagrams seem very similar for the Simpsons and the Lambs sources at the short ranges. Considering however the loss contour diagrams, a closer match between an independent and the original source for the Lambs sequence can be observed.

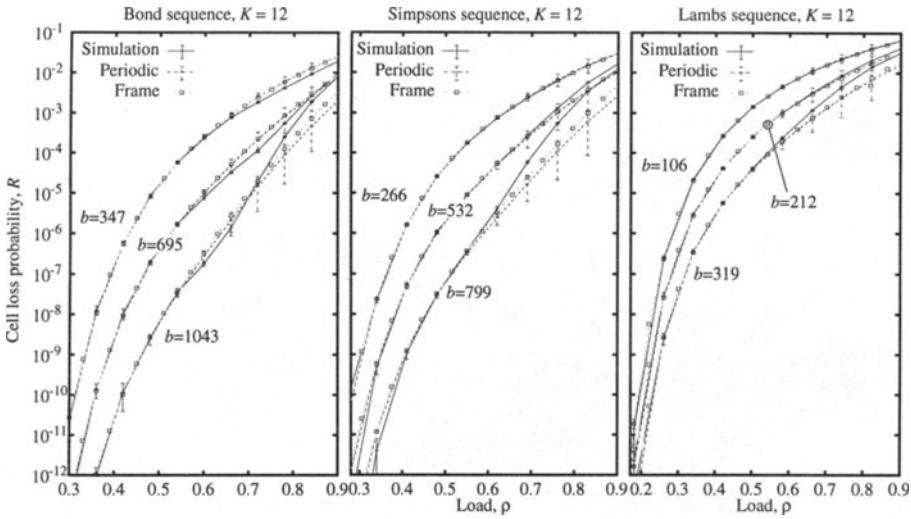

Figure 4 Multiplexer loss probabilities R as function of load ρ for various buffer sizes b

For a number of 12 sources and buffer lengths considered, the periodic property of the sources does not dominate multiplexer behaviour. For all sequences, the simple 'Frame' model gives results well within the uncertainty bounds of the 'Periodic' model. It is our experience that difference in results between these models may be large when only a few sources are multiplexed. This effect seems to diminish as the number of sources increases.

6 CONCLUSIONS AND FURTHER WORK

We have in this paper investigated correlation properties of MPEG VBR video sources, and what impact these properties will have on cell loss probabilities of multiplexed sources. Such sources exhibit properties at short time scales which in some respects resemble those of uncorrelated sources, a phenomenon which is found consistent with the nature of the MPEG coding algorithm. At larger time scales, the sources were however seen to exhibit strong autocorrelation properties. By employing Markov models utilizing only certain periodic properties of the sources, an investigation of which properties would dominate with respect to loss probability of multiplexed frame aligned sources was made possible. It was found that for a number of 12 sources, the short term properties were dominant for buffer lengths up to about 10 times the expected frame lengths of the video sources considered. In low loss regions, no effect of longer term correlation properties was seen at buffer lengths up to 15 times the mean frame lengths, with the same number of sources. Concluding that the longer range properties of the sources are not important in the above considered regions, may however be a hasty conclusion. The cell loss probability measure, even though widely applied, is not the best performance measure of the multiplexer. In order to determine the effect on services, the time distribution of cell losses will be significant, and correlation properties may certainly influence these distributions. An investigation of these questions will be an interesting topic for further research.

Acknowledgements

The author would like to thank Oliver Rose at the University of Würzburg for making the video traces available. Several useful discussions with professors Peder J. Emstad and Tore Riksaasen at the Norwegian University of Science and Technology is also gratefully acknowledged. The traces are available at: ftp-info3.informatik.uni-wuerzburg.de:/pub/MPEG

7 REFERENCES

Andreassen, R., Emstad, P., Riksaasen, T. (1995) Cell losses of multiplexed VBR MPEG sources in an ATM multiplexer, in *Proc. 12th Nordic Teletraffic Seminar*, 83-95

Beran, J., Sherman, R., Taqqy, M. S., Willinger, W. (1995) Long Range Dependence in Variable-Bit-Rate Video Traffic, *IEEE Transactions on Networking*, **43**, (2/3/4), 1566-1579

Garret, M.W. (1993) *Contributions towards real-time services on packet switched networks*, Ph.D. Thesis Columbia University.

Garret, M.W., Willinger, W. (1994) Analysis, Modelling and Generation of Self-Similar VBR Video Traffic, in *Proc. ACM Sigcomm 94*, ACM press, 269-280.

Huang, C., Devetsikiotis, M., Lambadaris, I., Kaye, A. R. (1995) Modeling and Simulation of Self-Similar Variable Bit Rate Compressed Video: A Unified Approach, in *Proc. ACM Sigcomm 95*, 114-125

Hübner, F. (1994) Dimensioning of a Peak cell Rate Monitor Algorithm Using Discrete-Time Analysis, in *Proc. ITC 14*, 1415-1424

Landry, R., Stavrakakis, I. (1994) Non-Deterministic Periodic Packet Streams and Their Impact on Finite-Capacity Multiplexer, in *Proc. IEEE Infocom 94*, 224-231

Lin, A. Y. M., Silvester, J. A. (1990) Queueing analysis of an ATM Switch with Multichannel Transmission Groups, *Performance Evaluation Review*, **18**, (1), 96-105

Lucantoni, D.M., Neuts, M. F., Reibman, A. R. (1994) Methods for Performance evaluation of VBR Video Traffic Models, *IEEE/ACM Transactions on Networking*, **2**, (2), 176-180

Reininger, D., Melamed, B., Raychaudhuri, D. (1994) Variable Bit Rate MPEG Video: Characteristics, Modeling and Multiplexing, in *Proc. ITC 14*, 295-306

Rose, O. (1995) *Statistical Properties of MPEG video traffic and their impact on traffic modelling in ATM systems*, (University of Würzburg Institute of Computer Science Research Report Series), Report No. 101

Østerbø, O. (1995) *Some Important Queueing Models for ATM*, (TF Rapport 22/95), Kjeller, Norwegian Telecom Research