ISO/IEC JTC1/SC29/WG11
MPEG 92/
Document AVC-296
July 1992

CCITT SGXV
Working Party XV/1
Experts Group for ATM Video Coding

SOURCE: Australia
TITLE: Cell Loss Characteristics for Statistically Multiplexed Video Sources
PURPOSE: Information and Proposal

Abstract

An understanding of cell loss characteristics is essential to design effective video codecs for ATM networks. Cell loss models currently used within the CCITT SGXV Experts Group consider only loss characteristics at the cell level. This implies that longer term characteristics of loss, which result from longer term correlations in the video source rate, are ignored. This document considers longer term characteristics which arise when streams from variable rate video codecs are statistically multiplexed. From the results it is clear that these will have a significant impact on codec performance.

1. Introduction

The Experts Group is considering the use of Variable Bit Rate (VBR) video coding to provide higher quality than conventional constant bit rate (CBR) coded video for a given equivalent network usage. However, to realise this advantage it is necessary to make use of statistical multiplexing of the variable rate sources to improve network utilisation. When statistical multiplexing is employed, cell loss due to switch buffer overflow results. The average cell loss ratio (CLR) can be kept to arbitrarily low levels by reducing network utilisation, however, to obtain reasonable network utilisation and hence provide some gain from statistical multiplexing, cell loss must be non-zero.

In this document we use a simplified model of the video sources and the multiplexing process which makes it possible to either calculate or determine by simple experiments various cell loss characteristics. We present a number of useful measures including the average CLR, the average congestion time (the length of time over which the offered rate is greater than the network capacity and hence the cell loss ratio is high), the congestion time distribution and the average CLR during congestion. These measures are extremely useful in the design of cell loss resilient variable rate video codecs. Numerical examples calculated for statistically multiplexed video sources are provided.

2. Results

A detailed discussion of the results, the assumptions which were made and the experiments performed is given in the Appendix. In summary the main results are:

- Congestion which leads to high cell loss occurs in bursts.
- Bursts of high cell loss can last from several frames up to tens of frames, depending on the correlation in the source statistics and the average CLR on the network.
- The duration of intervals with no cell loss depend almost in inverse proportion to the average cell loss on the network. The average durations range from about 1 hour down to about a second.
- During congestion, cell loss is very high and the level of the loss is only mildly dependent on the average CLR on the network. Instantaneous CLRs of 0.01 are not uncommon.

3. Conclusions

The model of cell loss currently being used within the Experts group [8] does not represent the long term burst characteristics of cell loss which will result when VBR video sources are statistically multiplexed. The analysis and experiments presented in this document indicate that, during intervals of congestion which can last over several frames, cell loss can be of the order of 0.01. We believe that VBR codecs should be tested under these cell loss conditions to ensure that adequate end to end performance is achieved.

Appendix. Detailed Discussion

A.1 Source and Network Models

To investigate the cell loss characteristics, we assume a simplified model of both the sources and the multiplexing arrangement in the network.

A.1.1 Source Models

In [3] it has been argued using experimental evidence from variable rate video codecs that the rate distribution, in bits per frame, is approximately bell shaped and the auto-covariance of the bit-rate is approximately exponential. This suggests that the statistical characteristics of bit-rate generation from video codecs can be modelled as a discrete time first order autoregressive process (AR(1)). The auto-covariance of an AR(1) process is exponential and can be written as $C_k = a^k \cdot C_0$, where C_0 is the variance of rate and a is the model parameter. If we have N independent auto-regressive sources with identical parameters (mean=μ_0, standard deviation=σ_0 and a) then the model for the multiplexed sources is also auto-regressive of order one with parameter a, mean given by $N \cdot \mu_0$ and standard deviation given by $\sqrt{N} \cdot \sigma_0$.

To demonstrate the accuracy of the AR(1) model we have examined the output of two different variable rate video codecs. The first is the Alcatel real-time VBR codec, for which an extensive set of rate data is available [4]. The second is based on the CCITT Rec. H.261 standard, applied to video-conference source material[1][5]. Figure 1 shows the auto-correlation function for these codecs. The first three plots are from the Alcatel codec while the second video phone plot is from the H.261 based codec. Also included are exact exponential auto-correlation functions which apply to the AR(1) model with different parameter values. From these results it can be seen that the auto-correlation functions are not precisely modelled by a single exponential, except over short segments. However, the exponential model with parameter of 0.975 is approximately correct for the video-phone data, while values of about 0.98 and 0.99 are appropriate for the other two data sets[2].

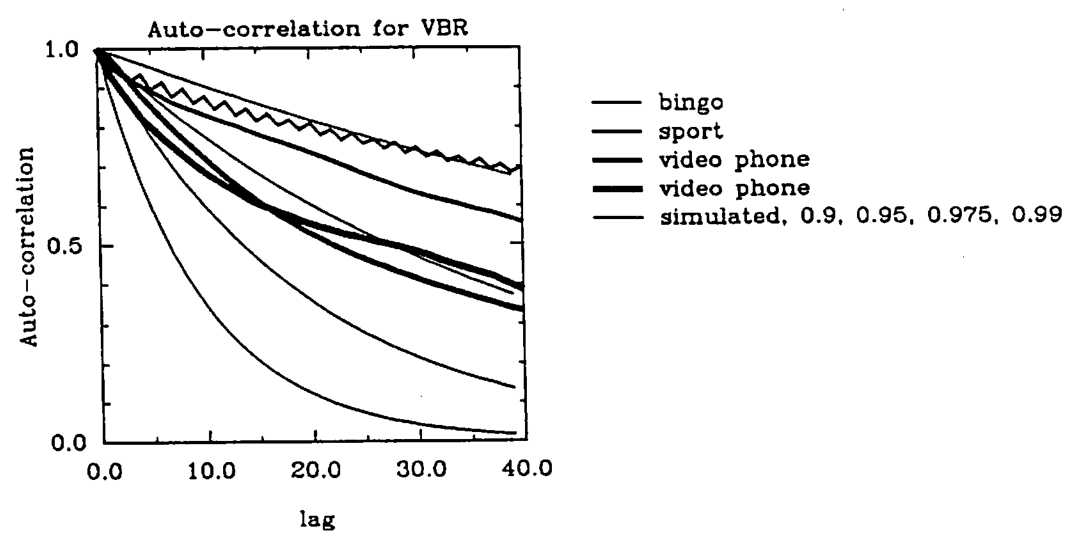

Figure 1. Auto-correlation function for VBR sources, along with the exponential model auto-correlation for various values of the parameter a.

---

[1]The variable rate output was obtained by keeping the quantizer step-size within the codec fixed.

[2]There is some doubt about the validity of the Alcatel data since it contains a number of unexpected spikes which suggest coder malfunctions. We have trimmed the data to remove these spikes in the current study.

A.1.2 Network/Switch Model

In developing models of cell loss processes, it is useful to distinguish between *cell scale* congestion and *burst scale* congestion [1,2]. Cell scale congestion is a result of the simultaneous arrival of cells from different sources over a short period of time. Burst scale congestion occurs when the sum of the rates of the sources is greater than the multiplexer capacity over a period much longer than the inter-cell arrival time. With relatively small buffers in the multiplexer cell loss due to cell scale congestion can be eliminated at the penalty of a non-zero, but acceptable delay [6]. For video sources the output rate can be correlated over many frames. Therefore it is impractical to use multiplex buffers to absorb the burst scale congestion, both due to the size of the required buffer and the associated delay incurred [7].

In this document we assume that multiplex buffers have been dimensioned to absorb cell scale congestion and consider purely burst scale congestion. This simplifies the analysis considerably since it allows us to use a continuous flow model of the sources and the multiplexer and we can assume that the multiplexer is memoryless.

Under the assumption of a memoryless multiplexer, the cell loss rate can be approximated as a simple function of the total rate of the multiplex of sources. Under ideal conditions this loss function, $L(R)$, would be zero for total rates below the network capacity and then increase linearly above the capacity with a slope of one [AVC-75], i.e.

$$L(R) = \begin{cases} 0, & R \le C \\ R - C, & R > C \end{cases}$$

where R is the total rate of the multiplexed sources and C is the network capacity. The instantaneous CLR is given by,

$$CLR_i = \frac{L(R)}{R}$$

The average CLR can be calculated as the average cell loss rate divided by the average cell transmission rate. The expression for average CLR in terms of $L(R)$, the probability density function of the multiplexed sources $p(R)$, the average rate of the multiplexed sources, μ_T, and the variance of the multiplexed sources, σ_T^2, is given in [AVC-75, AVC-106R] for both the general case and the case where $p(R)$ is a Gaussian density.

It should be noted that both the instantaneous and average CLRs strictly represent the CLR present in the multiplexer, and not the CLR experienced by individual sources. We make the assumption here that these two figures are identical. Analysis described in [1] supports this assumption, particularly in the case of identical sources.

Mean rate (Mb/s)	$CLR=10^{-2}$	$CLR=10^{-3}$	$CLR=10^{-4}$	$CLR=10^{-6}$
0.384	381	359	344	323
1.0	140	129	120	109
4.0	31	26	23	19

Table 1. Numbers of sources which can be accommodated on a 150 Mb/s network, assuming that the source standard deviation is equal to the mean[3].

Table 1 shows the number of VBR sources which can be accommodated for different CLRs and different source mean rate assumptions, assuming Gaussian distributed sources. It is worth noting

[3]Equal source mean and variance gives a variability which is typical of video material. It is approximately equivalent to allowing a peak to mean ratio of 3.

that the penalty, in terms of number of sources, incurred by going from a CLR of 10^{-4} to a CLR of 10^{-6} varies from 17.4 % to 6 % depending on the source mean rate assumptions. Hence, particularly for large numbers of sources, there is good reason to suggest that one can operate the network at reasonably low average CLRs without sacrificing significant network utilisation.

A.2. Bursty Nature of Congestion

It should be fairly clear that the model assumptions which have been made lead to the conclusion that cell loss occurs in bursts rather than being uniformly distributed over time. This is because the rate of the VBR sources is correlated over time and therefore when it is high it will tend to remain high for some period. The burstiness of cell loss is directly related to this correlation in the codec output rate.

In developing video coding techniques which can operate in the presence of cell loss it is important to understand both the frequency and the length of bursts of high cell loss, as well as the actual level of cell loss which occurs during bursts. In the following analysis we consider these aspects separately. This in itself makes the analysis an approximation since one should be considering the joint distributions of burst lengths and cell loss levels, however this separation gives a useful first approximation.

A.2.1 Burst length distribution

The burst length distribution is relatively straight-forward to determine analytically, unfortunately it is difficult to evaluate the multiple integral which results and so it is not of much use practically. Roberts, et al [2] have suggested approximating congestion as a two state Markov process, as shown in figure 2.

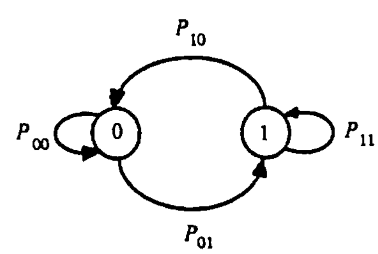

Figure 2. Two State Markov approximation to the congestion process. State 0 represents uncongested and state 1 represents congested.

The parameters of this process can be estimated directly from the assumed models for the source. We have

$$P_{11} = \frac{\int_{z=\gamma}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} \left\{ 1 - \aleph\left(\frac{\gamma - a \cdot z}{\sqrt{1-a^2}} \right) \right\} dz}{1 - \aleph(\gamma)}$$

where, $\aleph(.)$ is the standard normal distribution. Similarly,

$$P_{00} = \frac{\int_{z=-\infty}^{\gamma} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} \aleph\left(\frac{\gamma - a \cdot z}{\sqrt{1-a^2}} \right) dz}{\aleph(\gamma)}$$

Of course, $P_{10} = 1 - P_{11}$ and $P_{01} = 1 - P_{00}$. For this model the probability of exactly N consecutive congested frames, given that congestion has commenced, is given by $P(N) = (1 - P_{11})P_{11}^{N-1}$ and therefore

$$P(Burst \geq N \,/\, Burst) = 1 - (1 - P_{11}) \sum_{i=1}^{N-1} P_{11}^{i-1}$$

An equally straight-forward approach is to use experimental techniques with simulated data from an auto-regressive source. The simulation of the cell loss is trivial and statistics can be collected rather easily. Figure 3 provides the measured conditional distribution for a mean rate of 1 Mb/s, and for different values of the parameter a. Also shown are the distributions obtained from the Markov approximation. The approximation is reasonable accurate except when a approaches 1. As expected the probability of getting long burst lengths increases with increasing a. The length which the congestion interval exceeds with a probability of 0.1 is about 3 frames for a=0.9, 5 frames for a=0.95 and about 12 frames for a=0.99. In the case of a=0.99 even congestion intervals of 25 frames (i.e. about 1 second) are quite likely. Space precludes us from including the results for other operating conditions. However, the general trends are that the distribution widens with increasing average CLR, and increasing number of sources, though it is not very sensitive to either.

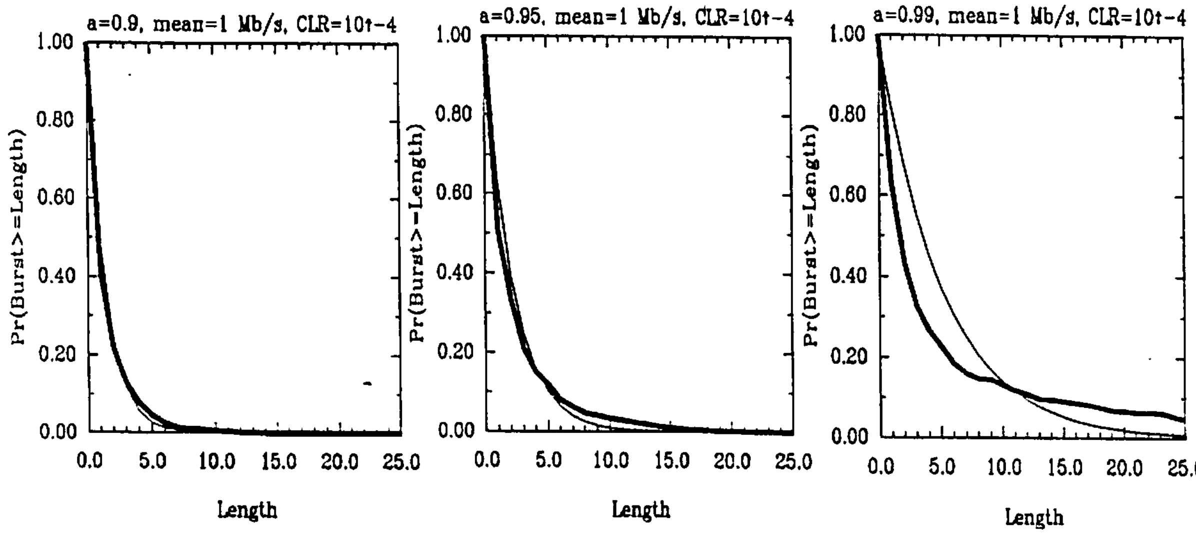

Figure 3. Burst length distribution as a function of the AR(1) parameter a. The thin lines represent the Markov approximation.

A.2.2 Average Congestion Length

For the Markov approximation the average congestion length is simply $1/(1-P_{11})$. Figure 4 plots the measured mean congestion length for a range of different source bit-rates, average CLRs and correlation parameters. Also plotted, for a parameter value of 0.95, are the estimated values using the 2-state Markov approximation. The results are close enough to make it difficult to separate them from the rest of the plots. We have found that the Markov approximation is very accurate over the complete range of parameter values investigated.

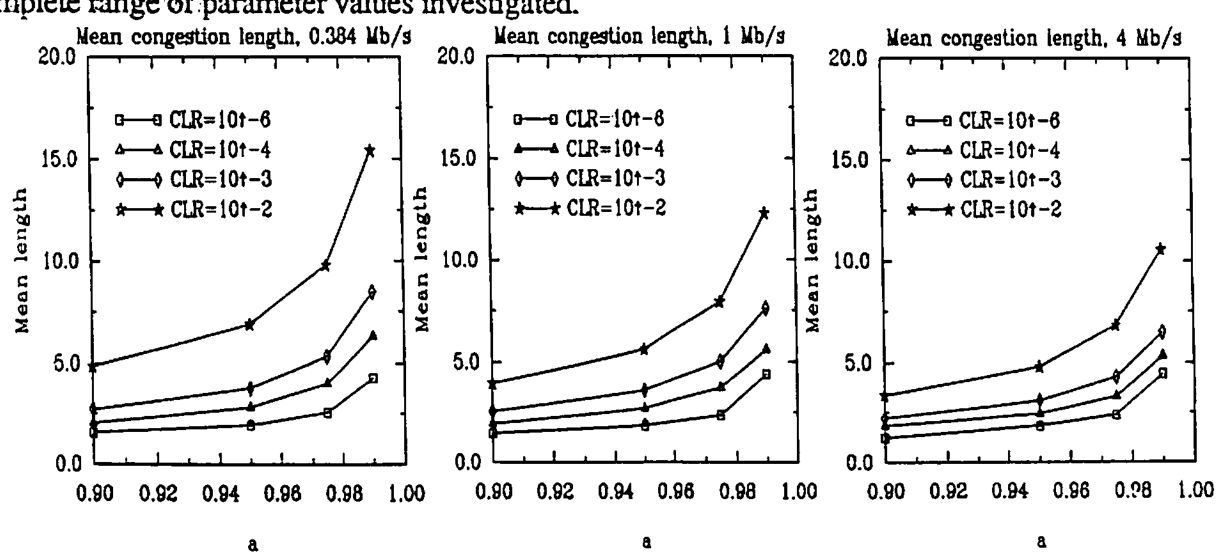

Figure 4. Average congested intervals for a range of conditions. The results from the 2-state Markov approximation are plotted for a=0.95 as solid circles.

As expected there is a clear dependency of average congestion interval on the correlation parameter. There is also a dependence on the average CLR and smaller average CLRs lead to shorter average congestion intervals. The dependency on average bit-rate (or number of sources multiplexed) is less significant though there is a clear trend showing more multiplexed sources leads to longer average congestion intervals.

Also of interest are average lengths of uncongested intervals. Table 2 gives some typical figures. Our results have indicated a trend for longer uncongested intervals when a larger number of sources are multiplexed, the average CLR is lower and when the source bit-rate is more correlated. The most significant variation is as a function of average CLR where the lengths vary roughly inversely (at least to the same order of magnitude). That is for an average CLR=10^{-6} average uncongested intervals are about two orders of magnitude greater than those for an average CLR=10^{-4}.

a	CLR=10^{-2}	CLR=10^{-3}	CLR=10^{-4}	CLR=10^{-6}
0.9	15	76	627	38910
0.95	22	108	862	48780
0.975	32	152	1219	61350
0.99	50	239	2109	95230

Table 2. Average uncongested intervals (number of frames) for different average CLRs and different correlation parameters. These figures are for a source mean rate of 1 Mb/s.

A.3. Cell Loss During Congestion

The instantaneous CLR distribution *given congestion* is determined by the rate distribution of the multiplex of sources and the capacity of the network. That is

$$P(CLR_i > X / R > C) = P(R > Z / R > C) = \begin{cases} 1, & C = Z \\ \dfrac{P(R > Z)}{P(R > C)}, & C < Z \end{cases}$$

where $Z=C/(1-X)$. Figure 5 illustrates the instantaneous CLR distribution during congestion, for a mean rate of 1 Mb/s and three different average CLRs. Interestingly the dependency on average CLR is minor. Similarly the distribution shows little variation for changes in the number of sources, with only a slight trend to widen with reducing numbers of sources. This result confirms the observation made by Norros [1] who calculated the average CLR during congestion and observed it was fairly independent of average CLR. The results do show that the CLR during congestion can be quite high, regardless of the average CLR, with values of 0.01 being relatively common.

REFERENCES

1. I. Norros and J.T. Virtamo, "Who Loses Cells in the Case of Burst Scale Congestion", ITC-13, 1991, pp. 829-833.
2. J.W. Roberts, J. Guibert and A. Simonian, "Network Performance Considerations in the Design of a VBR Codec", ITC-13, 1991, pp. 77-82.
3. B. Maglaris, et al, "Performance Models of Statistical Multiplexing in Packet Video Communications", IEEE Transactions on Communications, Vol. 37, No. 7, July 1988, pp. 834-844.
4. W. Verbiest and L. Pinnoo, "A Variable Bit Rate Video Codec for Asynchronous Transfer Mode Networks", IEEE Journal Select. Areas in Commun., vol. 7, pp. 761-770, June 1989.
5. W. B. S. Tan, N. Duong and J. Princen, "A Comparison Study of Variable Bit Rate Versus Fixed Rate Video Transmission", Proc. ABSSS-91, Sydney, Australia, July 1991.
6. J. W. Roberts, "Variable-Bit-Rate Traffic Control in B-ISDN", IEEE Communications Magazine, Vol. 29, No. 9, Sept. 1991, pp. 50-56.

7. W. Verbiest and L. Pinnoo, "The Imapct of the ATM Concept on Video Coding", IEEE Journal on Selected Areas on Communications, JSAC-6, No. 9, Dec. 1988, pp. 1623-1633.
8. "Cell loss Experiment Specification", CCITT SGXV Experts Group for ATM Video Coding, document AVC-197, January, 1992.

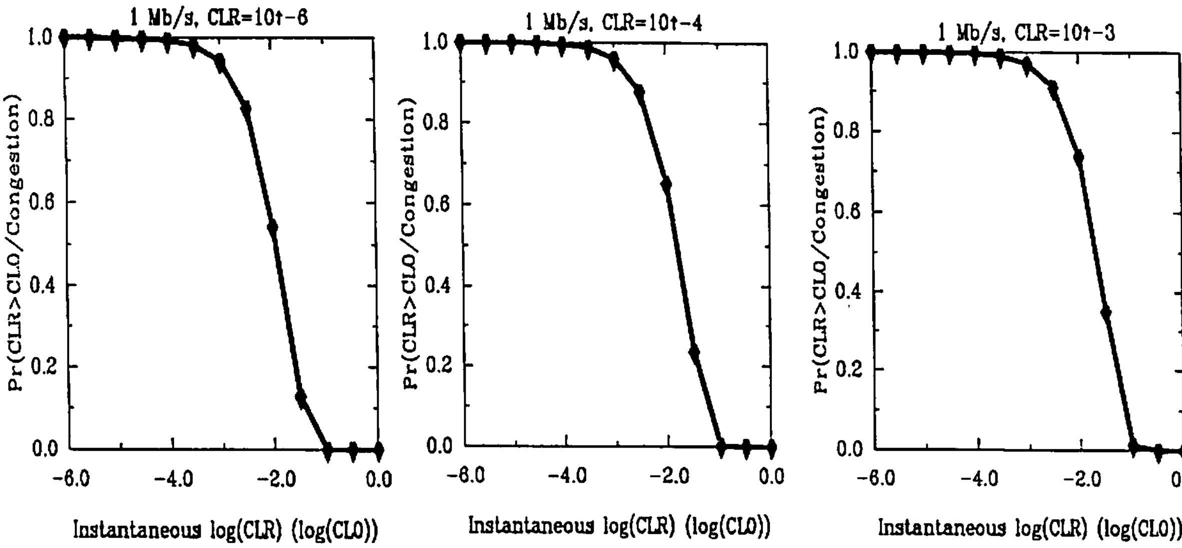

Figure 5. Conditional distribution of the instantaneous CLR for different average CLRs.