

The Correlation Structure for a Class of Scene-Based Video Models and Its Impact on the Dimensioning of Video Buffers

Marwan M. Krunz, *Member, IEEE*, and Arivu M. Ramasamy, *Student Member, IEEE*

Abstract— We analyze the autocorrelation structure for a class of scene-based MPEG video models at the GOP (course grain) and frame (fine grain) levels assuming an arbitrary scene-length distribution. At the GOP level, we establish the relationship between the scene-length statistics and the short-range/long-range dependence (SRD/LRD) of the underlying model. We formally show that when the *intra-scene* dynamics exhibit SRD, the overall model exhibits LRD if and only if the second moment of the scene length is infinite. Our results provide the theoretical foundation for several empirically derived scene-based models. We then study the impact of traffic correlations on the packet loss performance at a video buffer. Two popular families of scene-length distributions are investigated: Pareto and Weibull. In the case of Pareto distributed scene lengths, it is observed that the performance is rather insensitive to changes in the buffer size *even as the video model enters the SRD regime*. For Weibull distributed scene lengths, we observe that for small buffers the loss performance under a frame-level model can be larger than its GOP-level counterpart by orders of magnitude. In this case, the reliance on GOP-level models will result in very optimistic results.

Keywords— Video modeling, MPEG, traffic correlations, buffer design.

I. INTRODUCTION

In this paper, we investigate the correlation structure for a class of scene-based models that characterize variable bit rate (VBR) MPEG-coded video streams. Scene-based models form an important family of video models in which scene dynamics are explicitly incorporated. These models are particularly capable of capturing the multiple-time-scale variations in a VBR video source and, consequently, on providing accurate estimates of the queueing performance at a buffering node [13]. Several of these models have been proposed in the literature. Examples are given in [5], [9], [10], [13], [3], [19], [15], [16], [18] (also, see [11] for a survey of video models). In principle, a scene-based model could incorporate both inter- and intra-scene variations. However, intra-scene variations are often ignored to simplify the construction of the model [10], [5]. Scene-based models often differ in the *persistence* of the autocorrelations

that are manifested by the intra- and/or inter-scene components. Such differences are by no means trivial, as they have significant implications on the design and dimensioning of video buffers.

Most scene-based models have been developed for compression schemes in which all the frames in the video sequence are encoded in a uniform manner using the same mode(s) of compression. This results in *homogeneous* VBR sequences in which the fluctuations are primarily attributed to scene dynamics. In contrast, the MPEG algorithm applies different encoding techniques to produce three types of compressed frames (I , P , and B) that differ significantly in their bit rate characteristics. The complete MPEG sequence is obtained by interleaving frames of different types according to a Groups-of-Pictures (GOP) pattern, which specifies the sequence of P and B frames between two successive I frames. The GOP pattern is applied repeatedly to a video sequence, resulting in *heterogeneous* frame sizes and significant periodicity in the traffic pattern.

An important aspect of a video model is the form of its autocorrelation function (ACF). Traffic correlations, in general, are believed to have profound impact on the queueing performance at a buffering node. Supported by extensive statistical evidence, some researchers have further argued that network traffic, including VBR video, exhibits persistent correlations that can only be captured through LRD models [4], [17], [1], [6]. Other researchers, while acknowledging the presence of LRD in network traffic, argue that for *finite* buffers, long-term correlations have minor impact on the queueing performance [8], [7], [20]. Hence, they argue, the traffic can be sufficiently represented by Markovian models, which are typically easier to analyze than LRD models.

In this paper, we focus our attention on a general class of scene-based video models, which includes SRD and LRD models, and we try to answer some of the important questions related to the impact of correlations on the design of video buffers. The goal of this paper is neither to provide a new video model nor to advocate any of the existing ones. Instead, we aim at studying some generic statistical aspects that pertain to many of these models. As such, fitting of real data, which is an essential aspect of model construction, is

This work was supported by the National Science Foundation under grant ANI-9733143. Part of this paper was presented at the *International Workshop on Packet Video '99*, April 1999, New York. The authors are with the Department of Electrical and Computer Engineering, The University of Arizona, Tucson, Arizona 85721 (e-mail: krunz@ece.arizona.edu; mani@ece.arizona.edu).

not discussed in this paper, as this has been extensively addressed in the literature. To formally assess the impact of traffic correlations on the performance at a video buffer, we first analyze the ACF in a class of scene-based video models. Our analysis is then used in studying the impact of correlations on the dimensioning of video buffers. Figure 1 depicts one scenario for which our analysis is directly applicable (other scenarios exist as well). In here, variable-size video frames are generated on the fly and are fed into a fixed-bandwidth channel. For constant-quality video, the variations in the frame sizes can be quite significant. They can be reduced using rate-controlled encoding at the expense of variable quality. To limit its impact on video quality, rate control is applied in conjunction with sender-based traffic shaping (buffering). Therefore, the encoder still generates variable-size frames (i.e., near-VBR stream) that are fed into the buffer. Depending on the network bandwidth (C), the buffer may occasionally overflow, causing further degradation in the video quality at the receiver. One of our objectives here is to study the impact of traffic correlations on the performance at this buffer.




Fig. 1. Video buffer for smoothing the traffic at the sender.

The contributions of this paper are as follows. First, we derive the GOP-level ACF for a general class of scene-based models with an arbitrary scene-length distribution and frame-size statistics. The only restrictions we impose on this class are that inter-scene and intra-scene variations are mutually independent and that scene lengths constitute an *i.i.d.* process. Such assumptions are satisfied by most existing scene-based models (for which the ACF has not been previously reported). From the derived ACF, we establish the relationship between the scene-length distribution of a model and its SRD/LRD structure. Our results indicate that when the intra-scene dynamics exhibit SRD (as often the case), the overall video model is LRD only if the second moment of the scene length is infinite. Using the explicit relation between the ACF and the scene length distribution, one can determine an appropriate fit for this distribution without the need to directly measure scene lengths.

Based on our generic GOP-level MPEG model, we introduce a frame-level counterpart that incorporates the three types of MPEG frames. We derive the ACF for this frame-level model. Our results indicate that

due to the repetitive application of the GOP pattern, the pseudo-periodic frame-level ACF never drops off to zero. The non-zero-convergence result can be extended to other types of media streams that are interleaved in a deterministic manner (e.g., the interleaving of audio and video packets in MPEG-2).

Lastly, we study the impact of correlations in a scene-based model on the packet loss performance due to buffer overflow at the encoder. Two popular families of scene-length distributions are examined: Pareto and Weibull. In each family, we vary the level of correlations in the model and observe the resulting impact on the packet loss rate. We make several important remarks on this impact as inferred from both GOP- and frame-level models and with the assumption of finite- and infinite-buffer capacities.

II. GOP-LEVEL AUTOCORRELATIONS

In this section, we investigate the ACF for a scene-based video model at the GOP level. While the GOP notion is specific to MPEG video, our “GOP-level” analysis applies, in general, to VBR sequences in which frame sizes are homogeneous, i.e., produced by the same compression approach. For example, it applies to JPEG and H.261 video sequences, among others. Without loss of generality, we present our ideas in the context of MPEG video.

Consider an MPEG-coded video sequence. Let X_n be a random variable (rv) that models the “size” of the n th GOP in this sequence (i.e., the number of bits in that GOP). Let S_i be a discrete rv that models the length of the i th scene (measured in the number of GOPs). We assume that scene lengths are *i.i.d.* with common probability mass function f_s and cumulative distribution function F_s . Let S be a generic rv that describes the length of an arbitrary scene. Intuitively, GOPs that belong to the same scene are relatively close in size, while GOPs belonging to different scenes may have significantly different sizes. Accordingly, we model X_n by the sum of two random components:

$$X_n \stackrel{\text{a.s.}}{=} Y_n + Z_n \quad (1)$$

where “a.s.” refers to equality in the *almost surely* sense. The rv Y_n accounts for the average impact of scene dynamics on the bit rate. In essence, it represents the average GOP size within a given scene. If two GOPs i and j belong to the same scene, then $Y_i \stackrel{\text{a.s.}}{=} Y_j$; otherwise, Y_i and Y_j are *i.i.d.* The rv Z_n represents the difference between the size of the n th GOP and the mean GOP size in the underlying scene. By construction, $E[Z_n] = 0$. We assume that Y_n and Z_n are mutually independent.

The random processes $\{Y_n : n = 1, 2, \dots\}$ and $\{Z_n : n = 1, 2, \dots\}$ constitute two sequences of autocorrelated and identically distributed rvs. We as-

sume that both processes are second-order stationary, and we denote their corresponding ACFs at lag k by $\rho_Y(k)$ and $\rho_Z(k)$, respectively. Furthermore, we let σ_Y^2 and σ_Z^2 be the variances of Y_1 and Z_1 , respectively. The above formulation encompasses several scene-based models, including the ones in [5], [9], [10], [15], [16], [18].

Now consider the random process $\{X_n : n = 1, 2, \dots\}$. Its ACF at lag k is given by

$$\begin{aligned} \rho_X(k) &\triangleq \frac{E[(X_n - m)(X_{n+k} - m)]}{\sigma_X^2} \\ &= \frac{E[Y_n Y_{n+k}] + \sigma_Z^2 \rho_Z(k) - m^2}{\sigma_Y^2 + \sigma_Z^2} \end{aligned} \quad (2)$$

where $m \triangleq E[X_1] = E[Y_1]$. Consider the term $E[Y_n Y_{n+k}]$ for $k = 1, 2, \dots$, and an arbitrary n . The relationship between Y_n and Y_{n+k} depends on whether GOPs n and $n+k$ belong to the same scene. Let \hat{S} be the forward recurrence time that is associated with the scene length S . The pmf of \hat{S} is given by:

$$f_{\hat{S}}(i) \triangleq \Pr[\hat{S} = i] = \frac{\Pr[S \geq i]}{E[S]}, \quad i = 1, 2, \dots \quad (3)$$

Since n is chosen arbitrarily, the two GOPs belong to the same scene if $\hat{S} > k$. Otherwise, they belong to different scenes (and are, thus, independent). Consequently,

$$\begin{aligned} E[Y_n Y_{n+k}] &= \sum_{j=1}^{\infty} E[Y_n Y_{n+k} | \hat{S} = j] \cdot \Pr[\hat{S} = j] \\ &= m^2 \sum_{j=1}^k f_{\hat{S}}(j) + E[Y_1^2] \sum_{j=k+1}^{\infty} f_{\hat{S}}(j) \\ &= m^2 F_{\hat{S}}(k) + E[Y_1^2] (1 - F_{\hat{S}}(k)) \end{aligned} \quad (4)$$

where $F_{\hat{S}}$ is the CDF of \hat{S} . Accordingly, (2) can be written as

$$\rho_X(k) = \frac{\sigma_Y^2 [1 - F_{\hat{S}}(k)] + \sigma_Z^2 \rho_Z(k)}{\sigma_Y^2 + \sigma_Z^2} \quad (5)$$

In the absence of the noise process $\{Z_n : n = 1, 2, \dots\}$, (5) reduces to $\rho_X(k) = \Pr[\hat{S} > k]$, i.e., the ACF is simply given by the complementary distribution of \hat{S} .

Equation (5) can be used to construct a simple test for the LRD/SRD of a video model with a given scene-length distribution. Recall that a process exhibits LRD behavior if its ACF has an infinite sum. Taking the sum of $\rho_X(k)$ from $k = 0$ to ∞ , we have

$$\sum_{k=0}^{\infty} \rho_X(k) = \frac{\sigma_Y^2 E[\hat{S}] + \sigma_Z^2 \sum_{k=0}^{\infty} \rho_Z(k)}{(\sigma_Y^2 + \sigma_Z^2)} \quad (6)$$

It is easy to show that

$$E[\hat{S}] = \frac{1}{2} + \frac{E[S^2]}{2E[S]} \quad (7)$$

From (6) and (7), we arrive at the following result:

Proposition 1: The video model $\{X_n : n = 1, 2, \dots\}$ is long-range dependent if and only if at least one of the following conditions is satisfied:

1. The second moment of the scene length is infinite.
2. The noise process $\{Z_n : n = 1, 2, \dots\}$ is LRD.

Examples

I. Pareto Distribution

Some studies have reported the appropriateness of the Pareto distribution for modeling the scene duration [9], [13]. The complementary form of this distribution is given by $\Pr[S > k] = \left(\frac{w}{k}\right)^\alpha$ for $k \geq w$ where w and α are two positive parameters. Assuming that the noise process is SRD, then for $1 < \alpha < 2$, $E[S^2]$ is infinite and the model is LRD. A similar result has been provided for the superposition of ON/OFF sources in which the ON periods of one or more sources are Pareto distributed with $1 < \alpha < 2$ [2].

II. Frater's Scene-Length Distribution

In [5], Frater et al. introduced a model for JPEG video in which the scene duration has the following distribution:

$$f_s(k) \triangleq \Pr[S = k] = \frac{a}{k^n + b^2}, \quad k = 1, 2, \dots$$

where a , b , and n are three positive constants. The noise process was ignored. It is straightforward to show that $E[S^2]$ is finite for $n > 3$, and is infinite otherwise. In other words, Frater's video model is SRD if and only if $n > 3$. Two video sequences were examined in [5]: *Star Wars* and *Film*. Their corresponding n values were determined to be 2 and 3.8, respectively, implying that Frater's *model* for the first sequence exhibits LRD!

In typical scene-based modeling studies, the scene-length distribution is obtained by directly fitting the empirical scene lengths. The ACF in this case is obtained empirically using synthetic sequences. Due to the small number of long scenes (typically, in the order of few tens), accurate fitting of the tail of the empirical scene-length distribution is not feasible, despite the importance of this tail in determining the SRD/LRD structure of the model. We remedy this issue by providing an alternative modeling approach for the scene-length distribution. Ignoring the noise process, we have $\rho_X(k) = \Pr[\hat{S} > k]$. Thus, for $j = 1, 2, \dots$, we have

$$\Pr[\hat{S} = j] = \rho_X(j-1) - \rho_X(j) = \frac{\Pr[S \geq j]}{E[S]}, \quad (8)$$

from which we obtain the scene-length pmf in terms of the ACF:

$$\Pr[S = j] = \Pr[S \geq j] - \Pr[S \geq j + 1] = E[S] [\rho_X(j - 1) - 2\rho_X(j) + \rho_X(j + 1)] \quad (9)$$

To obtain $E[S]$ we set $j = 1$ in (8), which results in

$$E[S] = \frac{\Pr[S \geq 1]}{\rho_X(0) - \rho_X(1)} = \frac{1}{1 - \rho_X(1)} \quad (10)$$

Accurate modeling of the scene length can proceed by first obtaining an adequate fit for the empirical ACF (which can be done relatively with high accuracy), and then using this fit to derive the corresponding scene-length distribution based on (9).

III. FRAME-LEVEL AUTOCORRELATIONS

GOP-level modeling is sufficient for evaluating the queueing performance when the buffer size is relatively large. In this case, the drastic differences between the three types of MPEG frames are absorbed by the buffer. However, when the buffer is small (e.g., drains in less than a GOP time), these differences will have profound impact on the performance, and a frame-level model is needed to study this impact. In this section, we extend our generic GOP-level model to characterize the frame-level variations, and we derive the resulting ACF.

Consider the process $\{X_n : n = 1, 2, \dots\}$ that represents the GOP sequence. Let f_X be its marginal distribution. A GOP pattern is characterized by two parameters: the *I*-to-*I* frame distance (N) and the *I*-to-*P* frame distance (M). Not all MPEG sequences involve a repetitive GOP pattern. In fact, some encoders allow a new GOP to start before the completion of the previous one, typically in response to a large frame (i.e., the start of a high-action scene). However, for tractability purposes, we restrict our work to MPEG sequences that conform to repetitive GOPs, and many sequences do so in practice.

Denote the size of the k th frame in the MPEG sequence by U_k . Suppose that this frame belongs to the r th GOP. If the MPEG sequence starts with a complete GOP (i.e., the first frame is *I*), then $r = \lceil k/N \rceil$. However, this makes the process $\{U_n : n = 1, 2, \dots\}$ nonstationary, precluding any analysis of the correlation structure. Instead, we will allow the first GOP to be incomplete by randomly selecting the first frame in the sequence from any location in the GOP pattern, and continuing thereafter according to that pattern. This will have no effect on the long-term behavior of the model, but will ensure its stationarity. Accordingly, $\lceil k/N \rceil \leq r \leq \lceil k/N \rceil + 1$.

We use the following model for U_k :

$$U_k \triangleq \begin{cases} c_I X_r, & \text{if the } k\text{th frame is an } I \text{ frame} \\ c_P X_r, & \text{if the } k\text{th frame is a } P \text{ frame} \\ c_B X_r, & \text{if the } k\text{th frame is a } B \text{ frame} \end{cases} \quad (11)$$

The constants c_I , c_P , and c_B are obtained empirically as follows: $c_I = \frac{I_{avg}}{X_{avg}}$, $c_P = \frac{P_{avg}}{X_{avg}}$, and $c_B = \frac{B_{avg}}{X_{avg}}$, where I_{avg} , P_{avg} , and B_{avg} are the average (empirical) frame sizes for *I*, *P*, and *B* frames, respectively; and X_{avg} is the average size of a GOP. Note that $c_I + (N/M - 1)c_P + (N - N/M)c_B = 1$, ensuring that the sum of frame sizes in the r th GOP is equal to X_r . According to this model, *B* frames (also, *P* frames) that belong to the same GOP have the same size. So if the k th frame is of type *B*, then U_k represents the *average* size of *B* frames within the corresponding GOP. An example of the resulting sample path based on this model is shown in Figure 2.

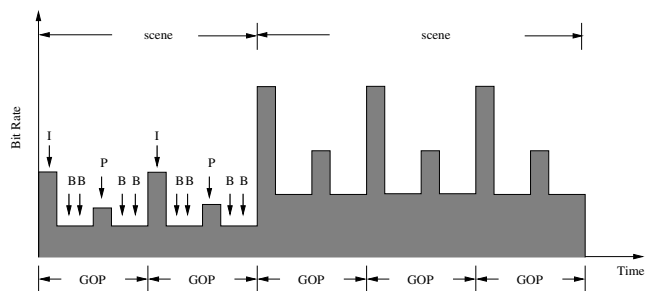



Fig. 2. Bit-rate variations in frame-level model ($N = 6$, $M = 3$).

It should be emphasized that the above frame-level model is indeed an approximation. In real video sequences, the ratio of the average size of, say, *P* frames in a GOP and the size of that GOP varies from one GOP to another. However, this variation is observed to be relatively small as illustrated in Figure 3, which depicts the time-varying coefficients $c_I(r)$, $c_P(r)$, and $c_B(r)$ as a function of r for the MPEG coded sequence *Lecture* [14]. Here, $c_T(r)$, $T \in \{I, P, B\}$, is the ratio of the *average* size of type- T frames in the r th GOP and the size of that GOP. The sample mean and standard deviation for each of these coefficients are shown in Table I. While a more elaborate frame-level characterization is possible, the attractiveness of (11) is that the coefficients c_I , c_P , and c_B can, in principle, be estimated a priori. Furthermore, a detailed frame-level characterization would not lend itself to the type of analysis presented in this paper.

According to the model in (11), the marginal distributions for the sizes of the three frame types are given in terms of f_X as follows: $f_I(x) = f_X(x/c_I)$, $f_P(x) = f_X(x/c_P)$, and $f_B(x) = f_X(x/c_B)$. Let U_I , U_P , and U_B be three generic rvs that indicate the sizes of arbitrary *I*, *P*, and *B* frames, respectively.


 Fig. 3. Time-varying coefficients $c_I(j)$, $c_P(j)$, and $c_B(j)$.

Frame Type (T)	Average $c_T(r)$	Std. of $c_T(r)$
I	0.5348	0.0412
P	0.2378	0.0215
B	0.0568	0.0067

TABLE I

SAMPLE MEAN AND STANDARD DEVIATION OF THE TIME-VARYING COEFFICIENTS (*Lecture* TRACE).

It readily follows that $E[U_I] = c_I m$, $E[U_P] = c_P m$, and $E[U_B] = c_B m$. Also, $\text{var}(U_I) \triangleq \sigma_I^2 = c_I^2 \sigma_X^2$, $\text{var}(U_P) \triangleq \sigma_P^2 = c_P^2 \sigma_X^2$, and $\text{var}(U_B) \triangleq \sigma_B^2 = c_B^2 \sigma_X^2$. Let $\rho_U(k)$ be the ACF of $\{U_n : n = 1, 2, \dots\}$ at lag k

$$\rho_U(k) = \frac{E[U_1 U_{1+k}] - \tilde{m}^2}{\sigma_U^2}$$

where $\tilde{m} \triangleq E[U_1] = \frac{m}{N}$ and $\sigma_U^2 \triangleq \text{var}(U_1)$ is given by

$$\frac{(\sigma_X^2 + m^2)}{N} \left[c_I^2 + \left(\frac{N}{M} - 1\right) c_P^2 + \left(N - \frac{N}{M}\right) c_B^2 \right] - \frac{m^2}{N^2}$$

Recall that according to our model, the first frame of an MPEG stream is selected randomly from the N frames of a GOP. Thereafter, the MPEG sequence proceeds according to the repetitive GOP pattern. Consider $E[U_1 U_{1+k}]$ for $k > 1$:

$$E[U_1 U_{1+k}] = \sum_{i=1}^N E[U_1 U_{1+k} / T_1 = i] \Pr[T_1 = i]$$

where T_i is a discrete rv that reflects the location (and consequently, the type) of the i th frame in the GOP pattern. The sample space of T_i is $\Omega_T = \{1, 2, \dots, N\}$. Thus, $T_j \stackrel{a.s.}{=} i$ means that the type of the j th frame is

the same as the type of the frame in the i th location of the GOP pattern. Because of the repetitive application of the GOP pattern, the process $\{T_n : n = 1, 2, \dots\}$ constitutes a deterministic Markov chain with transition probabilities $p_{ij} = \Pr[T_n = j / T_{n-1} = i] = 1$ if $j = i+1$ and $i = 1, \dots, N-1$ or if $i = N$ and $j = 1$, and zero otherwise. Our previous assumption related to the type of the first frame can now be stated formally by taking the initial distribution of the Markov chain to be its stationary distribution, i.e., $\pi_i \triangleq \Pr[T_1 = i] = 1/N$ for all $i \in \Omega_T$. Hence,

$$E[U_1 U_{1+k}] = \frac{1}{N} \sum_{i=1}^N E[U_1 U_{1+k} / T_1 = i] \quad (12)$$

Before proceeding with the computation of $E[U_1 U_{1+k} / T_1 = i]$, we need to define some related quantities. Let

$$g_N(i, k) \triangleq (i + k - 1) \bmod N$$

$$g_M(i, k) \triangleq (i + k - 1) \bmod M$$

where $i \in \Omega_T$ and k is a positive integer. Note that because N is a multiple of M , if $g_N(i, k) = 0$ then $g_M(i, k) = 0$ as well. Define the following two sets:

$$\Omega_P \triangleq \{1 + M, 1 + 2M, 1 + 3M, \dots, 1 + (N/M - 1)M\}$$

$$\Omega_B \triangleq \Omega_T - \{1\} - \Omega_P$$

Next, we define the following function $\eta(i, k)$:

- **Case 1:** $i = 1$
 - If $g_N(1, k) = 0$, then $\eta(1, k) \triangleq c_I^2$.
 - If $g_N(1, k) \neq 0$ but $g_M(1, k) = 0$, then $\eta(1, k) \triangleq c_I c_P$.
 - If $g_M(1, k) \neq 0$, then $\eta(1, k) \triangleq c_I c_B$.
- **Case 2:** $i \in \Omega_P$
 - If $g_N(i, k) = 0$, then $\eta(i, k) \triangleq c_I c_P$.
 - If $g_N(i, k) \neq 0$ but $g_M(i, k) = 0$, then $\eta(i, k) \triangleq c_P^2$.
 - If $g_M(i, k) \neq 0$, then $\eta(i, k) \triangleq c_P c_B$.
- **Case 3:** $i \in \Omega_B$
 - If $g_N(i, k) = 0$, then $\eta(i, k) \triangleq c_I c_B$.
 - If $g_N(i, k) \neq 0$ but $g_M(i, k) = 0$, then $\eta(i, k) \triangleq c_P c_B$.
 - If $g_M(i, k) \neq 0$, then $\eta(i, k) \triangleq c_B^2$.

It can be shown that $\sum_{i=1}^N \eta(i, k) =$

$$c_I^2 + c_P^2(N/M - 1) + c_B^2(N - N/M), \text{ if } k \bmod N = 0$$

$$2c_I c_P + (N - N/M)c_B^2 + (N/M - 2)c_P^2, \text{ if } k \bmod N \neq 0$$

$$\text{but } k \bmod M = 0 \quad (13)$$

$$2c_I c_B + 2(N/M - 1)c_P c_B + (N - 2N/M)c_B^2, \text{ if } k \bmod M \neq 0$$

Furthermore, it is easy to verify that $\sum_{k=1}^N \sum_{i=1}^N \eta(i, k) = 1$.

We now return to the problem of determining $E[U_1 U_{1+k} / T_1 = i]$. There are two cases to consider. First, when $i+k \leq N$, Frames 1 and $1+k$ must belong to the same (first) GOP, hence

$$E[U_1 U_{1+k} / T_1 = i] = \eta(i, k) E[X_1^2] = \eta(i, k) [\sigma_X^2 + m^2] \quad (14)$$

When $i+k > N$, Frames 1 and $1+k$ belong to different GOPs and possibly to different scenes. More specifically, the $(1+k)$ th frame belongs to the r th GOP, where $r \triangleq \lceil (i+k)/N \rceil > 1$. Thus,

$$E[U_1 U_{1+k} / T_1 = i] = \eta(i, k) E[X_1 X_{r-1}] = \eta(i, k) [\sigma_X^2 \rho_X(r-1) + m^2] \quad (15)$$

A. Computation of $E[U_1 U_{1+k}]$

Case I: $k = 1, 2, \dots, N-1$

From (12) and (13), and based on the previous discussion, we have

$$E[U_1 U_{1+k}] = \frac{1}{N} \sum_{i=1}^{N-k} \eta(i, k) (\sigma_X^2 + m^2) + \frac{1}{N} \sum_{i=N-k+1}^N (\sigma_X^2 \rho_X(r-1) + m^2)$$

where, as before, $r \triangleq \lceil (i+k)/N \rceil$. For $k = 1, \dots, N$ and $i = N-k+1, \dots, N$, $r = 2$. Thus,

$$E[U_1 U_{1+k}] = \frac{m^2}{N} \sum_{i=1}^N \eta(i, k) + \frac{\sigma_X^2}{N} \left(\sum_{i=1}^{N-k} \eta(i, k) + \rho_X(1) \sum_{i=N-k+1}^N \eta(i, k) \right)$$

Case II: $k \geq N$

Starting with (15), this case is further divided into two subcases.

Case II-A: $k = pN$ for $p = 1, 2, \dots$

In this case, Frames 1 and $1+k$ differ exactly by p GOPs, irrespective of the value of i . Thus, $r = \lceil i/N \rceil + p = 1+p$. Accordingly,

$$E[U_1 U_{1+k}] = \frac{1}{N} \sum_{i=1}^N \eta(i, k) [\sigma_X^2 \rho_X(p) + m^2] = \frac{\sigma_X^2 \rho_X(p) + m^2}{N} \sum_{i=1}^N \eta(i, k)$$

Finally,

$$\rho_U(k) = \frac{[(\sigma_X^2 \rho_X(p) + m^2) (\sum_{i=1}^N \eta(i, k)) / N] - \tilde{m}^2}{\sigma_U^2} \quad (16)$$

Case II-B: $k \neq pN$ and $k > N$

In this case, Frames 1 and $1+k$ may differ by either p GOPs or by $p+1$ GOPs, where $p \triangleq \lfloor k/N \rfloor - 1 = \lfloor k/N \rfloor$. More specifically, for $T_1 = 1, 2, \dots, i^*$, where $i^* = N \lfloor k/N \rfloor - k$, the 1st and $(1+k)$ th frames differ by p GOPs. For $T_1 = i^* + 1, \dots, N$, the two frames differ by $p+1$ GOPs. Thus,

$$E[U_1 U_{1+k}] = \frac{1}{N} \sum_{i=1}^{i^*} \eta(i, k) [\sigma_X^2 \rho_X(p) + m^2] + \frac{1}{N} \sum_{i=i^*+1}^N \eta(i, k) [\sigma_X^2 \rho_X(1+p) + m^2] \quad (17)$$

which concludes the derivation of the ACF for the process $\{U_n : n = 1, 2, \dots\}$.

B. Asymptotic Behavior of Frame-Level ACF

From the analytical form of ρ_U , one can examine its asymptotic behavior, shedding light on the LRD/SRD structure. As $k \rightarrow \infty$, $p \rightarrow \infty$ and $\rho_X(p) \rightarrow 0$, so that

$$\lim_{k \rightarrow \infty} \rho_U(k) = \frac{\frac{m^2}{N} \lim_{k \rightarrow \infty} \sum_{i=1}^N \eta(i, k) - \tilde{m}^2}{\sigma_U^2} \quad (18)$$

The limit of $\sum_{i=1}^N \eta(i, k)$ as $k \rightarrow \infty$ alternates between the three values given in (13). Substituting the values of \tilde{m} and σ_U^2 in (18), it is easy to see that $\lim_{k \rightarrow \infty} \rho_U(k)$ alternates between the following three values, depending on how k approaches infinity:

$$\xi_j^* = \frac{N \eta_j^* - 1}{N \eta_1^* (\sigma_X^2 / m^2 + 1) - 1}, \quad j = 1, 2, 3 \quad (19)$$

where η_1^* , η_2^* , and η_3^* are the three values in (13), respectively. In general, ξ_1^* , ξ_2^* , and ξ_3^* are non-zero, which justifies the persistent, periodic autocorrelations that are observed in empirical MPEG sequences. However, using (13) it can be shown that $\xi_1^* + (N/M - 1)\xi_2^* + (N - N/M)\xi_3^* = 0$, i.e., the sum of the autocorrelations over a GOP period converges to zero, as expected.

IV. VALIDATION OF ANALYTICAL RESULTS

In this section, we demonstrate the validity of our analytical expressions using three numerical examples. For simplicity, we ignore the noise process (the inter-scene variations). Our validation approach is based on

comparing the analytical ACF against the sample ACF of synthetically generated VBR sequences. In the first two examples, we investigate the ACF at the GOP level assuming gamma distributed GOP sizes with mean of 500 and standard deviation of 100.

In the first example, we use a shifted exponential scene-length distribution:

$$\Pr[S > k] = \Pr[\hat{S} > k] = e^{-\beta(k-1)}, \quad k = 1, 2, \dots \quad (20)$$

Note that in this case S and \hat{S} have the same distribution. We set $\beta = 1/49$, so that $E[S] = 50$. Ten synthetic traces were generated, and their sample ACFs were computed and averaged. The average ACF for the synthetic traces is plotted in Figure 4 along with its theoretical counterpart. There is a clear match between the two plots.

Next, we consider a subgeometric scene-length distribution of the form:

$$\Pr[S > k] = \alpha^{\sqrt{k}}, \quad k = 1, 2, \dots$$

for some $0 < \alpha < 1$. In this case, the ACF can be written recursively as follows:

$$\rho_X(k+1) = \rho_X(k) - \frac{\alpha^{\sqrt{k}}}{E[S]} \quad (21)$$

While a closed-form expression for $E[S] = \sum_{k=0}^{\infty} \alpha^{\sqrt{k}}$ is not available, it is easy to show that

$$\frac{2}{(\ln \alpha)^2} \leq E[S] \leq \frac{2}{(\ln \alpha)^2} + 1 \quad (22)$$

Setting $\alpha = 0.8$, we have $40.17 \leq E[S] \leq 41.17$. Thus, $E[S] \approx \frac{2}{(\ln \alpha)^2} + 0.5 = 40.67$. Figure 5 depicts the theoretical and empirical ACFs under a subgeometric scene-length distribution. At small and large lags, the plots match very well. At intermediate lags, there is a slight difference that is attributed to the large variance of the empirical autocorrelations and to other approximations in the generation of subgeometrically distributed random numbers.

Our last example is related to the frame-level ACF. Here, we use the same shifted exponential scene-length distribution as in the first example. We set $N = 12$, $M = 3$, $c_I = 5/22$, $c_P = 3/22$, and $c_B = 1/22$. The analytical and empirical ACFs are shown in Figure 6 for lags in the range 450 to 500. This range is chosen arbitrarily, and is representative of the behavior at large lags. The two ACFs almost match at all examined lags (similar trend is also observed at small lags). Note that although the scene-length distribution is exponential, the deterministic interleaving of three, drastically different processes (one for each frame type) induces strong correlations that determine

the asymptotic shape of the ACF. These correlations do not die out to zero as the lag goes to infinity, but instead they converge to $\eta_1^* = 0.8912$, $\eta_2^* = 0.710$, and $\eta_3^* = -0.3776$.




Fig. 4. ACF for GOP-level model with shifted exponential scene distribution ($\beta = 1/49$).




Fig. 5. ACF for the GOP sequence with a subgeometric scene-length distribution ($\alpha = 0.8$).

V. IMPACT OF CORRELATIONS ON BUFFER PERFORMANCE

In this section, we investigate the impact of traffic correlations on the queueing performance at a video buffer. The scenario we consider was depicted in Figure 1, where a VBR stream is fed into a buffer with a drain rate C . We study two families of scene-length distributions: Pareto and Weibull. While other distributions may also be used, lately these two distributions have been receiving much attention (see [9]).

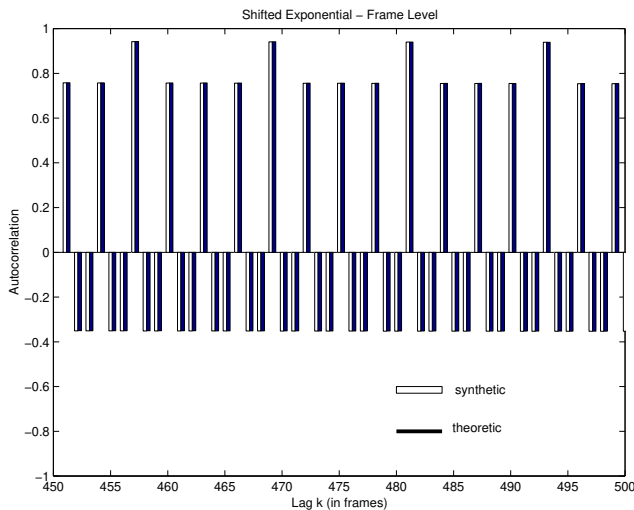


Fig. 6. ACF at the frame level based on a shifted exponential scene-length distribution ($\beta = 1/49$, $N = 12$, $M = 3$).

For each family of distributions, we examine the effect of varying the correlations persistence on the queueing performance. In line of the findings in [6], [10], we take the GOP size to be gamma distributed with scale and shape parameters λ and w . For simplicity, we take w to be integer valued.

For Pareto distributed scene lengths, our analysis is based on the work of Jelenkovic and Lazar [12] on the subexponential asymptotics of Markov-modulated random walks (see also [13]). As discussed in [13], Pareto distributions belong to the class of *regularly varying* distributions R_α , which have recently been the focus of several investigations. Consider a model with Pareto distributed scene lengths:

$$F_S(x) \triangleq \Pr[S \leq x] = 1 - \frac{1}{x^\alpha}, \quad x \geq 1 \text{ and } \alpha > 1$$

Since we assume that scene changes constitute a renewal process, our model is similar to the Markov renewal process that was studied in [13], with the exception that in [13] a 4-state discrete Markov chain governs the transitions between “regimes” (i.e., video scenes). In contrast, we consider a continuous and unbounded state space with transitions that are independent of the current state. In the discrete case, if the scene-length distribution is regularly varying, the asymptotic behavior of the queue length is given by [13]:

$$\Pr[Q > x] \sim \frac{\sum_{d_i > 0} \pi_i d_i^\alpha}{E[S](C - E[A])} \rho_X(x) \text{ as } x \rightarrow \infty \quad (23)$$

where Q is the steady-state queue length at renewal instants, π_i is the stationary probability of being in state i , A is the steady-state arrival rate at renewal times, and d_i is the mean drift rate while in state i (arrival

rate minus service rate). As before, $\rho_X(x)$ is the GOP-level ACF at lag x . It is assumed that the buffer capacity is infinite and that the system is “weakly stable”, i.e., $d_i > 0$ for at least one state i . In our case, because of the infinite sample space of the gamma distribution, the system is indeed weakly stable. Furthermore, because in our model transitions at scene boundaries are independent of the scene levels, π_i in (23) is just the pdf of the GOP size. Hence (23) becomes:

$$\Pr[Q > x] \sim \frac{\int_C^\infty (u - C)^\alpha f_\Gamma(u) du}{E[S](C - E[A])} \rho_X(x), \text{ as } x \rightarrow \infty \quad (24)$$

where f_Γ is the gamma pdf of the GOP size. For simplicity, we take α to be integer valued, $\alpha \geq 2$. Substituting the expressions for $E[S] = \alpha/(\alpha - 1)$, $E[A] = w/\lambda$, and $f_\Gamma(u)$ in (24), as $x \rightarrow \infty$ we have

$$\Pr[Q > x] \sim \frac{\int_C^\infty (u - C)^\alpha e^{-\lambda u} (\lambda u)^{w-1} \lambda du}{\frac{\alpha}{\alpha-1} (C - w/\lambda) (w-1)!} \rho_X(x)$$

With some manipulations, it can be shown that the above equation reduces to

$$\left(\frac{(\alpha - 1)e^{-\lambda C}}{\alpha \lambda^\alpha (C - w/\lambda)} \sum_{j=0}^{w-1} \frac{(w + \alpha - j - 1)! (C\lambda)^j}{(w - j - 1)! j!} \right) \rho_X(x) = \Psi(\lambda, w, C, \alpha) \cdot \rho_X(x)$$

where

$$\Psi(\lambda, w, C, \alpha) \triangleq \frac{(\alpha - 1)e^{-\lambda C}}{\alpha \lambda^\alpha (C - w/\lambda)} \sum_{j=0}^{w-1} \frac{(w + \alpha - j - 1)! (C\lambda)^j}{(w - j - 1)! j!}$$

The above expression directly relates the ACF of the GOP-level model to the queueing performance. For Pareto distributed scene lengths, the ACF $\rho_X(x)$ can be easily computed as follows:

$$\rho_X(x) = \Pr[\hat{S} > x] = \int_x^\infty \frac{\Pr[S \geq u]}{E[S]} du = \frac{x^{1-\alpha}}{\alpha}$$

for $\alpha > 1$ and $x > 0$. Figure 7 depicts the buffer overflow rate versus the buffer size. As expected, for a given buffer size the buffer overflow probability decreases as α increases. Interestingly, for $\alpha = 2, 3, 4$, the buffer overflow curve flattens fairly quickly despite the fact that the underlying model is SRD. While such behavior is already known for LRD models ($1 < \alpha < 2$), its presence under SRD models is surprising. Nonetheless, for large α (e.g., $\alpha = 6$), the buffer overflow curve starts to get steeper and the performance becomes more sensitive to changes in the video buffer size. It is worth mentioning that in the case of exponentially distributed scene lengths (i.e., a Markovian model), the buffer overflow probability plotted on a

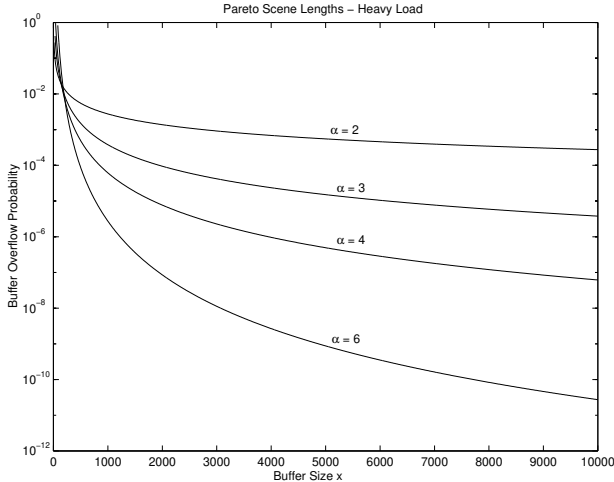


Fig. 7. Probability of buffer overflow when scene lengths are Pareto distributed ($\lambda = 0.05$, $w = 25$, load = 83%).

logarithmic scale drops off *linearly* with the buffer size, i.e., the curves in Figure 7 would have constant slopes.

Next, we consider the following class of discrete distributions:

$$\Pr[S > k] = \alpha^{\sqrt[n]{k}}, \text{ for } k = 1, 2, \dots, \text{ and } n = 1, 2, \dots \quad (25)$$

where $0 < \alpha < 1$. This is a special case of the general Weibull distribution $F(x) = 1 - e^{-\beta x^r}$, where $\beta > 0$ and $0 < r < 1$ (in (25) we set $\beta = -\ln \alpha$ and $r = 1/n$). When $n = 1$, the scene length distribution is geometric, whereas $n \geq 2$ gives rise to a subgeometric model. For a finite n , $E[S^2] < \infty$ and the corresponding video model is SRD (Proposition 1). As n increases the correlations become more persistent, and as $n \rightarrow \infty$, the model approaches the LRD regime. Figure 8 depicts the GOP-level ACF for $n = 1, 2, 3$ and $\alpha = 0.8$. Increasing α slows down the speed of convergence of the ACF. Note that for the underlying family of scene distributions, the ACF is obtained recursively using:

$$\rho_X(k+1) = \rho_X(k) - \frac{\alpha^{\sqrt[n]{k}}}{E[S]}, \quad k = 1, 2, \dots$$

where $E[S] \approx n! / (-\ln \alpha)^n$. We use simulations to evaluate the impact of correlations on the queueing performance. A disadvantage of simulations in this case is that they require extremely long traces to obtain any meaningful results (the more persistent the correlations, the longer the traces). This means that credible results can only be obtained for relatively large and moderate loss rates (above 10^{-4}). In our simulations, we assume that video frames are packetized into fixed-size packets (e.g., ATM cells). We investigate the packet loss rate (PLR) under GOP- and frame-level models assuming both finite and infinite buffer capacities. In the latter case, the PLR is estimated by the

percentage of packets that arrive at the buffer and find B or more packets in the queue, where B is the buffer size in the finite-buffer case.




Fig. 8. ACF for subgeometrically distributed scene lengths.

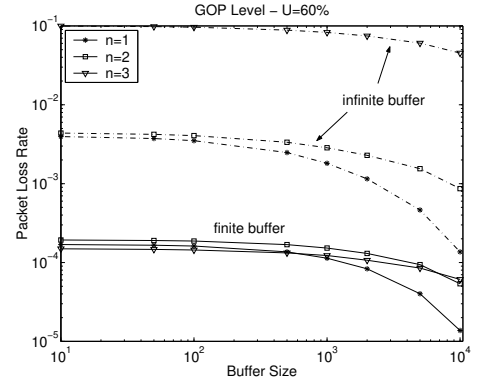
In our experiments, we fix the mean scene length at $E[S] = 50$ GOPs. Since $E[S]$ depends on α and n , α is adjusted whenever n is varied. For GOP-level results, we ran the simulations using synthetic traces of length 1,000,000 GOPs per trace. As before, GOP sizes are gamma distributed with shape parameter $w = 25$ and scale parameter $\lambda = 0.05$. For frame-level results, we set $N = 12$ and $M = 3$ with each trace consisting of 12,000,000 frames. A sufficient number of independent runs was used to ensure tight confidence intervals. To avoid cluttering the figures, we only show the average values of these runs.

Figure 9 depicts the PLR under two traffic loads ($U = 60\%$ and 80%) for GOP- and frame-level models and with finite- and infinite-capacity buffers. Based on these figures, several important observations can be made. First, the degree of correlations persistence, which is reflected in the value of n , has clear impact on the sensitivity of the PLR to changes in the buffer size; the larger the value of n the less sensitive is the performance. This means that for very large buffers, the degree of correlations persistence *does* matter. Interestingly, this trend is observed *for both finite- and infinite-capacity buffers*. Second, for small buffer sizes (less than 100 packets), the frame-level PLR is always larger than its GOP-level counterpart. The discrepancy between the two is more obvious when the buffer capacity is finite. Furthermore, this discrepancy is more pronounced at lower traffic loads ($U = 60\%$), where the difference can reach several orders of magnitude. As the buffer size increases, the discrepancy between the GOP- and frame-level results fades away. Based on our

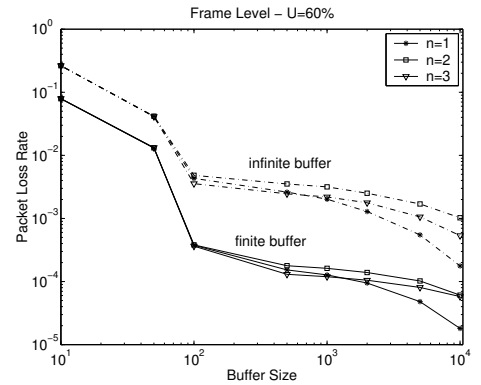
numbers, a buffer size of 100 packets amounts to a maximum queuing delay of $B/C = 100/C$, where C is the link bandwidth in packets/second. For $U = 80\%$ and average input rate of 500 packets/GOP (1000 packets/second), $C = 1000/U = 1250$ packets/second. Thus, the maximum delay is 80 msec, which is slightly less than the time to generate three frames. Hence, if the delay requirements of the video application are such that a delay of three frames at the encoder is tolerable, a GOP-level model is sufficiently accurate for use in performance evaluation and capacity planning studies. Otherwise, frame-level modeling is needed. Our last remark is related to the infinite-buffer results in Part (a) of Figure 9. As n goes from 2 to 3, the PLR increases suddenly by more than an order of magnitude. This trend was not observed at the higher load (Part (c)). Our justification of this phenomenon is that at high loads, buffer overflow is more frequent and is not only caused by the very rare events. In contrast, as we decrease the load, rare events (e.g., a high-action scene that lasts for a long period of time) become the primary cause of packet loss. Such events are directly related to the persistence of the autocorrelations, hence the observed trend. The greater impact of n at lower loads is analogous to its increased significance as the buffer size increases.

VI. CONCLUSIONS

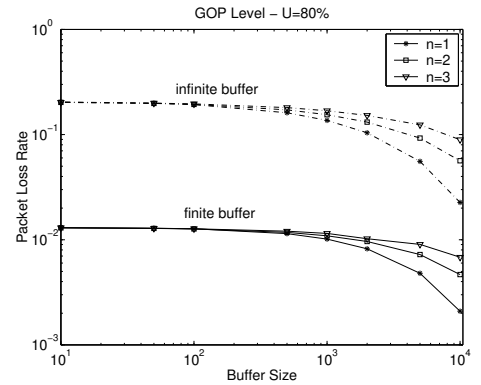
In this paper, we analyzed the ACF for a class of scene-based video models. Our analysis was performed at both GOP and frame levels, and was used to establish the relationship between the SRD/LRD structure of a model and its scene-length distribution. As a byproduct of this relationship, an efficient procedure for fitting the scene-length distribution was provided, which only requires fitting of the ACF. At the frame level, our results indicate that the repetitive application of the GOP pattern induces strong periodic components in the ACF. In fact, we showed that the frame-level ACF *does not* converge to zero as the frame lag goes to infinity. This, somehow surprising, result can be extended to composite processes in which two drastically different submodels are interleaved in a deterministic manner (e.g., composition of audio and video streams). The impact of correlations on the performance at a video buffer was studied via analysis and simulations for video models with Pareto and Weibull scene-length distributions. In the case of Pareto scene lengths, we observed that the insensitivity of the packet loss rate to changes in the buffer size extends beyond the LRD regime of the Pareto distribution ($1 < \alpha < 2$) to the SRD regime ($\alpha \geq 2$). Such insensitivity starts to change as α becomes large. For Weibull distributed scene lengths, several important observations can be made based on the simulation results. First, the more



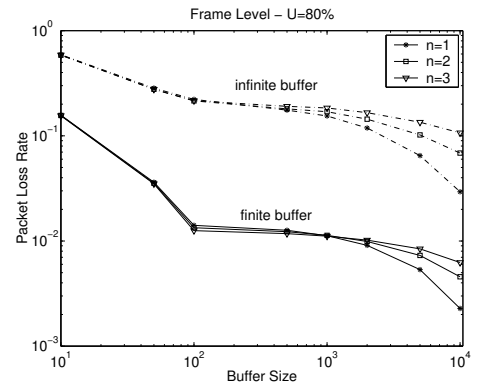
(a)



(b)



(c)



(d)

Fig. 9. PLR versus buffer size under Weibull distributed scene lengths.

persistent the correlations, the less sensitive the performance to changes in the buffer size. This trend was observed under both finite- and infinite-capacity buffers. Thus, for large, *finite* buffers, the degree of correlations persistence does matter. Second, for small buffer sizes (less than 100 packets), the frame-level performance is always worse than its GOP-level counterpart, with the discrepancy being more pronounced when the buffer capacity is finite and the traffic load is low. As the buffer size increases, this discrepancy fades away, and a GOP-level model becomes sufficient for analyzing the performance. Third, the impact of correlations persistence becomes more profound as the traffic load is decreased. Our work provides important guidelines that can be used in the design and dimensioning of video buffers and for efficient allocation of network bandwidth.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their constructive comments.

REFERENCES

- [1] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger. Long-range dependence in variable bit-rate video traffic. *IEEE Transactions on Communications*, 43:1566–1579, 1995.
- [2] O. J. Boxma and V. Dumas. Fluid queues with long-tailed activity distributions. technical report Report BS-R9705, Centrum voor Wiskunde en Informatica (CWI), 1997.
- [3] A. M. Dawood and M. Ghanbari. Content-based MPEG video traffic modeling. *IEEE Transactions on Multimedia*, 1(1):77–87, Mar. 1999.
- [4] H. J. Fowler and W. E. Leland. Local area network traffic characteristics, with implications for broadband network congestion management. *IEEE Journal on Selected Areas in Communications*, 9:1139–1149, 1991.
- [5] M. R. Frater, J. F. Arnold, and P. Tan. A new statistical model for traffic generated by VBR coders for television on the Broadband ISDN. *IEEE Trans. on Circuits and Systems for Video Technology*, 4(6):521–526, Dec. 1994.
- [6] M. W. Garrett and W. Willinger. Analysis, modeling, and generation of self-similar VBR video traffic. In *Proc. of the SIGCOMM '94 Conference*, pages 269–280, Sept. 1994.
- [7] M. Grossglauser and J.-C. Bolot. On the relevance of long-range dependence in network traffic. In *Proceedings of the ACM SIGCOMM '96 Conference*, Aug. 1996.
- [8] D. Heyman and T. Lakshman. What are the implications of long-range dependence for VBR video traffic engineering? *IEEE/ACM Transactions on Networking*, 4:301–317, June 1996.
- [9] D. P. Heyman and T. V. Lakshman. Source models for VBR broadcast-video traffic. *IEEE/ACM Transactions on Networking*, 4(1):40–48, Feb. 1996.
- [10] D. P. Heyman, A. Tabatabai, and T. V. Lakshman. Statistical analysis and simulation study of video teleconferencing traffic in ATM networks. *IEEE Trans. on Circuits and Systems for Video Technology*, 2(1):49–59, Mar. 1992.
- [11] M. Izquierdo and D. Reeves. A survey of statistical source models for variable bit-rate compressed video. Technical Report 97-10, Center for Advanced Computing and Communication, North Carolina State University, June 1997.
- [12] P. R. Jelenkovic and A. A. Lazar. Subexponential asymptotics of a Markov-modulated random walk with queueing applications. *Journal of Applied Probability*, 35(2), June 1998.
- [13] P. R. Jelenkovic, A. A. Lazar, and N. Semret. The effect of multiple time scales and subexponentiality in MPEG video streams on queueing behavior. *IEEE Journal on Selected Areas in Communications*, 15(6):1052–1071, Aug. 1997.
- [14] E. W. Knightly, D. Wrege, J. Liebeherr, and H. Zhang. Fundamental limits and tradeoffs of providing deterministic guarantees to VBR video traffic. In *Proc. of the ACM SIGMETRICS/PERFORMANCE '95 Conference*, pages 98–107, May 1995.
- [15] M. Krunz and S. K. Tripathi. On the characterization of VBR MPEG streams. In *Proceedings of the ACM SIGMETRICS '97 Conference*, pages 192–202, June 1997.
- [16] A. A. Lazar, G. Pacifici, and D. E. Pendarakis. Modeling video sources for real-time scheduling. *Multimedia Systems Journal*, 1(6):253–266, 1994.
- [17] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, Feb. 1994.
- [18] B. Melamed, D. Raychaudhuri, B. Sengupta, and J. Zdepiski. TES-based video source modeling for performance evaluation of integrated networks. *IEEE Transactions on Communications*, 42(10):2773–2777, Oct. 1994.
- [19] O. Rose. Simple and efficient models for variable bit rate MPEG video traffic. *Performance Evaluation*, 30:69–85, 1997.
- [20] B. Ryu and A. Elwalid. The importance of long-range dependence of VBR video traffic in ATM traffic engineering: Myths and realities. In *Proceedings of the ACM SIGCOMM '96 Conference*, pages 3–14, Aug. 1996.

Marwan M. Krunz received the B.S. in electrical engineering from Jordan University, Jordan, in 1990, and the M.S. and Ph.D. degrees in Electrical Engineering from Michigan State University in 1992 and 1995, respectively. From 1995 to 1997, he was a postdoctoral research associate with the Department of Computer Science and the Institute for Advanced Computer Studies at the University of Maryland, College Park. In January 1997,

he joined the Department of Electrical and Computer Engineering at the University of Arizona, where he is currently an Assistant Professor. His research interests are in teletraffic modeling, resource allocation in wireless networks, and QoS-based routing.

Dr. Krunz received the National Science Foundation CAREER Award in 1998. He is a Technical Editor for the IEEE Communications Interactive Magazine, and has served and continue to serve on the executive and program committees of several technical conferences.

Arivu Mani Ramasamy received the B.S. degree in electrical engineering from the Institute of Technology, Varanasi, India, in 1997, and the M.S. degree in electrical engineering from the University of Arizona, Tucson, in 1999. Since June 1999 he has been with Cisco Systems. His research interests are in quality of service, Web traffic analysis, and video modeling.