Proceedings of the Communications Theory Mini-Conference at GLOBECOM '94, San Francisco, California, pp. 120–126

A TES-BASED MODEL FOR COMPRESSED "STAR WARS" VIDEO

BENJAMIN MELAMED

NEC USA Inc.
4 Independence Way
Princeton, NJ 08540
Email: ben@ccrl.nj.nec.com, Fax: 609-951-2499

DIMITRIOS PENDARAKIS

Department of Electrical Engineering
and Center for Teletraffic Research
Columbia University
New York, NY 10027

Abstract

This paper presents a stochastic model for the variable bit rate (VBR) video of the entire "Star Wars" movie under a JPEG-like compression. The goal of this work is to develop a methodology for constructing a suite of high-fidelity source models of full-motion video as a traffic benchmark suitable for driving realistic Monte Carlo simulations of broadband telecommunications networks. To this end, a TES-based modeling methodology was employed, subject to fidelity criteria requiring the simultaneous capture of both first-order and second-order statistics. The entire movie was modeled as a Markov-Renewal-Modulated TES (MRMT) process with class-dependent holding times. The modeling results support the efficacy of the modeling methodology and suggest that high-fidelity compact source models of full-motion video can be developed by applying this methodology to other types of movies under a variety of coding schemes.

1 Introduction

This work has been motivated by the need to create accurate, yet simple and flexible, workload source models of realistic variable bit rate (VBR) compressed video. It is expected that video traffic will constitute a significant part of the load in future broadband Integrated Services Digital Networks (B-ISDN), especially of the Asynchronous Transfer Mode (ATM) variety. Such systems will be collectively referred to as broadband networks, for short.

The size of broadband networks, combined with the complexity and diversity of their projected traffic mix (voice, video, file transfers etc.), will doubtless render the task of network design and management a difficult one. Accurate modeling of broadband networks calls for careful performance analysis, mainly via simulation experiments. Such simulations must rely critically on high-fidelity traffic modeling, especially for bursty traffic. Intuitively, burstiness is present in a traffic process, if arrival points appear to form visual clusters on the time line. The mathematical underpinnings of burstiness are complex; two main factors contributing to traffic burstiness are the marginal distribution and autocorrelation function of traffic interarrival times, particularly strong positive short-term autocorrelations [3]. The impact of autocorrelation in traffic processes on queueing measures such as mean queue length, mean waiting times and loss probabilities in finite buffers, can be very dramatic, even in light traffic [6, 10, 14].

The studies cited above support the view that modeling a bursty traffic stream (one with autocorrelated interarrival times or bit rates) as a renewal process (independent variates) can entail a significant modeling risk. Worse still, such traffic models are bound to err on the optimistic side, if the modeled traffic has appreciable short-term positive autocorrelations. Unfortunately, this situation is prevalent in variable bit rate (VBR) coded video. The reason is quite straightforward: successive video frames differ very little as bitmaps, since their rate of generation is about 30 frames/second. Consequently, a coded video sequence from a single scene possesses an inherent temporal dependence, which is statistically manifested by strong short-term autocorrelations in the random bit rate process. Thus, autocorelated models of coded video are essential to high-fidelity modeling of broadband networks, in spite of their added complexity.

The bulk of previous work in video source modeling has been confined to short sequences of empirical records consisting of one or a few scenes; furthermore, aditional video source models have been designed to capture either the distribution function or the autocorrelation function of empirical bit rate records, but not both; see [3, 4] and references therein. Since traffic burstiness can be caused by the shape of the bit rate distribution as well as by its autocorrelation function, it is important to faithfully capture both first-order and second-order statistics of the empirical data.

This paper presents a novel approach to video modeling. Specifically, it makes two new main contributions. First, it presents a model of a *full-length movie* as opposed to a mere video clip. And second, the model proposed here relies on the TES modeling methodology (to be explained later) to accurately capture both first-order and second-order statistics of the empirical bit rate record; more specifically, the model approximates well the marginal distribution and the autocorrelation function (equivalently, spectral density) of the bit rate process, *simultaneously*, and the model sample path realizations bear a strong *"resemblance"* to the empirical record. Note, though, that time series "resemblance" is a subjective judgment which cannot be precisely quantified. While its exact meaning is best left at the intuitive level, we point out that such judgments are routinely made by applied statisticians in informal manner.

The model presented here is a Markov-Renewal-Modulated TES (MRMT) process, generalizing the well-known Markov-Modulated Poisson Process (MMPP), which had been successfully applied to model voice traffic in telephony [6]. The modeling procedure consisted of the following major steps. First, the entire movie was segmented into scenes. Second, scenes were partitioned into classes by a clustering algorithm and a scene tag was assigned to each scene. The random sequence of scene tags was modeled as a simple Markov chain with just four states (scene classes). Third, the bit rate process of a scene class was modeled as a TES process — one per class. Finally, scene durations were modeled as a class-dependent renewal process. The remainder of the paper describes the model fitted and shows that it complies well with the modeling requirements postulated above.

2 The Empirical Coded Video Data

The empirical coded VBR video data used in this study was the bit rate (frame size) record, obtained from compressing the movie "Star Wars" via a variant of the JPEG coding scheme [11]. This movie is a good representative of full-motion video since it contains a diverse collection of video sequences whose visual complexity and motion content ranges from low to high and very high. The data record consisted of some 171,000 frame sizes (in bytes per frame) — the equivalent of about two hours of video elapsed time at the rate of 24 frames a second. In view of the time complexity of generating high-quality full-motion coded video without hardware support, the coding scheme selected is similar, but not identical, to JPEG. The simplified variant codes only the luminance component, resulting in monochrome video. The digitized frame size is 480×504 pixels; the last 8 pixels on each 512-pixel line form the horizontal retrace interval, and those are dropped before coding. The frame is partitioned into blocks of 8×8 pixels, on which a discrete cosine transform (DCT) is performed. No differential or motion-compensated coding is included. A detailed description of the empirical bit rate record appears in [4].

3 Segmenting the Bit Rate Stream Into Scenes

Figure 1 displays a representative segment of the empirical bit rate record. The VBR sequence of bit rates, $\{X_n\}$, exhibits considerable burstiness as evidenced by the empirical record and the attendant high short-term empirical autocorrelations (see upper canvases in Figure 7).

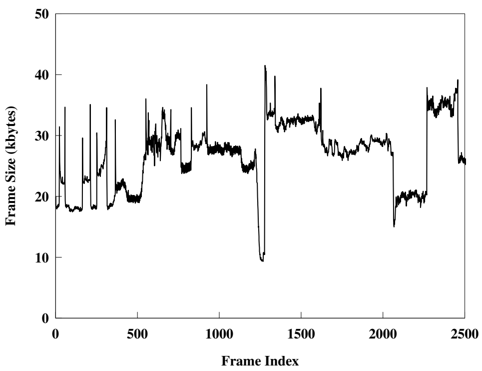

Figure 1: A Low Activity Segment of the Empirical Bit Rate Record.

The appearance of the data plotted in Figure 1 suggests a structure of punctuated stationarity. In other words, we postulate that the record is composed of stationary segments, with random transitions among segments. Segment boundaries are visually characterized by abrupt large magnitude changes in the local average of bit rates. It needs to be established, though, that these segments correspond to distinct *scenes* of the movie, i.e., to movie segments without significant shifts of view; such shifts result from abrupt camera zooming, panning, or cutting to a new visual context. By observing the actual movie frames and the corresponding bit rate sequence in lockstep, we were able to ascertain that scene changes do indeed coincide with abrupt sustained changes in bit rate magnitude.

In order to detect scene boundaries in the empirical bit rate record, we followed a methodology similar to that introduced in [12]. A scene change is declared in the bit rate stream when the absolute value of the difference between successive bit rates exceeds a certain threshold in a sustained manner. This approach did a good job of detecting scene changes in the empirical record.

4 Classifying Scenes into Classes

Having segmented the empirical bit rate sequence into individual scenes, the next objective is to partition the scenes into (stochastic) equivalence classes. The idea is to construct a parsimonious mathematical description of sufficient accuracy, in compliance with the three modeling criteria in Section 1. We, therefore, assume that the empirical record consists of scene realizations from a small number of stationary probability laws (classes), with class-dependent durations, and that transitions among scene classes are governed by a simple Markov chain.

We first decided to use either a scene bit rate mean or scene bit rate variance (or both) as partitioning criteria. We found, however, that scene bit rate mean alone was a good criterion for characterizing scene classes (see below).

As a partitioning methodolgy we chose *partitional clustering* [9]. The general abstract partitioning problem is formulated as follows. Given n elements in a d-dimensional metric space, find a partition of the elements into K equivalence classes (clusters), such that elements in each cluster are "closer" to each other than to elements in all other clusters. The value of K may or may not be prescribed. A clustering metric formalizing the notion of closeness has to be selected. The most common metric used in practice, and the one used in this study, is the Euclidean distance, $\eta(v, w) = \sqrt{\sum_{i=1}^{d}(v_i - w_i)^2}$, for vectors v and w in d-dimensional Euclidean space. The clustering heuristic proceeds iteratively in a trial-and-error fashion. The algorithm requires the specification of two auxiliary parameters: the *minimum distance between cluster centroids*, η_{min}, and the *maximum variation within a cluster*, E^2_{max}; see [9] for more details.

Figure 2 depicts the results of applying the clustering procedure to the set of empirical scene bit rates. Note that four visually distinct clusters are clearly discernible in Figure 2, each marked by a different character. The main point made by this figure is the significance of the mean scene bit rate as a clustering criterion, and the comparable insignificance of the corresponding variance. To see that, note that vertical lines have been drawn in figure 2 to demarcate approximate cluster boundaries by bit rate means; on the other hand, a similar segregation by bit rate variances is not evident.

This clustering pattern has persisted over a wide range of values of the auxiliary clustering parameters, leading us to conclude that the variance of the scene bit rate does not play an important role as a clustering criterion. The variance was subsequently dropped as a clustering criterion, and thereafter, scenes were clustered by bit rate mean only. Henceforth, we will refer to clusters of scenes as *classes* of scenes.

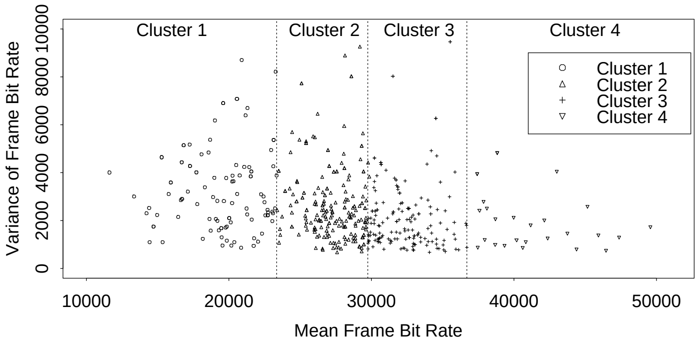

Figure 2: Clustering of a Subset of "Star Wars" Scenes by Mean and Variance Parameters.

5 The Modulating Markov Chain

Recall the modeling assumption, postulated in Section 4, that each class of scenes (cluster) is assumed to consist of scene realizations from the same stationary stochastic process. The random sequence of scene classes was assumed to be governed (modulated) by a simple Markov chain $\{M_n\}$.

The estimation of the 1-step transition matrix, $\mathbf{P}$ of $\{M_n\}$, is straightforward. Once the clustering ranges for the scene bit rate means were determined from the clustering procedure, it is a simple matter to assign a class tag to each scene in the empirical record, resulting in a realization, $\{c_n\}$, of class tags. This sequence is traversed, left to right, and class transitions are counted for each pair of classes. Normalization of these counts (division by the total number of class transitions) then yields the desired transition matrix, $\mathbf{P}$.

6 TES Modeling of Individual Scenes

This section provides a brief overview of TES processes and the TES (Transform-Expand-Sample) modeling methodology [7, 8, 15]. It then describes how TES processes were used as components of the "Star Wars" model to construct models of frame bit rates and scene durations for each scene class.

6.1 TES Processes

TES modeling differs from other modeling approaches in that it aims to fit a prescribed marginal distribution and a prescribed autocorrelation function *simultaneously*, via a stationary stochastic process. Thus, TES models are designed to capture the statistical signature of empirical time series vis-a-vis both first-order and second-order statistics.

The construction of TES models involves two random sequences in lockstep. The first sequence, called a *back-*

ground TES process, serves in an auxiliary role. The particular variant of background TES process used was the so-called TES$^+$ class, consisting of random sequences, $\{U_n^+\}_{n=0}^{\infty}$, of the form

$$U_n^+ \;=\; \begin{cases} U_0, & n=0 \\ \langle U_{n-1}^+ \;+\; V_n\rangle, & n>0 \end{cases} \qquad (1)$$

Here, U_0 is distributed uniformly on $[0,1)$; $\{V_n\}_{n=1}^{\infty}$ is a sequence of iid random variables, independent of U_0, called the *innovation sequence*; and angular brackets denote the modulo-1 (fractional part) operator $\langle x\rangle = x - \max\{\text{integer } n:\ n\le x\}$. The superscript notation in (1) is motivated by the fact that TES$^+$ processes can generate all lag-1 autocorrelations in the ranges [0,1]; another class, called TES$^-$, can generate all corresponding ones in the range [-1,0] (see [15]).

The second sequence,$\{X_n^+\}_{n=0}^{\infty}$, called a *foreground* TES process, is the target TES model. Foreground TES$^+$ sequences have the form

$$X_n^+ \;=\; D(U_n^+), \qquad (2)$$

where D is a transformation from $[0,1)$ to the reals, called a *distortion*.

Given an empirical time series, $\{Y_n\}_{n=0}^{N}$, one uses in practice a composite distortion of the form

$$D_{Y,\xi}(x) \;=\; \hat{H}_Y^{-1}(S_\xi(x)), \quad x\in[0,1). \qquad (3)$$

Here, the inner transformation, S_ξ, is a "smoothing" operation, called a *stitching transformation*, parameterized by $0\le\xi\le1$, and given by

$$S_\xi(y) \;=\; \begin{cases} y/\xi, & 0\le y\le\xi \\ (1-y)/(1-\xi), & \xi\le y<1. \end{cases} \qquad (4)$$

The outer transformation, $\hat{H}_Y^{-1}$, is the inverse of the empirical (histogram) distribution function computed from $\{Y_n\}$ for $x\in[0,1)$ as

$$\hat{H}_Y^{-1}(x) = \sum_{j=1}^{J} 1_{[\hat{C}_{j-1},\hat{C}_j)}(x)\,[l_j + (x-\hat{C}_{j-1})\,\frac{w_j}{\hat{p}_j}], \quad (5)$$

where 1_A is the indicator function of set A, J is the number of histogram cells, $[l_j, r_j)$ is the support of cell j with width $w_j = r_j - l_j > 0$, $\hat{p}_j$ is the probability estimator of cell j and $\{\hat{C}_i\}_{i=0}^{J}$ is the cdf of $\{p_j\}_{j=1}^{J}$, i.e., $\hat{C}_j = \sum_{i=1}^{j}\hat{p}_i,\ 1\le j\le J$ ($\hat{C}_0=0$ and $\hat{C}_J=1$).

The rationale for using TES processes stems from the following facts. First, all TES background sequences are stationary Markovian, and their marginal distribution is uniform on $[0,1)$, *regardless* of the probability law of the innovations $\{V_n\}$ selected [7]. Second, The *inversion method* [2] permits us, in principle, to transform any uniform variate to others with arbitrary distributions as follows: if U is uniform on $[0,1)$ and F is a prescribed distribution, then $X = F^{-1}(U)$ has distribution F. In particular, $F = \hat{H}_Y$ is just a special case. And third, for $0<\xi<1$, the effect of S_ξ is to render the sample paths of background TES sequences more "continuous-looking". Because stitching transformations preserve uniformity [15], the inversion method can still be applied to stitched background processes $\{S_\xi(U_n)\}$. It follows that any foreground TES sequence of the form

$$X_n^+ \;=\; \hat{H}_Y^{-1}(S_\xi(U_n^+)), \qquad (6)$$

obtained from any background sequence, is always guaranteed to have the empirical (histogram) distribution $\hat{H}_Y$, *regardless* of the innovation density f_V and stitching parameter ξ selected. The choice of a pair (f_V,ξ) will determine the dependence structure of (2), and in particular (6), including the associated autocorrelation function. Thus, TES modeling decouples the fitting of the empirical distribution from the fitting of the empirical autocorrelation function. Since the former is guaranteed, one can concentrate on the latter.

In practice, TES modeling consists in a heuristic search for a suitable pair (f_V,ξ), such that the ensuing TES model approximates well the leading empirical autocorrelations and gives rise to sample path realizations which adequately resemble the empirical record. Effective TES modeling requires software support. To construct TES models that fit empirical data, we used TEStool — a visual interactive software environment [5], designed to support TES modeling. TEStool allows the user to read in empirical time series and calculate their statistics (histogram, autocorrelation function and spectral density) in textual and graphical representations. It provides services to construct and modify TES models and to superimpose the corresponding TES statistics on their empirical counterparts. The search proceeds in an interactive style, guided by visual feedback: each model modification triggers a recalculation and redisplay of the corresponding statistics. TES model autocorrelations and spectral densities are calculated numerically from fast and accurate formulas developed in [7, 8]. TEStool was used extensively to model compressed video; see, e.g., [13, 16, 17, 18] and references therein.

6.2 TES Models of Individual Scene Classes

The first-order properties of each individual class of scenes are encapsulated by the corresponding empirical histogram of bit rates. Figure 3 shows that all these histograms exhibit a bell-shape pattern.

To estimate the empirical autocorrelation of the bit rate process for each scene class, we used the estimator in [1]

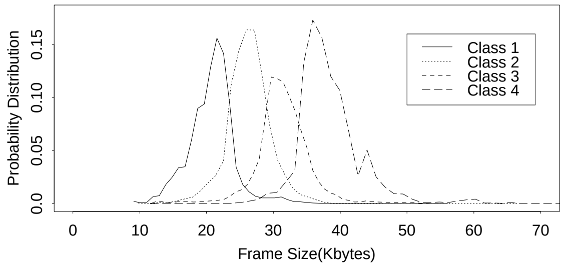

Figure 3: Empirical Histograms of Bit Rates for Individual Classes.

which for an empirical sequence, $\{Y_n\}_{n=0}^{N}$, is given by

$$\hat{\rho}(\tau) = \frac{1}{\hat{\sigma}_H\,\hat{\sigma}_T(N-\tau)} \sum_{j=0}^{N-1-\tau} (Y_j - \hat{m}_H)(Y_{j+\tau} - \hat{m}_T), \tag{7}$$

where

$$\hat{m}_H = \frac{1}{N-\tau}\sum_{j=0}^{N-1-\tau} Y_j, \quad \hat{\sigma}_H^2 = \frac{1}{N-\tau}\sum_{j=0}^{N-1-\tau} (Y_j - \hat{m}_H)^2,$$

$$\hat{m}_T = \frac{1}{N-\tau}\sum_{j=\tau}^{N-1} Y_j, \quad \hat{\sigma}_T^2 = \frac{1}{N-\tau}\sum_{j=\tau}^{N-1} (Y_j - \hat{m}_T)^2,$$

all depend on τ. Note that product terms in (7) include only samples from the *same* scene, since two successive scenes belonging to the same class are most likely not adjacent. The empirical autocorrealtion functions per scene class are depicted in Figure 4.

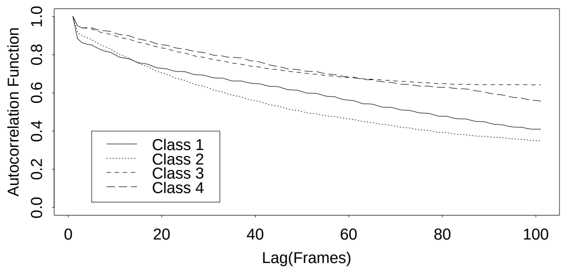

Figure 4: Empirical Autocorrelation Functions of Bit Rates for Individual classes.

Given the empirical histogram and autocorrelation function of bit rates for each scene class, we proceeded to model each as a TES$^+$ sequence, $\{X_n^+\}$, of the form (6), via the TES methodology described in Section 6.1.

6.3 Modeling Scene Durations

Recall that the MRMT model uses a modulating Markov chain to govern the sequence of scene classes. It remains to model scene durations as a stochastic sequence, to complete the overall "Star Wars" model.

Figure 5 displays the empirical distribution of scene durations (in number of frames) for each scene class. A visual inspection suggests that these distributions are quite similar, so that a class-independent duration model is a possibility. However, for greater modeling accuracy, a class-dependent duration model was constructed. We will elaborate on this point in the next section.

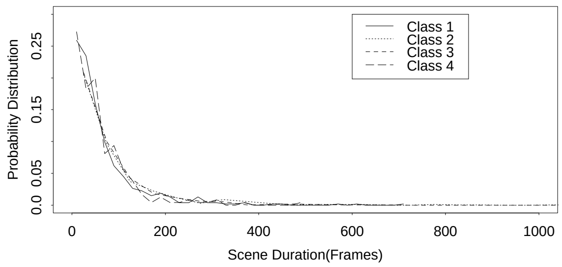

Figure 5: Empirical Distribution Functions of Durations for Individual classes.

Figure 6 displays the autocorrelation functions of scene durations for each scene class. The salient feature of this figure is the low magnitude of positive-lag autocorrelations for all classes. Indeed, for lags longer than 5, the autocorrelations practically vanish.

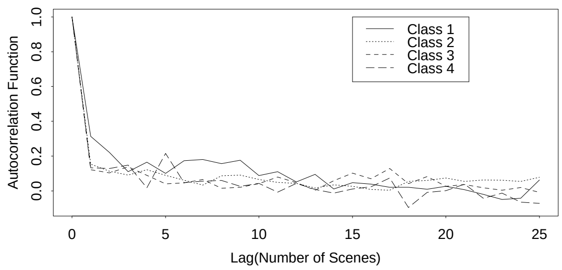

Figure 6: Empirical Autocorrelation Functions of Durations for Individual classes.

To simplify the overall model without sacrificing fidelity, we decided to model scene durations as class-dependent iid random variables (renewal sequences.)

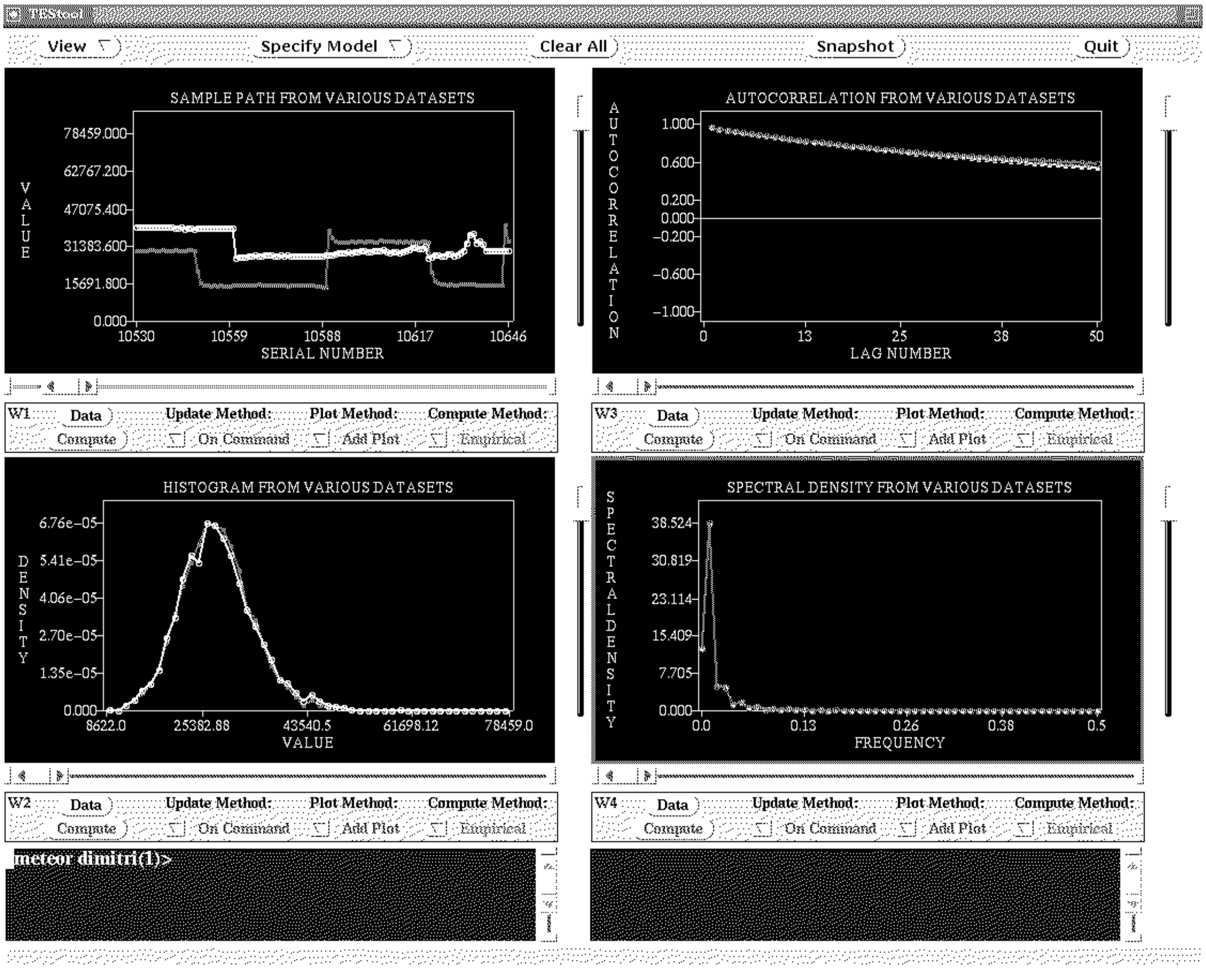

Figure 7: Comparison of Statistics of the Empirical and Simulated "Star-Wars" Bit Rate Process.

7 Results

In order to validate our model, we ran a Monte Carlo simulation of the model to generate a realization of the VBR process with the statistical signature of the original "Star Wars" compressed video. Simulation statistics were then compared to the original ones. Both the simulated time series and the original record consisted of 171,000 frames each. The simulation proceeds as follows. On startup, an initial scene class is generated from the stationary distribution of the modulating Markov chain, $\{M_n\}$ (see Section 5); this determines both the distribution of scene duration and the TES process used for this scene. The scene is then simulated and when its duration elapses, a new scene class is generated via the transition matrix, $\mathbf{P}$, of the modulating Markov chain, an independent duration is sampled, and so forth.

Figure 7 displays a screen image of TEStool, consisting of four canvases, each superimposing simulation-based statistics (bold curves) on their empirical counterparts (fainter curves). The upper-left canvas displays a subset of the original record superimposed on its simulation-based counterpart. The lower-left canvas depicts the empirical histogram and its simulation-based counterpart, computed from the data in the upper-left canvas. Similarly, the upper-right and lower-right canvases contain the corresponding autocorelation functions and spectral densities, respectively. Observe that all statistics exhibit excellent agreement between the empirical data and their simulation-based counterparts. In fact, the histogram, autocorrelation function and spectral density curves are barely distinguishable. In addition, the simulated time series bears a strong qualitative resemblance to the empirical record. Even though such a subjective judgment cannot be precisely quantified, the results shown in Figure 7 support the claim that our MRMT model fits the empirical record well, both quantitatively and qualitatively. Informally, the TES processes which model individual scenes capture short-term behavior, while the modulating Markov chain governing scene changes captures longer-term behavior.

Some additional remarks on modeling considerations are warranted at this point. First, the model is quite parsimonious in view of the fact that four scene classes appear to be the minimal number to yield a good fit. In particular, experimentation with fewer classes revealed deficiencies in the corresponding fit, while more classes brought no perceptible improvement (at the cost of higher complexity). Thus, the clustering analysis was extremely valuable in identifying a parsimonious model. Second, recall that Section 6.3 suggested that the distribution of scene durations appeared independent of scene classes. An initial attempt in this vein led to surprisingly poor fit to the empirical histogram. This problem disappeared, once scene durations were modeled as class-dependent renewal processes.

Our results suggest that the modeling methodology proposed here can be used to create parsimonious high-fidelity models of compressed video traffic, which capture the statistical signature of whole movies. Our goal is to create MRMT models for a suite of representative movies under a variety of coding algorithms. In particular, a TES-based model for MPEG coding is described in [18]; additional TES-based models are described in [13, 17] for other coding algorithms, including H.261. These can be used to drive realistic Monte Carlo simulations of broadband communications networks in an effective way.

Acknowledgments

The empirical data were provided to us courtesy of Mark Garrett, Bell Communications Research (Bellcore) and Martin Vetterli, University of California, Berkeley.

References

[1] Cox, D.R. and Lewis, P.A.W. *The Statistical Analysis of Series of Events*, Methuen, 1968.

[2] Devroye, L. *Non-Uniform Random Variate Generation*, Springer-Verlag, 1986.

[3] Frost, V. and Melamed, B. "Traffic Modeling for Telecommunications Networks", *Communications Magazine*, 1994 (to appear).

[4] Garrett, M.W. "Contributions Toward Real-Time Services on Packet Switched Networks", Ph.D. Dissertation, Columbia University, May 1993.

[5] Geist, D. and Melamed, B. "TEStool: An Environment for Visual Interactive Modeling of Autocorrelated Traffic", *Proceedings of the 1992 International Conference on Communications* (Chicago, Illinois), Vol. 3: 1285-1289, 1992.

[6] Heffes, H. and Lucantoni, D.M. "A Markov Modulated Characterization of Packetized Voice and Data and Related Statistical Multiplexer Performance", *IEEE J. on Selected Areas in Communications* Vol. SAC-4, 856–868, 1986.

[7] Jagerman, D.L. and Melamed, B. "The Transition and Autocorrelation structure of TES Processes Part I: General Theory". *Stochastic Models*, Vol. 8(2), 193-219, 1992.

[8] Jagerman, D.L. and Melamed, B. "The Transition and Autocorrelation structure of TES Processes Part II: Special Cases". *Stochastic Models*, Vol. 8(3), 499-527, 1992.

[9] Jain, A.K. *Algorithms for Clustering Data*, Prentice Hall, 1988.

[10] Latouche, G. "An Exponential Semi-Markov Process, With Applications to Queueing Theory", *Stochastic Models*, Vol. 1, No. 2, 137–169, 1985.

[11] "JPEG Technical Specification: Revision (Draft), Joint Photographic Experts Group, ISO/IEC JTC1/SC2/WG8, CCITT SGVIII", August, 1990.

[12] Lazar, A.A., Pacifici G. and Pendarakis, D.E. "Modeling of Video Sources for Real Time Scheduling", *Technical Report 324-93-03*, CTR, Columbia University, 1993.

[13] Lee, D.S., Melamed, B., Reibman, A. and Sengupta, B. "TES Modeling for Analysis of a Video Multiplexer", *Performance Evaluation*, Vol. 16, 21–34, 1992.

[14] Livny, M., Melamed B. and Tsiolis, A.K. "The Impact of Autocorrelation on Queuing Systems". *Management Science*, Vol. 39, 322–339, 1993.

[15] Melamed, B. "TES: A Class of Methods for Generating Autocorrelated Uniform Variates". *ORSA J. on Computing*, Vol. 3, No. 4, 317–329, 1991.

[16] Melamed, B., Raychaudhuri, D., Sengupta B. and Zdepski, J. "TES-Based Traffic Modeling for Performance Evaluation of Integrated Networks," *Proceedings of INFOCOM '92*, Vol. 1, 75–84, Florence, Italy, 1992.

[17] Melamed, B. and Sengupta, B. "TES Modeling of Video Traffic". *IEICE Transactions on Communications*, Vol. E75-B, No. 12, 1292–1300, 1992.

[18] Reininger, D., Melamed, B. and Raychaudhuri, D. "Variable Bit Rate MPEG Video: Characterization, Modeling and Multiplexing", *Proceedings of the 14-th International Teletraffic Congress*, Antibes Juan-les-Pins, France, Vol. 1a, 295–306, June 1994.