Analytical Investigation of the Bias Effect in Variance-Type Estimators for Inference of Long-Range Dependence*

Marwan Krunz† and Ibrahim Matta‡
† Department of Electrical & Computer Engineering
University of Arizona, Tucson, AZ 85721
Email: krunz@ece.arizona.edu

‡ Department of Computer Science
Boston University, Boston, MA 02215
Email: matta@cs.bu.edu

Last Revised: March 8, 2002

Abstract

Since the publication of the Bellcore measurements in the early nineties, long-range dependence (LRD) has been in the center of a continuous debate within the teletraffic community. While researchers largely acknowledge the significance of the LRD phenomenon, they still disagree on two issues: (1) the utility of LRD models in buffer dimensioning and bandwidth allocation, and (2) the ability of commonly used statistical tools to differentiate between true LRD and other potential interpretations of it (e.g., non-stationarity). This paper is related to the second issue. More specifically, our objective is to *analytically* demonstrate the limitations of variance-type indicators of LRD. Our work is not meant to advocate a particular modeling philosophy (be it LRD or SRD), but rather to emphasize the potential misidentification caused by the inherent bias in variance-type estimators. Such misidentification could lead one to wrongly conclude the presence of an LRD structure in a sequence that is known to be SRD. Our approach is based on deriving simple analytical expressions for the slope of the aggregated variance in three autocorrelated traffic models: a class of SRD (but non-Markovian) $M/G/\infty$ processes, the discrete autoregressive of order one model (SRD Markovian), and the fractional ARIMA process (LRD). Our main result is that a variance-type estimator often indicates, falsely, the existence of an LRD structure (i.e., $H > 0.5$) in synthetically generated traces from the two SRD models. The bias in this estimator, however, diminishes monotonically with the length of the trace. We provide some guidelines on selecting the minimum trace length so that the bias is negligible. We also contrast the VT estimator with three other estimation techniques.

*An abridged version of this paper appeared in the *Proceedings of the IEEE INFOCOM 2001 Conference*, Anchorage, Alaska, April 2001. The work of M. Krunz was supported by the National Science Foundation under grants ANI 9733143, CCR 9979310, and ANI 0095626. The work of I. Matta was supported by the National Science Foundation under grant ANI 0095988.

1 Introduction

During the last decade, several studies supported by extensive measurements concluded the likely presence of long-range dependence (LRD)[1] in various types of network traffic, including LAN [9, 18], WAN [25], variable-bite-rate (VBR) video traffic [4, 11, 14], and WWW traffic [6]. Yet, several aspects of the LRD phenomenon are still being debated within the research community. One of these aspects is related to the relevance of traffic correlations to the dimensioning of network resources (buffer and bandwidth). While researchers, in general, agree on the importance of correlations, they still disagree on how much of these correlations should be captured in a *practical* traffic model [12, 13, 29]. Earlier traffic models are Markovian in nature, with an autocorrelation function (ACF) that drops off exponentially. Examples of these models are the autoregressive moving-average (ARMA) models, Markov Arrival processes (MAP), Markov modulated processes, etc. (see [10, 2, 19] for surveys of traffic models). Markovian models exhibit an exponentially decaying autocorrelation structure, which makes them *short-range dependent* (SRD). An SRD model is one for which the ACF is summable, i.e., $\sum_k \rho_k < \infty$. Note, however, that an SRD model is not necessarily Markovian. In fact, several types of non-Markovian SRD models have been recently studied in the literature, including the $M/G/\infty$ process [23, 16] and a class of (subexponential) path-based Markov renewal processes [15, 17]. The interest in such models has to do with their ability to produce a wide range of correlation structures, including, as extreme cases, both Markovian and LRD structures. LRD models on the other hand exhibit a slowly decaying ACF (typically as a power function) to the extent that the correlations now have an infinite sum, i.e., $\sum_k \rho_k = \infty$.

Foremost, the statistical evidence supporting LRD was based on the estimated value of the Hurst (H) parameter. This parameter, in fact, gauges the self-similarity, rather than the LRD, of a process. However, it is common practice to talk about H as a measure of LRD, since the "derivative" of a second-order self-similar process (i.e., the "rate" process) exhibits an LRD structure when the H value of the "parent" process is between 0.5 and 1. Several statistical techniques have been used to estimate the H parameter, including [3]: the variance-time (VT) test, the R/S statistic, periodogram-based analysis, and more recently wavelet-based techniques [32, 1, 28]. In this paper, we focus on the VT test, which is one of the main tests used to discover the LRD phenomenon in network traffic. We show that such a test can give misleading indication about the true SRD/LRD nature of a time series, despite the availability of many data points. Although the bias in this test has been previously examined (albeit, empirically) and words of caution on its use have been reiterated [7, 30], it is still being applied in traffic analysis and modeling. One reason for that is that previous investigations of the VT test primarily focused on its bias in the presence of nonstationarities or trends (e.g., [31, 21, 20, 27]), giving the impression that the VT test does not necessarily suffer from any problems if the traffic is stationary. Our goal in this paper is to take the bias issue one step further and prove its presence even under stationary models.

In [21] the authors investigated the inherent difficulty associated with estimating the Hurst parameter. They found that the estimated H value depends on several factors, including the sample size, the time scale, and the level shifts (jumps in the mean). The effects of certain types of nonstationarity on the estimation of the H parameter was studied in [31] for variance-type estimators. It was suggested that using the first

[1]Throughout the paper, we use 'LRD' both as a noun and as an adjective.

difference of the variance, rather than the variance itself, in the VT test can help distinguish between LRD and nonstationarities due to slow deterministic trends or level shifts. Other types of nonstationarity (e.g., polynomial trends) were investigated in [20] for several LRD tests. For the VT test, the authors in [20] found that while the VT test is accurate when applied to a pure LRD process, namely the fractional Gaussian noise (FGN), it gives misleading results when FGN is contaminated by certain types of nonstationarity. In [27] the authors investigated the robustness of the wavelet-based Abry-Veitch estimator [32] against nonstationarities in the mean and/or variance of an LRD process. Our work is distinct from previous works in two key aspects. First, we are particularly interested in distinguishing between LRD and SRD behaviors *under the assumption of stationarity.* Second, we use *analytical* arguments to investigate the inherent bias in variance-type estimators. All of the processes considered in this paper are stationarity, and they do not exhibit any deterministic trends. We provide analytical expressions for the bias in three popular processes with different correlation structures: $M/G/\infty$ process (SRD but non-Markovian), fractional ARIMA (LRD), and the discrete autoregressive of order one model (SRD Markovian). Focusing on three specific models with contrasting autocorrelation behaviors does not diminish the impact of our results since, in fact, the analysis uses only the autocorrelation structures of these models (i.e., does not depend on the marginal distributions). So one can generalize the treatment to other models with known autocorrelation structures. It should be emphasized that our work is not meant to advocate one model over another, but to illustrate the caveats in using the VT test for inference of LRD and to provide guidelines on the required number of data points for which the test is credible.

The rest of the paper is structured as follows. In Section 2 we briefly describe the three processes that are used in our study. The aggregated variance for each of these processes is derived in Section 3. In the same section, we discuss the limitations of the VT test. In Section 4 we contrast the VT test with three other estimators, namely the Abry-Veitch wavelet estimator, the R/S statistic, and the absolute-value method. The paper is concluded in Section 5.

2 Autocorrelated Processes

2.1 $M/G/\infty$ Input Process

The $M/G/\infty$ process is the busy-server process of a discrete-time $M/G/\infty$ queue. It can be constructed as follows (see [16, 23] for details). Start with a discrete-time $M/G/\infty$ queue. During time slot $[n, n+1)$ ($n = 0, 1, \ldots$), ξ_{n+1} new customers arrive into the system. Customer j, $j = 1, \ldots, \xi_{n+1}$, is presented to its own server, which begins its service by the start of slot $[n+1, n+2)$, with a service time $\sigma_{n+1,j}$ (in number of slots). Let b_n denote the number of busy servers, or equivalently, the number of customers present in the system at the beginning of time slot $[n, n+1)$, with b_0 being the initial number of customers present in the system. It is assumed that the $\mathbb{N}$–valued random variables (rvs) b_0, $\{\xi_{n+1},\ n = 0, 1, \ldots\}$, $\{\sigma_{n,j},\ n = 1, 2, \ldots;\ j = 1, 2, \ldots\}$ and $\{\sigma_{0,j},\ j = 1, 2, \ldots\}$ satisfy the following assumptions: (i) they are mutually independent; (ii) $\{\xi_{n+1},\ n = 0, 1, \ldots\}$ are *i.i.d.* Poisson rvs with parameter $\lambda > 0$; (iii) $\{\sigma_{n,j},\ n = 1, \ldots;\ j = 1, 2, \ldots\}$ are *i.i.d.* rvs with common pmf G on $\{1, 2, \ldots\}$. Let σ be a generic $\mathbb{N}$–valued rv distributed according to the pmf G; assume that $\mathbf{E}[\sigma] < \infty$. Then, the $M/G/\infty$ input process is simply the busy-server process $\{b_n,\ n = 0, 1, \ldots\}$.

Although $\{b_n,\ n = 0, 1, \ldots\}$ is in general *not* a (strictly) stationary process, it does admit a stationary and ergodic version, $\{b_n^\star,\ n = 0, 1, \ldots\}$, that can be constructed by taking: (i) b_0 to be Poisson distributed with parameter $\lambda \mathbf{E}[\sigma]$; (ii) $\{\sigma_{0,j},\ j = 1, 2, \ldots\}$ to be *i.i.d.* rvs distributed according to the *forward recurrence time* $\hat{\sigma}$ associated with σ. The pmf of $\hat{\sigma}$ is given by

$$\mathbf{P}[\hat{\sigma} = r] \stackrel{\text{def}}{=} \frac{\mathbf{P}[\sigma \geq r]}{\mathbf{E}[\sigma]}, \quad r = 1, 2, \ldots \tag{1}$$

Based on the above construction, several useful properties of the stationary version $\{b_n^\star,\ n = 0, 1, \ldots\}$ are readily obtained [22]:

(i) For each $n = 0, 1, \ldots$, the rv $b_n^\star$ is a Poisson rv with parameter $\lambda \mathbf{E}[\sigma]$;

(ii) The ACF of $\{b_n^\star,\ n = 0, 1, \ldots\}$ is given by

$$\rho_k = \mathbf{P}[\hat{\sigma} > k], \quad k = 0, 1, \ldots \tag{2}$$

By varying G, the process $\{b_n^*,\ n = 0, 1, \ldots\}$ can display various forms of positive autocorrelations, the extent of which is controlled by the tail behavior of G.

To close this section, we point out that the process $\{b_n^\star,\ n = 0, 1, \ldots\}$ can induce *both* SRD and LRD behaviors: From (2), it follows readily [24] that

$$\sum_{k=0}^{\infty} \rho_k = \mathbf{E}[\hat{\sigma}] = \frac{1}{2} + \frac{\mathbf{E}[\sigma^2]}{2\mathbf{E}[\sigma]}. \tag{3}$$

Consequently, the process $\{b_n^*,\ n = 0, 1, \ldots\}$ is LRD (resp. SRD) *if and only if* $\mathbf{E}[\sigma^2]$ is infinite (resp. finite). In particular, the $M/G/\infty$ input traffic will be LRD when G is Pareto, with a shape parameter in the interval $(1, 2)$ [5].

2.2 Discrete Autoregressive of Order One Process

The DAR(1) process is a popular Markovian (hence, SRD) model that has been used to characterize video teleconferencing traffic [8]. This process can exhibit any arbitrary marginal distribution. Its autocorrelation structure is similar to that of the common AR(1) process. To generate a DAR(1) process, we start with two mutually independent random sequences $\{V_n : n = 1, 2, \ldots\}$ and $\{Y_n : n = 1, 2, \ldots\}$. The sample space for $\{V_n : n = 1, 2, \ldots\}$ is $\{0, 1\}$, and its marginal distribution is given by:

$$\Pr[V_n = i] = \begin{cases} r, & \text{if } i = 1 \\ 1 - r, & \text{if } i = 0 \end{cases}$$

for $n = 1, 2, \ldots$. The process $\{Y_n : n = 1, 2, \ldots\}$ is renewal with an arbitrary but countable sample space $\mathcal{S}_Y$. Its marginal distribution is defined by:

$$\Pr[Y_n = i] \stackrel{\text{def}}{=} \pi_i, \text{ for all } i \in \mathcal{S}_Y$$

Then, the DAR(1) process $\{X_n : n = 1, 2, \ldots\}$ is defined as follows:

$$X_n = V_n X_{n-1} + (1 - V_n) Y_n, \; n = 1, 2, \ldots \tag{4}$$

It is easy to show that $\{X_n : n = 1, 2, \ldots\}$ constitutes a Markov chain with an autocorrelation structure of the form $\rho_k = r^k$ for $k = 0, 1, \ldots$.

2.3 Fractional ARIMA Process

The last process that we will examine is the popular fractional ARIMA$(0, d, 0)$ process. This LRD Gaussian process was proposed as a basis for modeling VBR video traffic [11]. Its ACF is given by

$$\rho_k = \frac{d(1+d)\cdots(k-1+d)}{(1-d)(2-d)\cdots(k-d)}, \quad k = 1, 2, \ldots \tag{5}$$

where $0 < d < 0.5$ is the *fractional differencing parameter* given by $d = H - 1/2$. As $k \to \infty$, ρ_k behaves as $k^{-\alpha}$, where $\alpha = 2 - 2H$. See [11] for details on how to generate synthetic F-ARIMA traces.

3 Analysis of Aggregated Variance

Consider a second-order stationary process $\{X_n : n = 1, 2, \ldots\}$ with mean $\overline{X}$ and variance v. Let $C_k \stackrel{\text{def}}{=} \operatorname{cov}(X_n, X_{n+k}) = E\left[(X_n - \overline{X})(X_{n+k} - \overline{X})\right]$. The ACF is defined as $\rho_k \stackrel{\text{def}}{=} C_k/v$, for $k = 0, 1, \ldots$. For $m = 1, 2, \ldots$, let

$$X_n^{(m)} \stackrel{\text{def}}{=} \frac{\sum_{i=nm-m+1}^{nm} X_i}{m}, \quad n = 1, 2, \ldots \tag{6}$$

so that $\{X_n^{(m)}\}$ is an averaged version of $\{X_n\}$, with the averaging taken over non-overlapping blocks of length m. The variance of the new time series is given by:

$$v_m \stackrel{\text{def}}{=} \operatorname{var}(X_n^{(m)}) = \frac{v}{m} + \frac{2}{m^2} \sum_{p=1}^{m-1} \sum_{q=1}^{p} C_q \tag{7}$$

We will refer to v_m as the aggregated variance at level m. If $\{X_n\}$ is an LRD process, then it must satisfy $m v_m \to \infty$ as $m \to \infty$. More specifically, for an LRD process $v_m \sim m^{-\alpha}$ when m is large, where $0 < \alpha < 1$ is the same parameter defined above. For an SRD process, $\alpha \geq 1$. To test whether a given time series is LRD or not, the empirical VT test proceeds by plotting $\log(v_m/v)$ versus $\log m$ for various aggregation levels m. The asymptotic slope of the plot is then taken as an estimate of $-\alpha$. If $\alpha < 1$, the empirical sequence is believed to exhibit LRD. As an example, the VT plot for the *Star Wars* trace is shown in Figure 1. Its asymptotic slope, ignoring aggregation levels smaller than 100, is estimated by least-square method to be 0.43, roughly in agreement with the numbers in [4] and [11].

Next, we obtain analytical expressions for the slope of the VT plot in the three examined processes.

3.1 Aggregated Variance in the $M/G/\infty$ Model

Let $\{X_n(N) : n = 1, 2, \ldots\}$ be a subclass of $M/G/\infty$ processes that is parameterized by N and that possesses the following ACF form:

$$\rho_k = e^{-\beta \sqrt[N]{k}}, \quad k = 0, 1, 2, \ldots \tag{8}$$

where $\beta > 0$ and $N = 1, 2, 3, \ldots$. For fixed β and N, it is easy to see that $\{X_n(N) : n = 1, 2, \ldots\}$ is an SRD process, since

$$\sum_{k=0}^{\infty} \rho_k = 1 + \sum_{k=1}^{\infty} \rho_k \le 1 + \int_0^{\infty} e^{-\beta \sqrt[N]{t}} dt = 1 + \frac{N!}{\beta^N} < \infty \tag{9}$$

Our interest in this subclass of $M/G/\infty$ processes stems from the fact that it offers a rich spectrum of autocorrelation forms, which on the two extremes give rise to Markovian ($N = 1$) and LRD ($N = \infty$) forms. Naturally, the regime in the middle ($2 \le N < \infty$) is of particular interest to us. The ACF in (8) results in a Weibull-like service-time distribution (G).

Now consider the aggregated process $\{X_n^{(m)}(N) : n = 1, 2, \ldots\}$ for $m = 1, 2, \ldots$. The *normalized* aggregated variance of this process can be written as [5]:

$$\widetilde{v}_m \stackrel{\text{def}}{=} \frac{v_m}{v_1} = \frac{1}{m} + \frac{2}{m^2} \sum_{k=1}^{m} (m-k)\rho_k \tag{10}$$

Since m is discrete, the instantaneous slope of the curve that describes $\log \widetilde{v}_m$ as a function of $\log m$ is defined by the first difference:

$$s_m \stackrel{\text{def}}{=} \frac{\log \widetilde{v}_{m+1} - \log \widetilde{v}_m}{\log(m+1) - \log m} \tag{11}$$

Without loss of generality, we assume that all logarithms are to the base ten. Note that in the empirical VT test, s_m is replaced by its average value that is obtained using least square fitting. We now derive almost exact expressions for $\widetilde{v}_m$ and s_m. To do that, we allow m to take any nonnegative real value. To distinguish it from its discrete-time counterpart, we indicate the variance of the aggregated series in the continuous case by $\widetilde{v^*}$, which is given by [5]:

$$\widetilde{v^*}_m = \frac{2}{m^2} \int_0^m (m-h)\rho_h dh = \frac{2}{m^2} \int_0^m (m-h) e^{-\beta \sqrt[N]{h}} dh \tag{12}$$

Equation (12) can be written as follows:

$$\widetilde{v^*}_m = \frac{2}{m^2} \left[m \int_0^m e^{-\beta \sqrt[N]{t}} dt - \int_0^m t e^{-\beta \sqrt[N]{t}} dt \right] \tag{13}$$

Consider the two indefinite integrals $\int e^{-\beta \sqrt[N]{t}} dt$ and $\int t e^{-\beta \sqrt[N]{t}} dt$. Their solutions are given by:

$$\int e^{-\beta \sqrt[N]{t}} dt = -e^{-\beta \sqrt[N]{t}} N! \sum_{k=1}^{N} \frac{t^{(N-k)/N}}{\beta^k (N-k)!} \tag{14}$$

$$\int t e^{-\beta \sqrt[N]{t}} dt = -N t^2 \left(\beta t^{1/N}\right)^{-2N} \Gamma(2N, \beta t^{1/N}) \tag{15}$$

where $\Gamma(x,y)$ is the incomplete Gamma function, defined as:

$$\Gamma(x,y) \stackrel{\text{def}}{=} \int_y^\infty u^{x-1}e^{-u}du \tag{16}$$

Accordingly,

$$\begin{aligned}\Gamma(2N,\beta t^{1/N}) &= \int_{\beta t^{1/N}}^\infty u^{2N-1}e^{-u}du \\ &= e^{-\beta t^{1/N}}\left((\beta t^{1/N})^{2N-1} + (2N-1)(\beta t^{1/N})^{2N-2} + \cdots + (2N-1)!\right)\end{aligned} \tag{17}$$

Substituting the last expression for $\Gamma(2N,\beta t^{1/N})$ in (15), and after some manipulations, we end up with

$$\int te^{-\beta \sqrt[N]{t}}dt = -N(2N-1)!e^{-\beta \sqrt[N]{t}}\sum_{k=1}^{2N}\frac{t^{(2N-k)/N}}{\beta^k(2N-k)!} \tag{18}$$

From (14) and (18), we obtain the solution for the two *definite* integrals in (13):

$$\int_0^m e^{-\beta \sqrt[N]{t}}dt = \left(-N!e^{-\beta \sqrt[N]{m}}\sum_{k=1}^{N}\frac{m^{(N-k)/N}}{\beta^k(N-k)!}\right) + \frac{N!}{\beta^N} \tag{19}$$

$$\int_0^m te^{-\beta \sqrt[N]{t}}dt = \left(-N(2N-1)!e^{-\beta \sqrt[N]{m}}\sum_{k=1}^{2N}\frac{m^{(2N-k)/N}}{\beta^k(2N-k)!}\right) + \frac{N(2N-1)!}{\beta^{2N}} \tag{20}$$

Using the above two equations, $\widetilde{v^*}_m$ in (13) can be written as:

$$\widetilde{v^*}_m = e^{-\beta \sqrt[N]{m}}\left((2N)!\sum_{k=1}^{2N}\frac{m^{-k/N}}{\beta^k(2N-k)!} - 2(N!)\sum_{k=1}^{N}\frac{m^{-k/N}}{\beta^k(N-k)!}\right) + \frac{2(N!)}{m\beta^N} - \frac{(2N)!}{m^2\beta^{2N}} \tag{21}$$

Note that as $m\to\infty$, $\widetilde{v^*}_m \sim 2(N!)/(m\beta^N) \sim \mathcal{O}(1/m)$, as expected (since the process is SRD). Figure 2 depicts $\widetilde{v^*}_m$, obtained using (21), versus m when $\rho_k = e^{-\beta\sqrt{k}}$ ($N=2$). This special case of the $M/G/\infty$ process has been shown to accurately characterize the behavior of variable-bit-rate video [16]. The plot depicts clear convexity, which is the source of the bias in the VT test. Figure 3 depicts $\widetilde{v^*}_m$ for various values of N with $\beta = 0.06$. We found these plots almost indistinguishable from their discrete counterparts (not shown) obtained directly from (10). It is interesting to note that $\widetilde{v^*}_m$ depends on β and m only through the term $\tilde{m} \stackrel{\text{def}}{=} \beta\sqrt[N]{m}$. More specifically, (21) can be rewritten as:

$$\widetilde{v^*}_m = e^{-\tilde{m}}\left((2N)!\sum_{k=2}^{2N}\frac{1}{\tilde{m}^k(2N-k)!} - 2(N!)\sum_{k=2}^{N}\frac{1}{\tilde{m}^k(N-k)!}\right) + \frac{2(N!)}{\tilde{m}^N} - \frac{(2N)!}{\tilde{m}^{2N}} \tag{22}$$

Now that we have obtained an expression for $\widetilde{v^*}_m$, we proceed to derive the instantaneous slope of $\widetilde{v^*}_m$, which is defined as follows:

$$s^*_m \stackrel{\text{def}}{=} \frac{d(\log\widetilde{v^*}_m)}{d(\log m)} = \frac{m}{\widetilde{v^*}_m}\frac{d\widetilde{v^*}_m}{dm} \tag{23}$$

With some basic algebraic manipulations, it can be shown that:

$$s_m^* = \left[e^{-\widetilde{m}} \left(\frac{-(2N)!}{N} \sum_{k=2}^{2N} \frac{\widetilde{m}+k}{\widetilde{m}^k(2N-k)!} + 2(N-1)! \sum_{k=2}^{N} \frac{\widetilde{m}+k}{\widetilde{m}^k(N-k)!} \right) - \frac{2(N!)}{\widetilde{m}^N} + \frac{2(2N)!}{\widetilde{m}^{2N}} \right] / \widetilde{v^*}_m \tag{24}$$

As $m \to \infty$, $s_m^* \to -1$, as expected. From the concavity of $\widetilde{v^*}_m$, it readily follows that

$$\mid s_m^* \mid < \mid s_m \mid < \mid s_{m+1}^* \mid \tag{25}$$

In the special case of $N = 2$, (24) reduces to:

$$s^*_{m|_{N=2}} = \frac{e^{-\widetilde{m}} \left[\frac{-20}{\widetilde{m}^2} - \frac{48}{\widetilde{m}^3} - \frac{48}{\widetilde{m}^4} - \frac{4}{\widetilde{m}} \right] - \frac{4}{\widetilde{m}^2} + \frac{48}{\widetilde{m}^4}}{e^{-\widetilde{m}} \left[\frac{8}{\widetilde{m}^2} + \frac{24}{\widetilde{m}^3} + \frac{24}{\widetilde{m}^4} \right] + \frac{4}{\widetilde{m}^2} - \frac{24}{\widetilde{m}^4}} \tag{26}$$

Figure 4 depicts $-s_m^*$ versus the aggregation level m for four values of N ($\beta = 0.06$). For $N > 1$, s_m^* converges very slowly to -1. In fact, even at an aggregation level of $m = 8000$ and $N = 2$, $\mid s_m^* \mid$ is still smaller than 0.8. The speed of convergence decreases rapidly as N increases. For example, when $N = 4$, even at an aggregation level as high as 8000, $\mid s_m^* \mid$ does not exceed 0.1, far from indicating any SRD structure. Had we not known in advance that the underlying process is SRD, we would have mistakenly decided (based on the VT test) that the data exhibit LRD behavior. Figure 5 depicts the impact of various values of β. Clearly, the larger the value of β, the faster is the convergence of s_m^* to -1.

The plot of $-s_m^*$ versus $\widetilde{m}$ is shown in Figure 6 for $N = 2$. From this figure, it can been seen that the absolute value of the slope of the analytically obtained VT plot is *always* less than one for a finite $\widetilde{m}$. This critical observation implies that *when applied to traces of an SRD $M/G/\infty$ process with $\rho_k = e^{-\beta \sqrt[N]{k}}$, the VT test will always indicate, wrongly, the presence of an LRD structure irrespective of the length of these traces.* Only when such traces are of infinite length, the slope of the VT plot will be -1. If for the sake of empirical approximation, one is to take $\mid s_m^* \mid \geq 0.95$ as an indication of SRD, then in the case of $N = 2$ we must have $\widetilde{m} \geq 11.2$. If $\beta = 0.05$ (which is a typical value for video sequences fitted using an $M/G/\infty$ model with $\rho_k = e^{-\beta\sqrt{k}}$ [16]), then we need at least 50176 data points to correctly infer that the data exhibit SRD.

3.2 Aggregated Variance in the DAR(1) Model

Next, we consider the DAR(1) process. Substituting the expression for the ACF, $\rho_k = r^k$, in (10), and after some straightforward manipulations, we obtain:

$$\widetilde{v}_m = \frac{1}{m} + \frac{2}{m^2} \left(\frac{r(r^m - mr + m - 1)}{(r-1)^2} \right) \tag{27}$$

As $m \to \infty$, $\widetilde{v}_m \sim (1 + 2r/(1-r))/m$, which is, as expected, $\mathcal{O}(1/m)$. By substituting the values for $\widetilde{v}_m$ and $\widetilde{v}_{m+1}$ in (11), we can plot the first-order difference s_m versus m, as shown in Figure 7. Note that when r is close to one, the convergence of s_m to -1 becomes very slow. We will come back to this issue later in this section.

Next, we provide a closed-form expression for s_m^*, the continuous version of s_m, which was defined in

(23). By substituting $\rho_h = r^h$ in (12) and after some manipulations, we arrive at the following expression for $\widetilde{v^*}_m$:

$$\widetilde{v^*}_m = \frac{2}{m^2}\left[\frac{r^m - m\ln r - 1}{(\ln r)^2}\right] \tag{28}$$

where ln(.) is the natural logarithm. As $m \to \infty$ (with $r < 1$), $\widetilde{v^*}_m \sim -2/(m \ln r)$, which is $\mathcal{O}(1/m)$. As in the case of the $M/G/\infty$ model, the VT plots for the DAR(1) model in the continuous case are almost indistinguishable from their discrete-parameter counterparts. For brevity, we only show the plots in the continuous case (Figure 8).

Differentiating $\widetilde{v^*}_m$ in (28) with respect to m, we obtain

$$\frac{d\widetilde{v^*}_m}{dm} = \frac{2}{(\ln r)^2}\frac{(m\ln r - 2)r^m + m\ln r + 2}{m^3}$$

Hence, from (23) s^*_m for the DAR(1) model is given by

$$s^*_m = \frac{(m\ln r - 2)r^m + m\ln r + 2}{r^m - m\ln r - 1} \tag{29}$$

As $m \to \infty$, $s^*_m \to -1$, as expected. The speed of convergence of s^*_m is this case is rather fast due to the fast decay of the geometric terms in (29). To get an idea about how many data points are sufficient to infer SRD/LRD, we first rewrite (29) in terms of the variable $x \stackrel{\text{def}}{=} r^m$ as follows:

$$s^*_m = \frac{(\ln x - 2)x + \ln x + 2}{x - \ln x - 1} \tag{30}$$

Figure 9 depicts the plot of s^*_m as a function of x. As $x \to 0$, s^*_m converges to -1. However, as $x \to 1$, s^*_m approaches zero! So the utility of the VT test as an indicator of the SRD structure of the DAR(1) model, or any Markovian model to that extent, depends on the value of $x = r^m$. For a fixed $r < 1$, the number of points in a Markov-based trace must be large enough to ensure a sufficiently large m, so that r^m is close to zero. For example, to ensure that $|s^*_m| \geq 0.95$, we must have $m \geq -20.95/\ln r$. In this case, if $r = 0.999$, then we need an aggregation level $m \geq 20936$ (i.e., about 21,000 points *per block*). The size of the data trace should be at least ten times this number to give a meaningful sample estimation of the variance $\widetilde{v^*}_m$.

3.3 Aggregated Variance in the F-ARIMA Model

Finally, consider the F-ARIMA process described before. We first provide a simple recursive approach for computing $\tilde{v}_m$ for this process. First, we define the sums $X_m \stackrel{\text{def}}{=} \sum_{k=1}^{m} \rho_k$ and $Y_m \stackrel{\text{def}}{=} \sum_{k=1}^{m} k\rho_k$, for $m \geq 1$. Equation 10 can now be written as follows:

$$\tilde{v}_m = \frac{1}{m} + \frac{2}{m^2}(mX_m - Y_m). \tag{31}$$

Since $X_m = X_{m-1} + \rho_m$, $Y_m = Y_{m-1} + m\rho_m$, and $\rho_m = (m-1+d)/(m-d)\rho_{m-1}$, (31) can be computed recursively starting from $X_1 = Y_1 = \rho_1 = d/(1-d)$. Figure 10 depicts the VT plots for various values of d. It is interesting to note the linearity of the plots, with slopes that barely change with the aggregation level. (Contrast these plots with their nonlinear counterparts in Figure 2 and 8 for the $M/G/\infty$ and the

DAR(1) models, respectively). Moreover, these plots seem to be distinctly different from the empirical VT plot for the original *Star Wars* trace (Figure 1). This says that from an aggregated variance standpoint, the $M/G/\infty$ model (non-Markovian SRD) is more appropriate than the F-ARIMA model (LRD) in characterizing the JPEG-coded *Star Wars* sequence. The slope of the VT plot for the F-ARIMA model is shown in Figure 11 as a function of m (obtained using (11)).

So far, we have examined the behavior of the aggregated variance analytically, without involving any statistical estimation. One may question whether the trends observed in the previous figures still hold when the *empirical* VT test is used. To verify this point, we applied the empirical VT test to synthetic traces from the $M/G/\infty$ and F-ARIMA models. Figures 12 and 13 depict the results for two representative traces. For the $M/G/\infty$ trace, we set $\beta = 0.076$ and $N = 2$, which give a good fit for the empirical ACF of the *Star Wars* JPEG-coded sequence [16]. For the F-ARIMA trace, we took $d = 0.3$. The $M/G/\infty$ trace consists of 1,000,000 data points, while the F-ARIMA has 500,000 points (the computational complexity involved in generating $M/G/\infty$ traces grows linearly with the trace length, while this complexity grows quadratically in the case of F-ARIMA traces). Figure 12 indicates asymptotic slopes of -0.79 and -0.75 for aggregation levels in the ranges $[10^{3.5}, 10^{4.5}]$ and $[10^4, 10^5]$, respectively. This would suggest that the underlying data exhibit LRD. However, we know that the data were generated from an SRD $M/G/\infty$ model! Despite the length of the $M/G/\infty$ trace, the VT test may wrongly suggest the presence of LRD in this trace (if one is not careful in interpretting its outcome). Note that the concavity of the VT plot (which is predicted from the analysis) is not so apparent at large values of m, mainly because of the statistical inaccuracy in estimating v_m. Hence, it would be difficult to simply rely on visual inspection to determine the inappropriateness of the VT test by monitoring the concavity in the empirical VT plot. For the F-ARIMA trace (Figure 13), the slope of the VT plot is estimated at -0.58. Although v_m (also, the ACF) of a F-ARIMA is expected to behave as $k^{-0.4}$ when $k \to \infty$, it takes extremely long time to reach this asymptotic behavior. Figure 14 depicts the ACFs for the two tested traces along with the ACF for the real *Star Wars* trace.

4 Comparison with Other LRD Tests

We now contrast the VT test with three other methods: the wavelet-based Abry-Veitch (AV) test [32], the R/S method, and the absolute-value method (see [30] for descriptions of the latter two). The AV method is a semi-parametric estimator that is based on the discrete wavelet transform. It estimates both parameters α and c_f that appear in the spectral representation of an LRD process at the origin; namely, $f(v) \sim c_f|v|^{-\alpha}$ as $v \to 0$ (H is related to α through $H = (1 + \alpha)/2$). The method enjoys significant computational advantages along with an amenability for real-time implementation [28]. It is known to provide less biased estimates than traditional tests, although the bias has been studied mainly within the LRD regime (and for real traces). We used the Matlab code written by D. Veitch (see http://www.emulab.ee.mu.oz.au/~darryl) to apply the AV test to a 1,000,000-point-long SRD trace generated from an $M/G/\infty$ process with $\rho_k = e^{-0.076\sqrt{k}}$ (the same trace that was used in Figure 12). Figure 15 depicts the *Logscale Diagram* (LD), plotted over the complete range of available scales (the 95% confidence intervals are indicated by the short bars). The LD is essentially a log-log plot of the estimated variances of the wavelet details. For each scale j, the variance of the details (μ_j) is estimated as the average of the squared details at that scale.

So the LD is a plot of $y_j \stackrel{\text{def}}{=} \log_2 \mu_j$ versus j. An estimate of α is given by the slope of a weighted linear regression of the y_j's over an "appropriate" range of scales (j_1, j_2). To choose the optimal j_1, call it j_1^*, the Matlab code produces a graph of the goodness-of-fit function, $Q(j_1)$, versus j_1 with j_2 fixed at its maximum possible value. An example of such a graph is shown in Figure 16. For small scales, $Q(j_1)$ increases with j_1, indicating a better fit. This trend continues up to the scale j_1^*, after which $Q(j_1)$ barely changes with j_1. Regression is performed over the scales from j_1^* up to the maximum available scale. Based on this regression, $H \approx 0.645$ with a 95% confidence interval [0.566,0.725]. This says that the test indicates, wrongly, the presence of LRD in the $M/G/\infty$ trace. In other words, *the AV method is also biased when the underlying traffic is SRD.*

Since the estimated α (and consequently, H) in the AV test depends on several parameters, most notably the regression region (j_1, j_2) and the *vanishing moment* (M_{van}), we repeated the AV test using several values of (j_1, j_2) and M_{van}. Estimates for H along with their confidence intervals are shown in Table 1. Changing the values of j_1 and j_2 does not seem to help; in fact, the estimates deviate further from SRD. However, using $M_{van} = 4$ is found to yield a highly accurate estimate of H (for $(j_1, j_2) = (13, 16)$ and $M_{van} = 4$, $H \approx 0.491$). This could indicate that with appropriate selection of M_{van}, the AV test may be able to produce unbiased estimates of H even when the underlying traffic is slowly converging SRD. The appropriate selection of M_{van} is currently under investigation.

The next test that we examined is the absolute-value method. According to this method, one plots the sample mean of the sum of the absolute values of $\{X_n^{(m)} : n = 1, 2, \ldots\}$ for various values of the aggregation level m. The plot is fitted by a line, and the slope of that line is taken as an estimate of $H - 1$. Figure 17 depicts the resulting plot for the $M/G/\infty$ trace being examined. The H parameter is estimated at about 0.4932, in line with the true SRD identity of the trace.

Finally, we consider the R/S statistic, which is one of the earliest tests used for estimating the H parameter. The pox plot used in this test is shown in Figure 18 for the same SRD $M/G/\infty$ trace. Using linear regression, H is estimated at about 0.76, indicating a significant estimation bias.

5 Conclusions

Evidence supporting the existence of LRD in network traffic has been based on statistical techniques for estimating the Hurst parameter. In this paper, we examined the reliability of the VT test. We analyzed the aggregated variance in three, differently correlated random processes. Our main result is that this technique in inherently biased, and can often lead to incorrect conclusions about the true correlation structure of the examined data. This is true even in the absence of shifts in the mean of the process. Our finding can have significant implications on capacity planning and buffer engineering practices in QoS-based networks. The bias in the VT test gradually diminishes with the size of the data. For the examined models, we provided some guidelines on the required number of data points that are needed to render the bias insignificant. The VT test is not the only one that suffers from a bias; in fact, when the wavelet-based AV test and the R/S statistic were applied to a very long SRD $M/G/\infty$ trace, they also gave biased estimates of H. However, with proper selection of the vanishing moment, the AV test was able to produce an unbiased estimate. The absolute-value test gave an unbiased estimate for the same examined trace. Our future work will focus on producing more reliable variance-type tests for inference of LRD. One such attempt is found in [26].

References

[1] Patrice Abry and Darryl Veitch. Wavelet analysis of long-range dependent traffic. *IEEE Transactions on Information Theory*, 44(1):2–15, 1998.

[2] A. Adas. Traffic models in broadband networks. *IEEE Communications*, 35(7):82–89, July 1997.

[3] J. Beran. *Statistics for long-memory processes.* Chapman and Hall, New York, 1994.

[4] 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.

[5] D. R. Cox. Long-range dependence: A review. In H. A. David and H. T. David, editors, *Statistics: An Appraisal*, pages 55–74. The Iowa State University Press, Ames, Iowa, 1984.

[6] Mark Crovella and Azer Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[7] N. G. Duffield, J. T. Lewis, N. O'Connell, R. Russel, and F. Toomey. Statistical issues raised by Bellcore data. In *IEE 11'th UK Teletraffic Symposium*, March 1994.

[8] A. Elwalid, D. Heyman, T. V. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13(6):1004–1016, August 1995.

[9] 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.

[10] Victor S. Frost and Benjamin Melamed. Traffic modeling for telecommunications networks. *IEEE Communications Magazine*, 32(3):70–81, March 1994.

[11] Mark W. Garrett and Walter Willinger. Analysis, modeling, and generation of self-similar VBR video traffic. In *Proc. of the SIGCOMM '94 Conference*, pages 269–280, September 1994.

[12] M. Grossglauser and Jean-Chrysostome Bolot. On the relevance of long-range dependence in network traffic. *IEEE/ACM Transactions on Networking*, 7(5):629–640, October 1999.

[13] D. Heyman and T.V. Lakshman. What are the implications of long-range dependence for VBR video traffic engineering? *IEEE/ACM Transactions on Networking*, 4:301–317, June 1996.

[14] Changcheng Huang, Michael Devetsikiotis, Ioannis Lambadaris, and A. Kaye. Modeling and simulation of self-similar variable bit rate compressed video: A unified approach. In *Proceedings of the SIGCOMM '95 Conference*, pages 114–125, 1995.

[15] Predrag R. Jelenkovic, Aurel A. Lazar, and Nemo 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, August 1997.

[16] Marwan Krunz and Armand Makowski. Modeling video traffic using M/G/∞ input processes: A compromise between Markovian and LRD models. *IEEE Journal on Selected Areas in Communications*, 16(5):733–748, June 1998.

[17] Marwan Krunz and Arivu Mani Ramasamy. The correlation structure for a class of scene-based video models and its impact on the dimensioning of video buffers. *IEEE Transactions on Multimedia*, 2(1):27–36, March 2000.

[18] 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, February 1994.

[19] H. Michiel and K. Laevens. Teletraffic engineering in a broadband era. *Proceedings of the IEEE*, 85:2007–2033, 1997.

[20] Sandor Molnar and Trang Dinh Dang. Pifalls in long range dependence testing and estimation. In *Proceedings of the IEEE GLOBECOM 2000*, San Francisco, California, 2000.

[21] Sandor Molnar, Attila Vidacs, and Arne A. Nilsson. Bottlenecks on the way towards fractal characterization of network traffic: Estimation and interpretation of the Hurst parameter. In *Proceedings of PMCCN '97*, pages 125–144, Tsukuba, Japan, 1997.

[22] Minothi Parulekar. *Buffer Engineering for Self-Similar Traffic*. PhD thesis, University of Maryland, College Park, 1997.

[23] Minothi Parulekar and Armand Makowski. M/G/∞ input processes: A versatile class of models for network traffic. In *Proceedings of IEEE INFOCOM '97*, April 1997.

[24] Minothi Parulekar and Armand M. Makowski. Tail probabilities for M/G/∞ input processes (I): Preliminary asymptotics. *Queueing Systems - Theory & Applications*, 27(3/4):271–296, 1997.

[25] V. Paxson and Sally Floyd. Wide area traffic: The failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1993.

[26] Walter Rosenkrantz and J. Horowitz. Statistical analysis of variance-time plots used to estimate parameters of a long range dependent process, 1999. Private Communications.

[27] Matthew Roughan and Darryl Veitch. Measuring long-range dependence under changing traffic conditions. Available from `http://www.emulab.ee.mu.oz.au/~darryl`.

[28] Matthew Roughan, Darryl Veitch, and Patrice Abry. Real-time estimation of the parameters of long-range dependence. *IEEE/ACM Transactions on Networking*, 8(4):467–478, August 2000.

[29] Bong Ryu and Anwar 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, August 1996.

[30] M. S. Taqqu, V. Teverovsky, and W. Willinger. Estimators for long-range dependence: An empirical study. *Fractals*, 3(4):785–798, 1995.

[31] V. Teverovsky and M. S. Taqqu. Testing for long-range dependence in the presence of shifting means or a slowly declining trend, using a variance-type estimator. *Time Series Analysis*, 18(3):279–304, May 1997.

[32] Darryl Veitch and Patrice Abry. A wavelet-based joint estimator of the parameters of long-range dependence. *IEEE Transactions on Information Theory*, 45(3):878–897, April 1999.

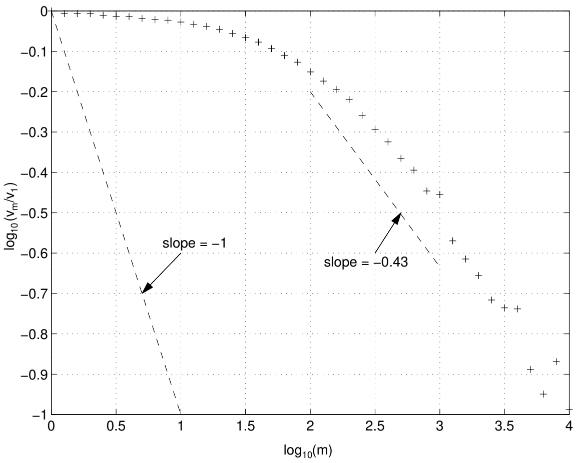

Figure 1: Empirical VT plot for the *Star Wars* trace.

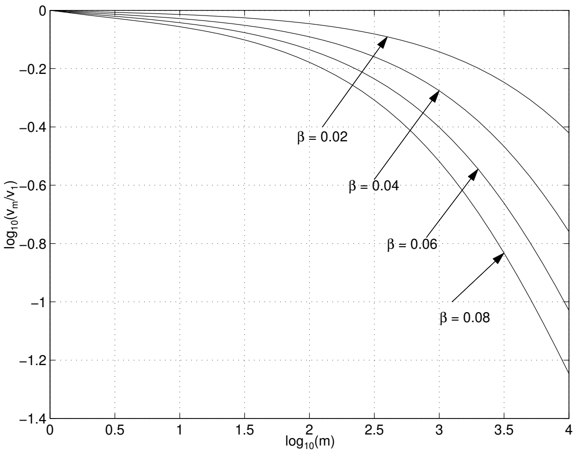

Figure 2: Analytically obtained VT plot for the $M/G/\infty$ model ($N = 2$).

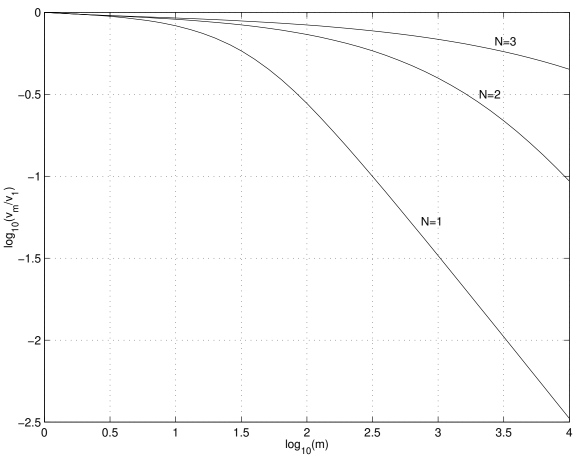

Figure 3: Analytically obtained VT plot for the $M/G/\infty$ model ($N = 1, 2, 3$).

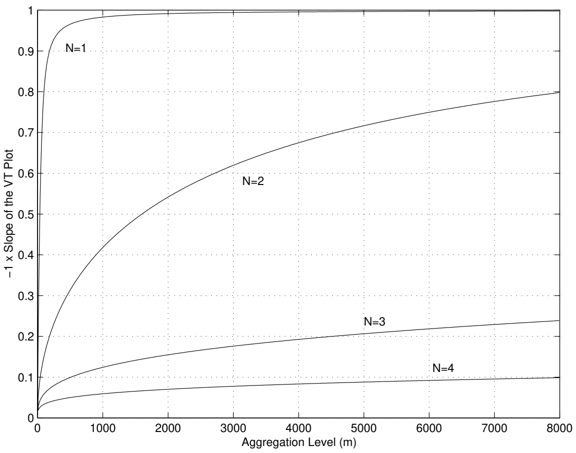

Figure 4: $-s_m^*$ versus m for the $M/G/\infty$ process ($N = 1, 2, 3, 4$).

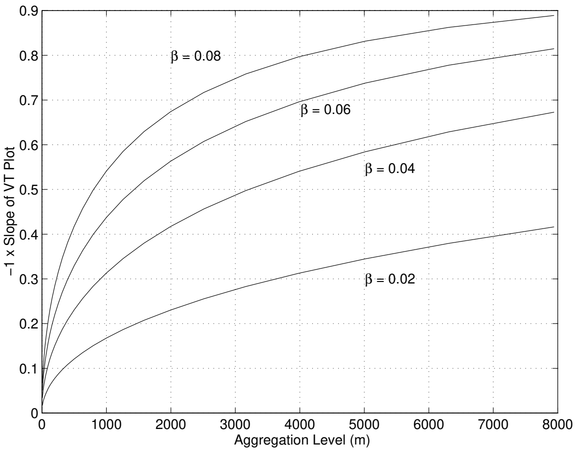

Figure 5: $-s_m^*$ versus m for the $M/G/\infty$ process with $N = 2$.

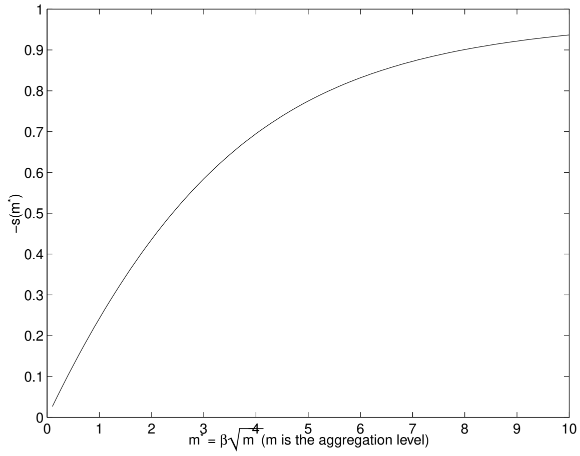

Figure 6: $-s_m^*$ versus $\widetilde{m} = \beta\sqrt{m}$ for the $M/G/\infty$ process.

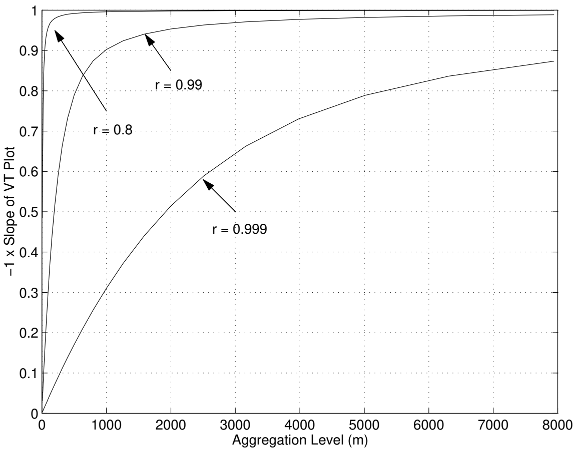

Figure 7: $-s_m$ versus m for the DAR(1) model.

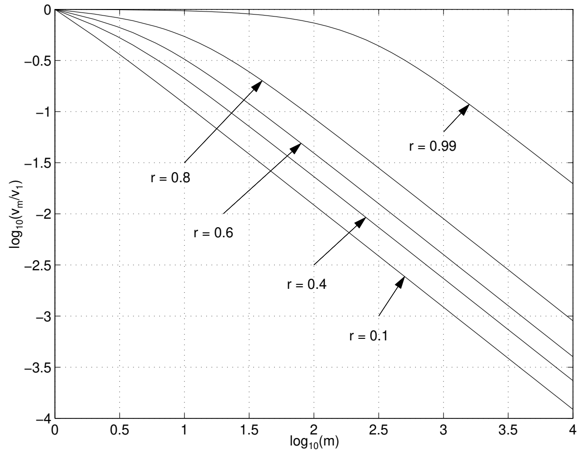

Figure 8: Analytically obtained VT plot for the DAR(1) model ($\rho_k = r^k$, $0 < r < 1$).

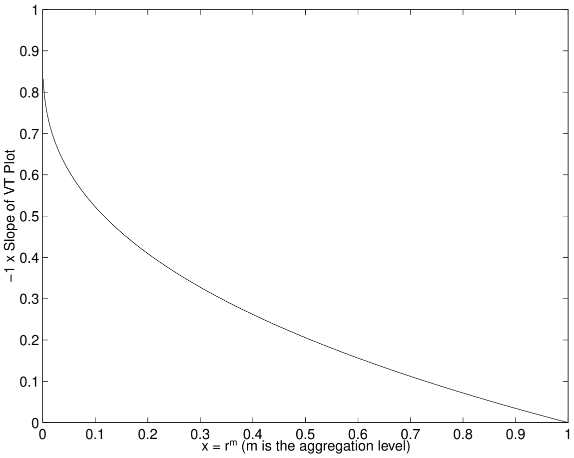

Figure 9: $-s_m^*$ versus $x = r^m$ in the DAR(1) model.

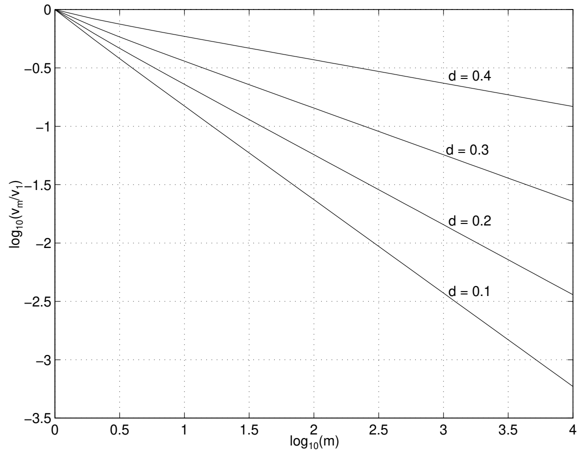

Figure 10: Analytically obtained VT plot for the LRD F-ARIMA$(0, d, 0)$ model.

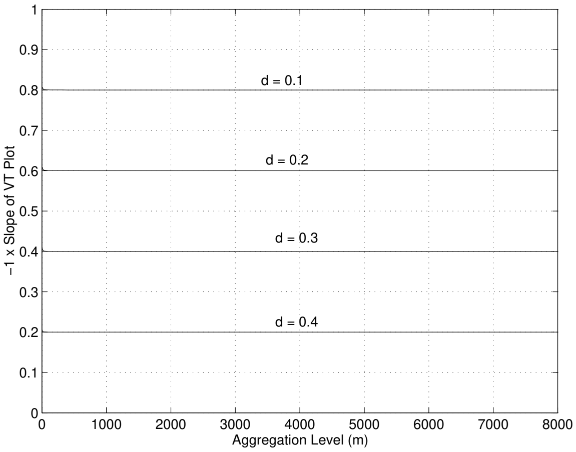

Figure 11: $-s_m$ versus aggregation level m in the F-ARIMA$(0, d, 0)$ model.

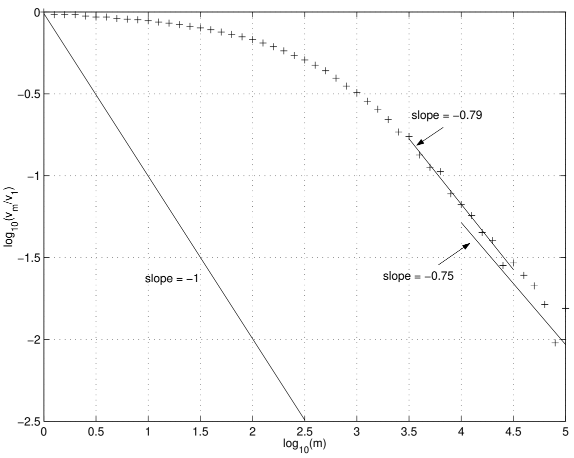

Figure 12: Empirical VT plot for a 1,000,000-long synthetic $M/G/\infty$ trace (for m in the range $[10^{3.5}, 10^{4.5}]$, the slope of the LMSE fit is -0.79, while for m in the range $[10^4, 10^5]$, the slope is -0.75).

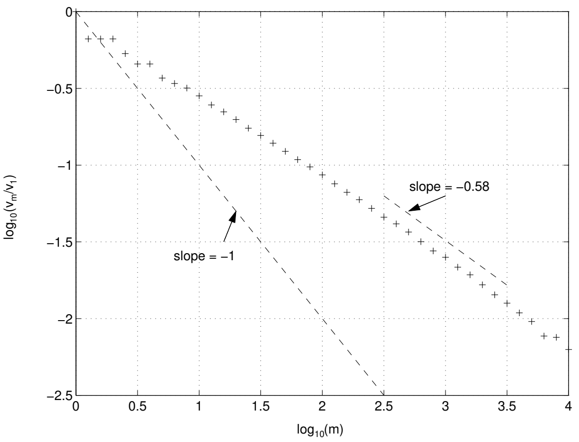

Figure 13: Empirical VT plot for a synthetic F-ARIMA trace.

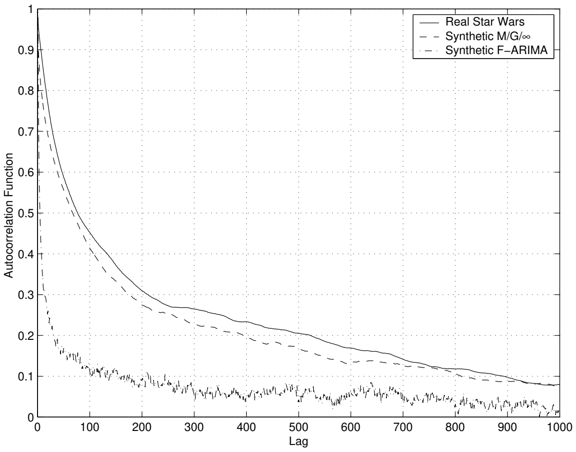

Figure 14: Autocorrelation functions for three tested traces.

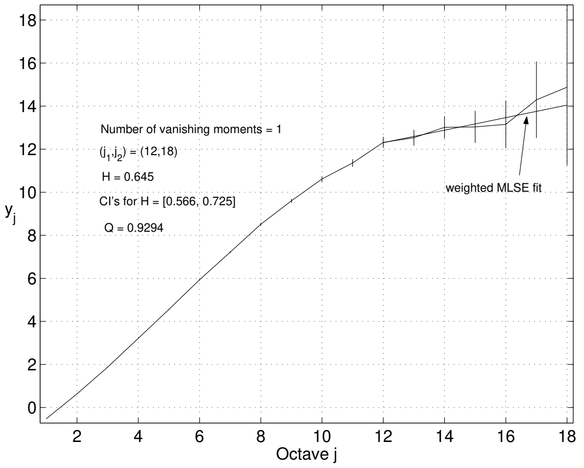

Figure 15: Logscale diagram in the AV test for a one-million-point $M/G/\infty$ trace with $\rho_k = e^{-0.076\sqrt{k}}$.

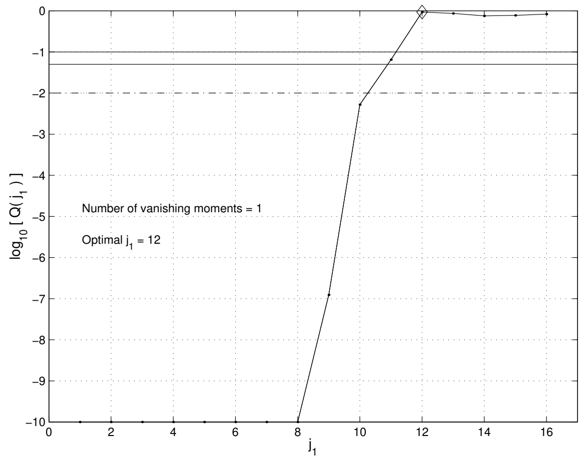

Figure 16: Goodness-of-fit statistic $Q(j_1)$ versus the lower limit of the scaling range (j_1).

Vanishing Moment	Range (j_1, j_2)	Estimated H	CI for H
1	(12,18)*	0.645	[0.566, 0.725]
1	(12,17)	0.641	[0.560, 0.722]
1	(12,16)	0.629	[0.542, 0.717]
1	(10,18)	0.801	[0.766, 0.836]
1	(10,17)	0.801	[0.766, 0.837]
2	(12,17)*	0.643	[0.557, 0.729]
2	(12,16)	0.647	[0.556, 0.738]
2	(12,15)	0.619	[0.515, 0.724]
2	(10,17)	0.805	[0.768, 0.841]
2	(10,16)	0.809	[0.772, 0.847]
3	(12,17)*	0.661	[0.569, 0.753]
3	(12,16)	0.661	[0.567, 0.755]
3	(12,15)	0.638	[0.532, 0.745]
3	(10,17)	0.812	[0.775, 0.850]
3	(10,16)	0.814	[0.776, 0.852]
4	(13,16)*	0.491	[0.319, 0.663]
4	(13,15)	0.473	[0.267, 0.680]
4	(11,16)	0.701	[0.641, 0.762]
4	(11,15)	0.714	[0.650, 0.778]

Table 1: Estimated value of H and its confidence interval (CI) based on the AV test. Estimation is done by means of weighted LSE fitting over the range of scales (j_1, j_2). For a given vanishing moment, the range for which j_1 is "optimal" and j_2 is maximal is indicated by a '*'.

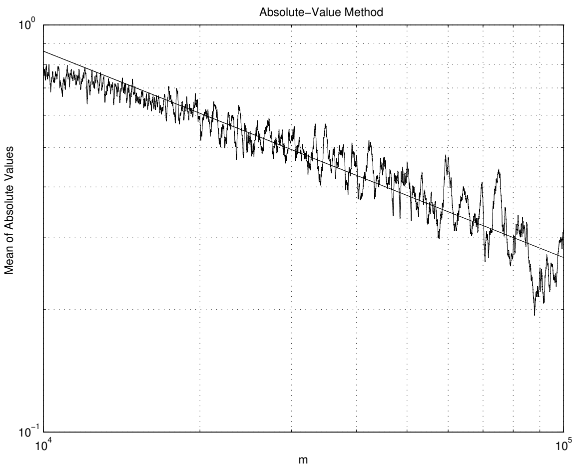

Figure 17: Estimation of H based on the absolute-value method for a 1,000,000-point-long SRD $M/G/\infty$ trace ($H \approx 0.4932$).

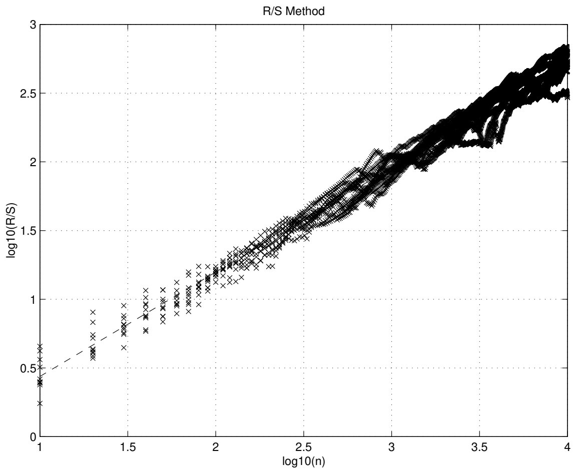

Figure 18: Estimation of H based on R/S method for a 1,000,000-point-long SRD $M/G/\infty$ trace (block size $= 100,000$, $H \approx 0.76$).

Biographical Sketches

Marwan M. Krunz

Marwan M. Krunz received his Ph.D. degree in Electrical Engineering from Michigan State University, Michigan, in 1995. From 1995 to 1997, he was a Postdoctoral Research Associate with the Department of Computer Science and the Institute for Advanced Computer Studies (UMIACS), University of Maryland, College Park. In January 1997, he joined the University of Arizona, where he is currently an Associate Professor of Electrical and Computer Engineering. Dr. Krunz's research interests lie in the field of computer networks, especially in its performance and traffic control aspects. His recent work has focused on the provisioning of quality of service (QoS) over wireless links, QoS routing (path selection, state aggregation), traffic modeling (video, WWW), bandwidth allocation, video-on-demand systems, and power-aware protocols for ad hoc networks. He has published more than 50 journal articles and refereed conference papers on these topics.

Dr. Krunz is a recipient of the National Science Foundation CAREER Award (1998-2002). He currently serves on the editorial board for the IEEE/ACM Transactions on Networking, the Computer Communications Journal, and the IEEE Communications Interactive Magazine. He was the Technical Program Co-chair for the 9th Hot Interconnects Symposium (Stanford University, August 2001). He serves as the Publicity Co-chair for the ACM/IEEE MOBICOM 2002 and the IEEE INFOCOM 2003 Conferences. He has served as the Tutorial Co-chair for the IEEE INFOCOM 2001, the Panel Chair for the IEEE IPCCC 2000 Conference, and the Panel Co-chair for the IEEE INFOCOM '99 Conference. Dr. Krunz has served and continue to serve on the technical program committees of many international conferences. He is a member of IEEE and ACM.

Ibrahim Matta

Ibrahim Matta received his Ph.D. in Computer Science from the University of Maryland at College Park. He is an Assistant Professor at the Computer Science Department of Boston University. His research involves the design and analysis of quality-of-service and wireless architectures and protocols. His recent projects investigate quality-of-service routing, Internet topology and traffic analysis, and Internet traffic controllers. Dr. Matta received the National Science Foundation Career Award in 1997. He is a member of IEEE and ACM.