On the Limitations of the Variance-Time Test for Inference of Long-Range Dependence

Marwan Krunz
Department of Electrical & Computer Engineering
University of Arizona
Tucson, AZ 85721
krunz@ece.arizona.edu

Abstract— **The objective of this paper is to demonstrate the limitations of the variance-time (VT) test as a statistical tool for inferring long-range dependence (LRD) in network traffic. Since the early Bellcore studies [7], [16], LRD has been in the center of a continuous debate within the teletraffic community. The controversy is typically focused on the utility of LRD models to predict the performance at network buffers. Our work here is not meant to advocate one modeling approach over another, but to point out (experimentally and theoretically) to the caveats in using the VT test as a tool for detecting LRD. We do that by deriving simple analytical expressions for the slope of the aggregated variance in three autocorrelated traffic models: $M/G/\infty$ process (short-range dependent (SRD) but non-Markovian), the discrete autoregressive of order one model (SRD Markovian), and the fractional ARIMA process (LRD). Our main result is that the VT test 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 the VT test, 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.**

I. Introduction

One of the controversial issues in current teletraffic research is 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 many of these correlations should be captured in a traffic model [10], [11], [23]. 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 [8], [1], [17] 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 [19], [14] and a class of (subexponential) path-based Markov renewal processes [12], [15]. The interest in such models is related to their ability to produce a wide range of correlation structures, including, as extreme cases, both Markovian and LRD structures.

This work was supported in part by the National Science Foundation under grants ANI 9733143 and CCR 9979310.

The debate on the significance of traffic correlations was spurred by studies conducted at Bellcore and elsewhere, which indicated the existence of persistent correlations in various types of network traffic, including LAN Ethernet [7], [16], WAN [21], and variable-bite-rate (VBR) video traffic [3], [9]. It was argued that such persistence can be adequately captured by LRD models, for which the ACF decays slowly (typically as a power function) to the extent that the correlations now have an infinite sum, i.e., $\sum_k \rho_k = \infty$. The LRD phenomenon has long been observed in other domains such as hydrology and economics [13]. In teletraffic studies, advocates of LRD argue that such a phenomenon has significant impact on network performance, and thus must be accounted for when dimensioning network resources. On the other side, advocates of Markovian modeling, while acknowledging the presence of such a phenomenon, argue that for networks with *finite* buffers it is sufficient to incorporate correlations up to some *critical* lag that depends on the buffer size and the bandwidth capacity [10], [11], [23].

Foremost, the statistical evidence supporting LRD was based on the estimated value of the Hurst (H) parameter, which is between 0.5 and 1 for LRD models. Several statistical tests have been used to estimate the H parameter, including the R/S statistic, periodogram-based analysis (i.e., Whittle's estimator), and the VT test (see [2], for example). In this paper, we consider the VT test. Our main goal is study the accuracy of this test. In particular, we show that such a test can give misleading indication about the true SRD/LRD nature of a given time series, despite the availability of many data points. The bias in the VT estimate has been noted before [5], [24] through empirical studies, but the focus has been mainly on the bias caused by jumps in the mean of the SRD process or by slow deterministic trends [25]. Instead, our focus here is on SRD processes (often, non-Markovian) with a constant mean and no trends. We provide analytical evidence of the bias in such processes. Our results are applied to three 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). It should be emphasized that our work is not meant to advocate one model over another, but to point out to the caveates 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 II we briefly describe the three processes that are used in our study. The aggregated variance for each of these processes is derived in Section III. In the same section, we discuss the limitations of the VT test. The paper is concluded in Section IV.

II. Autocorrelated Processes

A. $M/G/\infty$ 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 [14], [19] 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 [18]:

(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 $\{b_n^\star,\ n = 0, 1, \ldots\}$ 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 [20] 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 $(1, 2)$ [4].

B. 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 [6]. 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 $k = 0, 1, \ldots$.

C. 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 [9]. 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 $k^{-\beta}$, where $\beta = 2 - 2H$. See [9] for details on how to generate synthetic F-ARIMA traces.

III. Analysis of Aggregated Variance

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

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

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:

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

We will refer to v_m as the aggregated variance at level m. If $\{X_n\}$ is an LRD process, then it must satisfy $mv_m \to \infty$ as $m \to \infty$. More specifically, for an LRD process $v_m \sim m^{-\beta}$ when m is large, where $0 < \beta < 1$ is the same parameter defined above. For an SRD process, $\beta \geq 1$. 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 $-\beta$. If $\beta < 1$, 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 [3] and [9].

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

A. Aggregated Variance in the $M/G/\infty$ Model

Let $\{X_n : n = 1, 2, \ldots\}$ be an $M/G/\infty$ process with an ACF of the form $\rho_k = e^{-\beta\sqrt{k}}$. A process with such an ACF has been used in modeling JPEG video traffic [14]. It is clearly a SRD process since $\sum_k \rho_k < \infty$. With this choice of the ACF, the service-time distribution (G) has a Weibull-like form. Now consider the aggregated process $\{X_n^{(m)} : n = 1, 2, \ldots\}$ for $m = 1, 2, \ldots$. The *normalized* aggregated variance of this process can be written as [4]:

$$\widetilde{v}_m \stackrel{\text{def}}{=} \frac{v_m}{v_1} = \frac{1}{m} + \frac{2}{m^2}\sum_{k=1}^{m}(m-k)\rho_k \tag{8}$$

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:

$$s_m \stackrel{\text{def}}{=} \frac{\log \widetilde{v}_{m+1} - \log \widetilde{v}_m}{\log(m+1) - \log m} \tag{9}$$

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. Figure 2 depicts $\widetilde{v}_m$, obtained using (8), versus m when $\rho_k = e^{-\beta\sqrt{k}}$. Using (9), the corresponding s_m is shown in Figure 3. Clearly, s_m converges very slowly to -1. In fact, even at an aggregation level of $m = 8000$, $-s_m$ is still smaller than one. Have 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.

Instead of relying on (8) and (9) to obtain s_m, we now provide an almost exact closed-form expression for 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 [4]:

$$\widetilde{v}^*{}_m = \frac{2}{m^2}\int_0^m (m-h)\rho_h dh = \frac{2}{m^2}\int_0^m (m-h)e^{-\beta\sqrt{h}}dh \tag{10}$$

Equation (10) can be written as follows:

$$\begin{aligned}\widetilde{v}^*{}_m &= \frac{2}{m^2}\left[m\int_0^m e^{-\beta\sqrt{t}}dt - \int_0^m te^{-\beta\sqrt{t}}dt\right]\\ &= e^{-\beta\sqrt{m}}\left[\frac{8}{\beta^2 m} + \frac{24}{\beta^3 m\sqrt{m}} + \frac{24}{\beta^4 m^2}\right.\\ &\quad \left. + \frac{4}{\beta^2 m} - \frac{24}{\beta^4 m^2}\right]\end{aligned} \tag{11}$$

Note that as $m \to \infty$, $\widetilde{v}^*{}_m \sim 4/(m\beta^2) = \mathcal{O}(1/m)$, as expected. The plots of $\widetilde{v}^*{}_m$ versus $\log m$ (not shown) were found indistinguishable from those in Figure 2. Hence, it sufficies to use $\widetilde{v}^*{}_m$ in place of $\widetilde{v}_m$. Next, we obtain the slope of $\widetilde{v}^*{}_m$ defined as:

$$s_m^* \stackrel{\text{def}}{=} \frac{d(\log \widetilde{v}^*{}_m)}{d(\log m)} = \frac{m}{\widetilde{v}^*{}_m}\frac{d\widetilde{v}^*{}_m}{dm} \tag{12}$$

With some basic algebraic manipulations, it is easy to show that:

$$\begin{aligned}\frac{d\widetilde{v}^*{}_m}{dm} &= e^{-\beta\sqrt{m}}\left[\frac{-20}{\beta^2 m^2} - \frac{48}{\beta^3 m^2\sqrt{m}} - \frac{48}{\beta^4 m^3}\right.\\ &\quad \left. - \frac{4}{\beta m\sqrt{m}}\right] - \frac{4}{\beta^2 m^2} + \frac{48}{\beta^4 m^3}\end{aligned} \tag{13}$$

from which, we conclude that

$$s_m^* = \frac{e^{-\beta\sqrt{m}}\left[\frac{-20}{\beta^2 m} - \frac{48}{\beta^3 m\sqrt{m}} - \frac{48}{\beta^4 m^2} - \frac{4}{\beta\sqrt{m}}\right] - \frac{4}{\beta^2 m} + \frac{48}{\beta^4 m^2}}{e^{-\beta\sqrt{m}}\left[\frac{8}{\beta^2 m} + \frac{24}{\beta^3 m\sqrt{m}} + \frac{24}{\beta^4 m^2}\right] + \frac{4}{\beta^2 m} - \frac{24}{\beta^4 m^2}} \tag{14}$$

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

$$\mid s_m^* \mid < \mid s_m \mid < \mid s_{m+1}^* \mid \tag{15}$$

Let $\widetilde{m} \stackrel{\text{def}}{=} \beta\sqrt{m}$. Then, (14) can be written as a function of $\widetilde{m}$:

$$s_m^* = \frac{e^{-\widetilde{m}}\left[\frac{-20}{\widetilde{m}^2} - \frac{48}{\widetilde{m}^3} - \frac{48}{\widetilde{m}^4} - \frac{4}{\widetilde{m}}\right] - \frac{4}{\widetilde{m}^2} + \frac{48}{\widetilde{m}^4}}{e^{-\widetilde{m}}\left[\frac{8}{\widetilde{m}^2} + \frac{24}{\widetilde{m}^3} + \frac{24}{\widetilde{m}^4}\right] + \frac{4}{\widetilde{m}^2} - \frac{24}{\widetilde{m}^4}} \tag{16}$$

The plot of s_m^* versus $\widetilde{m}$ is shown in Figure 4. 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 $\rho_k = e^{-\beta\sqrt{k}}$, 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 $| s^*_m | \geq 0.95$ as an indication of SRD, then in this case 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 $\rho_k = e^{-\beta\sqrt{k}}$ [14]), then we need at least 50176 data points to correctly infer that the data exhibit SRD.

B. 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 (8), and after some straightforward manipulations, we obtain:

$$\widetilde{v}_m = \frac{1}{m} + \frac{2}{m^2}\left(\frac{r(r^m - mr + m - 1)}{(r-1)^2}\right) \tag{17}$$

As $m \to \infty$, $\widetilde{v}_m \sim (1 + 2r/(1-r))/m$, which is, as expected, $\mathcal{O}(m)$. By substituting the values for $\widetilde{v}_m$ and $\widetilde{v}_{m+1}$ in (9), we can plot the first-order difference s_m versus m, as shown in Figure 5. 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 (12). By substituting $\rho_h = r^h$ in (10) and after some manipulations, we arrive at the following expression for $\widetilde{v^*}_m$:

$$\widetilde{v^*}_m = \frac{2}{m^2}\left[\frac{r^m - m\ln r - 1}{(\ln r)^2}\right] \tag{18}$$

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 dicrete-parameter counterparts. For brevity, we only show the plots in the continuous case (Figure 6).

Differentiating $\widetilde{v^*}_m$ in (18) 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 (12) s^*_m for the DAR(1) model is given by

$$s^*_m = \frac{(m\ln r - 2)r^m + m\ln r + 2}{r^m - m\ln r - 1} \tag{19}$$

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 (19). To get an idea about how many data points are sufficient to infer SRD/LRD, we first rewrite (19) in terms of the variable $x \stackrel{\text{def}}{=} r^m$ as follows:

$$s^*_m = \frac{(\ln x - 2)x + \ln x + 2}{x - \ln x - 1} \tag{20}$$

Figure 7 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 sufficient 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 $\widetilde{v^*}_m$.

C. Aggregated Variance in the F-ARIMA Model

Finally, consider the F-ARIMA process described before. We first provide a simple recursive approach for computing $\widetilde{v}_m$ 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 8 can now be written as follows:

$$\widetilde{v}_m = \frac{1}{m} + \frac{2}{m^2}(mX_m - Y_m). \tag{21}$$

Since $X_m = X_{m-1} + \rho_m$, $Y_m = Y_{m-1} + m\rho_m$, and $\rho_m = (m-1+d)/(m-d)\rho_{m-1}$, (21) can be computed recursively starting from $X_1 = Y_1 = \rho_1 = d/(1-d)$. Figure 8 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 6 for the $M/G/\infty$ and DAR(1) models). 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 9 as a function of m (obtained using (9)).

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 10 and 11 depict the results for two representative traces. For the $M/G/\infty$ trace, we set $\beta = 0.076$ in $\rho_k = \exp(-\beta\sqrt{k})$, which produced a good fit for the empirical ACF of the *Star Wars* sequence [14]. For the F-ARIMA trace, we took $d = 0.3$. The $M/G/\infty$ trace consisted of 1,000,000 data points, while the F-ARIMA had 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 10 for the $M/G/\infty$ trace indicates an asymptotic slope of about 0.59 (starting with aggregation level of $10^{2.5} = 317$), which according to the VT test, should 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 (longer than any of the empirical video records ever tested), the VT test is wrongly suggesting the presence of

LRD in this trace. The trends in both figures are in line with our previous analysis of the aggregated variance. For the F-ARIMA trace (Figure 11), the slope of the VT plot is estimated at -0.58, which is surprisingly close to that of the $M/G/\infty$ trace. 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 12 depicts the negative of the slope of the VT plot as a function of the first aggregation level (S) for the original *Star Wars* sequence and the $M/G/\infty$ trace (only aggregated variances with aggregation levels $m \geq S$ are used in fitting the VT plot). For the $M/G/\infty$ trace, the slope increases very slowly with S, and expectedly tends to -1 as $S \to \infty$. However, the convergence is quite slow, that even for traces 1,000,000-long, the VT plot can give misleading conclusions about the exsitence of SRD/LRD structure in the data.

IV. 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. As a byproduct of our study, we noted that from an aggregated variance standpoint, a non-Markovian SRD $M/G/\infty$ model with ACF of the form $\rho_k = \exp(-\beta\sqrt{k})$ is more appropriate for modeling the JPEG-coded *Star Wars* trace than the LRD F-ARIMA model proposed in [9]. Our future work will focus on producing more reliable variance-type tests for inference of LRD. One such attempt is found in [22].

References

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

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

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

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

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] V. Klemes. The Hurst phenomenon: A puzzle? *Water Resources Res.*, 10:675–688, 1974.

[14] 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.

[15] 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.

[16] 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.

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

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

[19] 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.

[20] Minothi Parulekar and Armand M. Makowski. Tail probabilities for M/G/∞ input processes (I): Preliminary asymptotics. *Queueing Systems - Theory & Applications*, xx:xxxx, 1998.

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

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

[23] 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.

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

[25] 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.

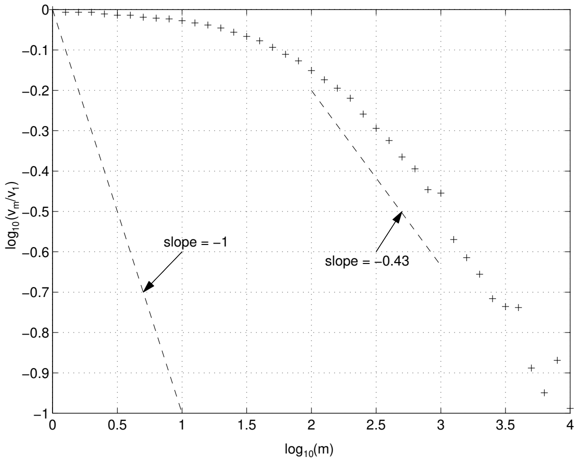

Fig. 1. Empirical VT plot for the *Star Wars* trace.

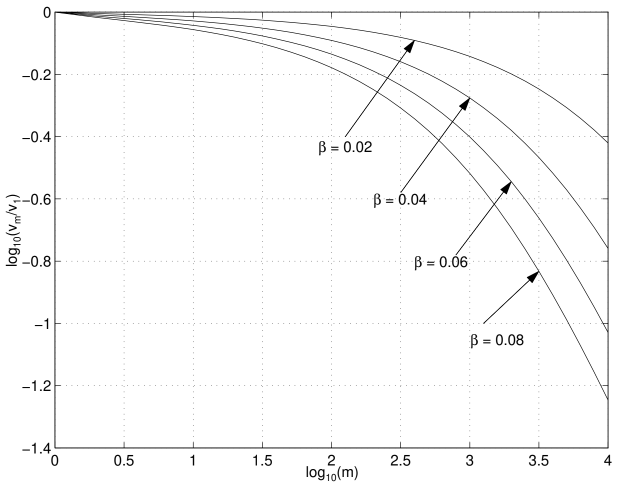

Fig. 2. Analytically obtained VT plot for the SRD $M/G/\infty$ model ($\rho_k = e^{-\beta\sqrt{k}}$).

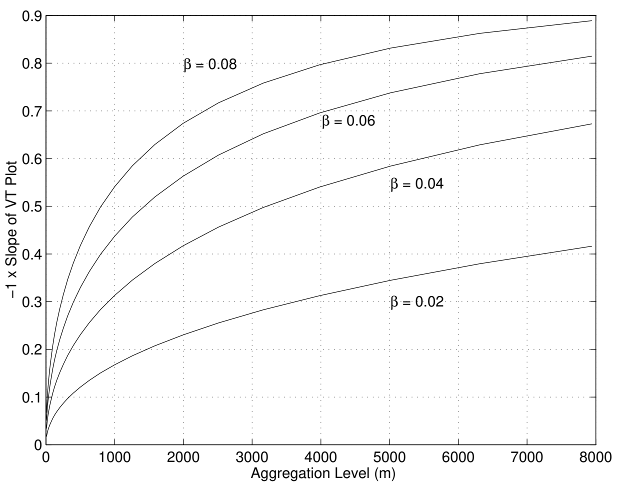

Fig. 3. Discrete-time instantaneous slope s_m versus m for the $M/G/\infty$ process.

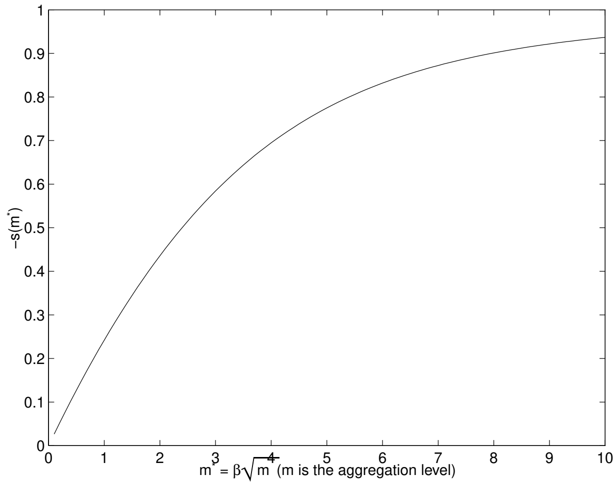

Fig. 4. $-s^*_m$ versus $\widetilde{m} = \beta\sqrt{m}$ for the $M/G/\infty$ process.

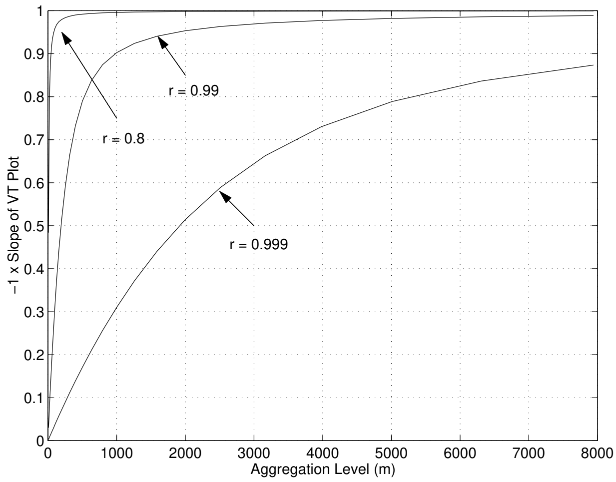

Fig. 5. $-s_m$ versus m for the DAR(1) model.

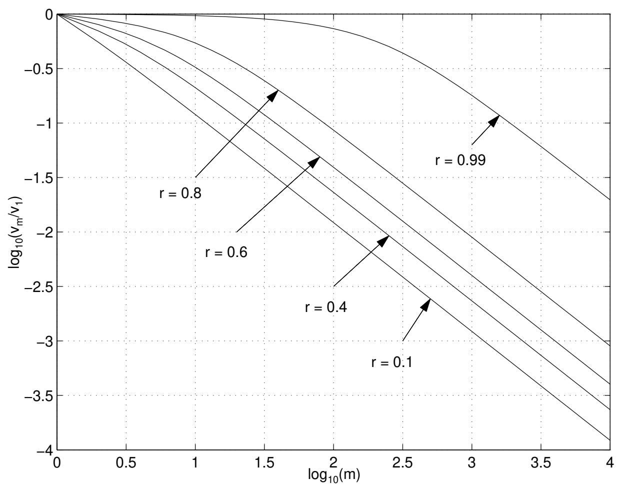

Fig. 6. Analytically obtained VT plot for the DAR(1) model ($\rho_k = r^k$, $0 < r < 1$).

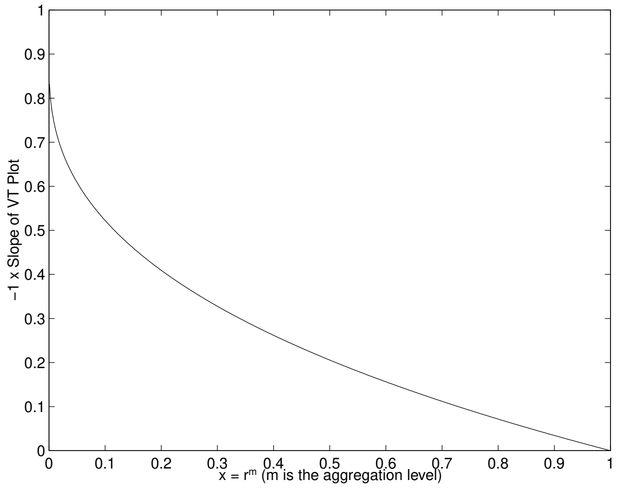

Fig. 7. $-s_m^*$ versus $x = r^m$ in the DAR(1) model.

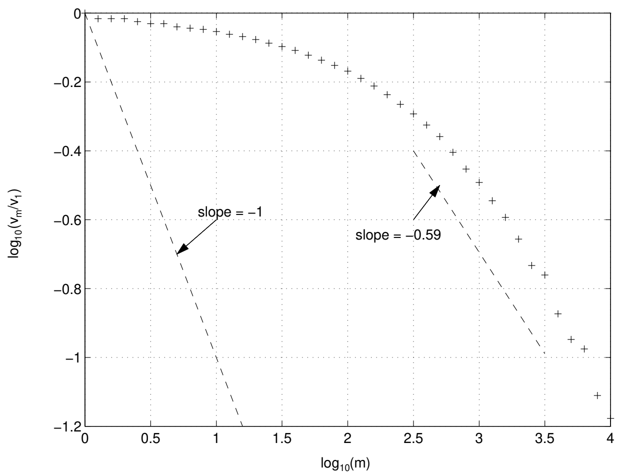

Fig. 10. Empirical VT plot for a synthetic $M/G/\infty$ trace.

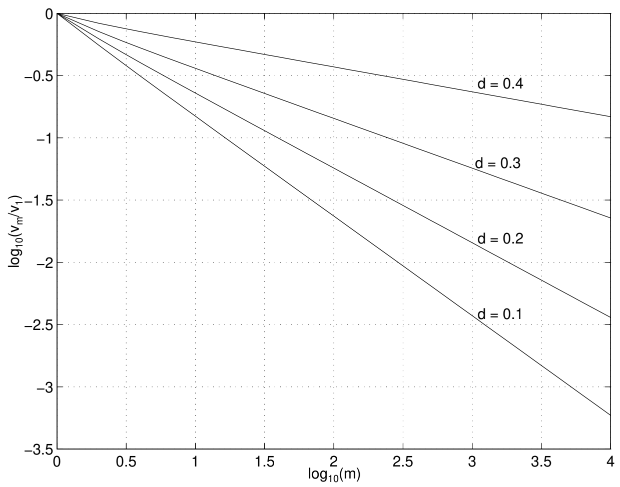

Fig. 8. Analytically obtained VT plot for the LRD F-ARIMA$(0, d, 0)$ model.

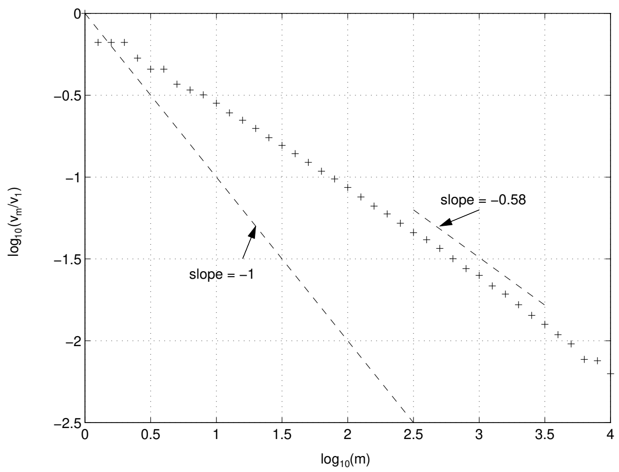

Fig. 11. Empirical VT plot for a synthetic F-ARIMA trace.

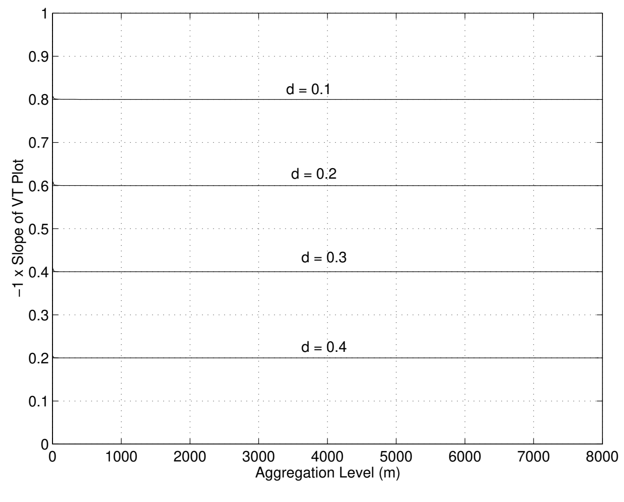

Fig. 9. $-s_m$ versus aggregation level m in the F-ARIMA$(0, d, 0)$ model.

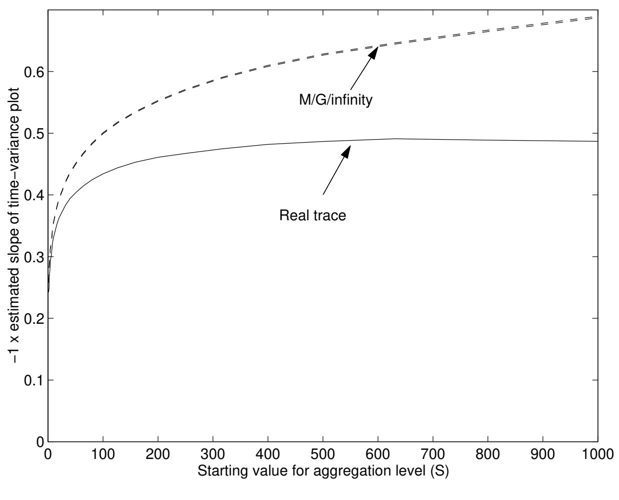

Fig. 12. Slope of empirical VT plot versus first aggregation level.