

Describing Network Traffic Using the Index of Variability

Georgios Y. Lazarou, Julie Baca, Victor S. Frost and Joseph B. Evans

Abstract—Commonly used measures of traffic burstiness do not capture the fluctuation of traffic variability over the entire range of time scales. In this paper, we present a measure of variability, called the *Index of Variability* ($H_v(\tau)$), that depicts the degree of variability (burstiness) of a typical network traffic process at each time scale and is analytically tractable for many traffic models. As an illustration, we derive the closed-form expressions of $H_v(\tau)$ for two traditional traffic models and generate a variety of 2D and 3D Index of Variability curves. These curves demonstrate that the Index of Variability can help in determining the complexities of the network traffic variability over the network performance relevant time scales. We then introduce a practical method for estimating the Index of Variability curve from a given traffic trace. Using this method, we estimate the Index of Variability curves for 12 long NLANR network traffic traces. The results indicate that the variability of real network traffic varies with time scales and that the Index of Variability has the ability to discern qualitative differences between traffic traces obtained from different networks. Thus, the Index of Variability offers the potential to gain insights into the dynamics of network traffic that existing tools do not offer.

I. INTRODUCTION

Many empirical studies have shown that Internet traffic exhibits high variability¹ [2], [3], [4], [5], i.e., traffic is bursty (variable) over a wide range of time scales in sharp contrast to the assumption that traffic burstiness exists only at short time scales while traffic is smooth at large time scales [4]. High variability in traffic has been shown to have a significant impact on network performance [4], [6]. The results from [6], [7], [8], [9] show that knowledge of the traffic characteristics on multiple time scales helps to improve the efficiency of traffic control mechanisms. More importantly, the design and provision of quality-of-service-guarantees over the

Internet requires the understanding of traffic characteristics, such as variability.

Since the publication of [4], the popular belief has been that the high variability in traffic is due to the *long-range dependence* (LRD) property of the traffic processes. In general, a (weakly) stationary discrete-time real-valued stochastic process $Y = \{Y_n, n = 0, 1, 2, \dots\}$ with mean $\mu = E[Y_n]$ and variance $\sigma^2 = E[(Y_n - \mu)^2] < \infty$ is long-range dependent if $\sum_{k=1}^{\infty} r(k) = \infty$, where $r(k)$ measures the correlation between samples of Y separated by k units of time. If $\sum_{k=1}^{\infty} r(k) < \infty$, then Y is said to exhibit *short-range dependence* (SRD).

Common traffic models with LRD are based on self-similar processes. In traffic modeling, the term self-similarity usually refers to the *asymptotically second order self-similar* or *mono-fractal* processes [10]. The definition of asymptotically second order self-similarity is as follow [4]: assume that Y has an autocorrelation function of the form $r(k) \sim k^{-\beta}L(k)$ as $k \rightarrow \infty$, where $0 < \beta < 1$ and the function L is slowly varying at infinity, i.e., $\lim_{k \rightarrow \infty} \frac{L(kx)}{L(k)} = 1 \forall x > 0$. For each $m = 1, 2, 3, \dots$, let $Y^{(m)} = \{Y_n^{(m)}, n = 1, 2, 3, \dots\}$ denote a new aggregated time series obtained by averaging the original series Y over non-overlapping blocks of size m , replacing each block by its sample mean. That is, for each $m = 1, 2, 3, \dots$, $Y^{(m)}$ is given by

$$Y_n^{(m)} = \frac{Y_{nm-m+1} + \dots + Y_{nm}}{m} \quad n \geq 1. \quad (1)$$

The new aggregated discrete-time stochastic process $Y^{(m)}$ is also (weakly) stationary with an autocorrelation function $r^{(m)}(k)$; then, Y is called asymptotically second order self-similar with self-similar parameter $H = 1 - \frac{\beta}{2}$ if for all k large enough, $r^{(m)}(k) \rightarrow r(k)$ as $m \rightarrow \infty$, that is, Y is asymptotically second-order self-similar if the corresponding aggregated processes $Y^{(m)}$ become indistinguishable from Y at least with respect to their autocorrelation functions. **By definition, asymptotically second order self-similarity implies LRD** [10].

The parameter H is called the *Hurst parameter*. For general self-similar processes, it measures the degree of “self-similarity”. For random processes suitable for modeling network traffic, the Hurst parameter is basi-

Part of this paper was presented at MSO’03 [1], ©IASTED 2003.

G.Y. Lazarou and J. Baca are with the Mississippi State University glaz@ece.msstate.edu, baca@cavs.msstate.edu.

V.S. Frost and J.B. Evans are with the University of Kansas, {frost,evans}@eecs.ku.edu.

¹Fluctuation of traffic as a function of time.

cally a measure of the speed of decay of the tail of the autocorrelation function. If $0.5 < H < 1$, then the process is LRD, and if $0 < H < 0.5$, then it is SRD. If $H = 0.5$ then the process can be either SRD or LRD. However, most processes relevant to traffic modeling with $H = 0.5$ are SRD, eg., Markov processes. Hence, H is widely used to capture the intensity of long-range dependence of a traffic process: the closer H is to 1 the more long-range dependent the traffic is, and vice versa [10].

Several methods exist for estimating H from a traffic trace. One of the most widely used is the *Aggregated Variance* method: for successive values of m that are equidistant on a log scale, the sample variance of $Y^{(m)}$ is plotted versus m on a log-log plot [11][12]. By fitting a least-square line to the points of the plot and then calculating its slope, an estimate of the Hurst parameter is obtained as $\hat{H} = 1 - \frac{|\text{slope}|}{2}$.

Another very popular method is based on wavelets [13]. Given a traffic trace Y_n , the Hurst parameter can be estimated as follows: for each scale j , the wavelet energy $\mu_j = \frac{1}{N_j} \sum_{k=1}^{n_j} d^2(j, k)$ is plotted versus j on a semi-log plot (i.e., $\log_2(\mu_j)$ vs. j). By fitting a least-square line to the points of the curve region that looks linear and then computing its slope α , H is estimated as $\hat{H} = \frac{\alpha+1}{2}$.

A. Need for a New Measure of Variability

Commonly used measures of traffic burstiness, such as the peak-to-mean ratio, the coefficient of variation of interarrival times, the indices of dispersion for intervals and counts, and the Hurst parameter, do not capture the fluctuation of variability over different time scales.

It is claimed in [4] that the Hurst parameter is a good measure of variability, and the higher the value of H , the burstier the traffic. The popular belief from early studies [6], [14], [15], [16], [17], [18] on the impact of LRD on network performance is that high values of the Hurst parameter are associated with poor queueing performance. However, later studies [8], [19] show examples in which larger values of H are associated with better queueing performance compared to smaller values of H . In addition, the results in [9] indicate that the queueing performance depends mostly on the variability over certain time scales rather than on the value of H .

Moreover, it is known [7] that different long-range dependent processes with the same value of the Hurst parameter can generate vastly different queueing behavior. Clearly, the single value Hurst parameter does not capture the fluctuation of the degree of traffic burstiness across time scales, regardless of whether the traffic process exhibits LRD or SRD. From the definition presented in the previous section, the Hurst parameter is defined




Fig. 1. $\log_2(\mu_j)$ versus scale j for the Auckland-IV traffic trace 20010301-310-0 (For information about the trace, see Section IV).

asymptotically (i.e., for large time scales) and hence conveys nothing about the variability of measured traffic over small or medium time scales; unless the traffic is exactly self-similar with known variances. Therefore, the Hurst parameter is an incomplete descriptor of traffic variability.

For many network traffic processes, the wavelet energy-scale or variance-time plots usually do not tend to straight lines, i.e., see Fig. 1. Usually many of these processes have piecewise fractal behavior with varying Hurst parameter over some small ranges of time scales [20]. Such processes are usually referred to as multifractal [21], [22].

Queueing performance greatly depends on traffic irregularities at small time scales which are typically attributed to the complex dynamics of data networks [7], [23]. Multifractal analysis based on the Legendre spectrum is often used to study the multiscaling behavior of traffic at small time scales [20], [24], [25], [26]. The process of estimating the Legendre spectrum involves higher order sample moments and negative values of moments. It is known [27] that higher order sample moments are not well-behaved and negative values of moments tend to be erratic. In addition, the Legendre spectrum is difficult to interpret [28]. The infinitely divisible cascade analysis represents an improvement over the multifractal analysis [29]. It relaxes the questionable assumption that moments behave as power-laws of scales [22].

Hence, there is a need for an intuitively appealing, conceptually simple, and mathematically rigorous measure which can capture the various scaling phenomena that are observed in data networks on both small and large scales [30]. In this paper, we present an alternative measure of variability, called the *Index of Variability*

$(H_v(\tau))$, that captures the degree of variability of a typical network traffic process at each time scale and is analytically tractable for many traffic models.

The remainder of this paper is organized as follows: Section II presents the definition of the *Index of Variability*. Section III presents the derivation of the closed-form expressions of $H_v(\tau)$ for two traditional traffic models. A variety of 2D and 3D Index of Variability curves are also presented for illustration. Section IV presents a practical method for estimating the Index of Variability curve from a given traffic trace and experimental results. The paper concludes in Section V.

II. INDEX OF VARIABILITY FOR PACKET TRAFFIC SEQUENCES

Let $N(t)$ denote the number of events (packet arrivals) of a stationary point process in the interval $(0, t]$. For each fixed time interval $\tau > 0$, an event count sequence $Y = \{Y_n(\tau), \tau > 0, n = 1, 2, \dots\}$ can be constructed from each point process, where

$$Y_n(\tau) = N[n\tau] - N[(n-1)\tau] \quad (2)$$

denotes the number of events that have occurred during the n^{th} time interval of duration τ . Clearly, Y is also (weakly) stationary for all $\tau > 0$. In this study, Y represents a network traffic trace where $Y_n(\tau)$ denotes the number of packets observed from an arbitrary point in the network during the n^{th} time interval of duration τ . We refer to τ as the *time scale* of the traffic trace, and it represents the length (i.e., 10ms, 1s, 10s, e.t.c.) of one sample of Y .

The expected number of events that have occurred during the interval $(0, t]$ is always: $E[N(t)] = \frac{t}{E[X]} = \lambda t$ where $E[X]$ is the expected interarrival time and λ is the mean event (packet) arrival rate. The index of dispersion for counts (IDC) is defined as: $IDC(t) \equiv \frac{\text{Var}[N(t)]}{E[N(t)]} = \frac{\text{Var}[N(t)]}{\lambda t}$. The IDC was defined such that it provides some comparison with the Poisson process, for which $IDC(t) = 1 \forall t$. Note that since the point process is stationary, IDC has the same value over any interval of length t ; thus, t can be viewed as the time scale τ of the traffic process Y defined in (2). Henceforth t will be used to denote generality and τ to denote time scales, i.e., the time length of each sample of the packet-count sequence Y .

An important feature of IDC is that it is mathematically equivalent to the Aggregated Variance method for estimating the Hurst parameter H of a self-similar process. For a self-similar process, plotting $\log(IDC(m\tau))$ against $\log(m)$ results in an asymptotic straight line with slope $2H - 1$. When Y is a long-range dependent process,

the slowly decaying variance property of LRD processes [4] with parameter $0 < \beta < 1$ is equivalent to an IDC curve² with an asymptotic straight line with slope $1 - \beta$, implying $0 < \text{slope} < 1$. When the IDC curve converges to an asymptotic straight line with $\text{slope} = 0$ for some $\tau < \infty$, then this is suggestive of Y been a short-range dependent process. Based on the above property of IDC, we define the following new measure of variability:

Definition 1: For a general stationary traffic process Y as defined by (2) whose $IDC(\tau)$ is continuous and differentiable over $(0, \infty)$, we refer to:

$$H_v(\tau) \equiv \frac{\frac{d(\log(IDC(\tau)))}{d(\log(\tau))} + 1}{2} \quad (3)$$

as the *Index of Variability* of Y for the time scale τ , where $\frac{d(\log(IDC(\tau)))}{d(\log(\tau))}$ is the local slope of the IDC curve at each τ when plotted in log-log coordinates.

Note that the index of variability is so defined in order for an asymptotically or second-order self-similar process $H_v(\tau) \rightarrow H \in (0.5, 1)$ as $\tau \rightarrow \infty$. If the process is exactly self-similar, then $H_v(\tau) = H \in (0.5, 1)$ for all $\tau > 0$, that is, if $\log(IDC(\tau))$ is linear with respect to $\log(\tau)$, then $H_v(\tau)$ reduces to H . The Index of Variability can be viewed as the Hurst parameter defined at each time scale.

In general³, the process Y exhibits significant variability for those time scales τ such that $0.5 < H_v(\tau) < 1$. When $\frac{d(\log(IDC(\tau)))}{d(\log(\tau))} \rightarrow 1$, then $H_v(\tau) \rightarrow 1$, implying very high variability. A plot of $H_v(\tau)$ versus τ would depict the behavior of the traffic process Y in terms of variability (burstiness) at each time scale τ (= 10ms, 100ms, 1s, ...).

Expanding the local slope of the IDC curve at each time scale, we obtain:

$$\begin{aligned} \frac{d(\log(IDC(\tau)))}{d(\log(\tau))} &= \frac{\tau}{IDC(\tau)} \frac{d(IDC(\tau))}{d\tau} \\ &= \frac{\tau}{\text{Var}[N(\tau)]} \frac{d(\text{Var}[N(\tau)])}{d\tau} - 1. \end{aligned} \quad (4)$$

Using the above in (3), we obtain a more convenient form of the Index of Variability:

$$H_v(\tau) = 0.5\tau \left(\frac{\frac{d\text{Var}[N(\tau)]}{d\tau}}{\text{Var}[N(\tau)]} \right) \quad (5)$$

$$= \frac{1}{2} \left\{ 1 + \tau \left(\frac{\frac{d(IDC(\tau))}{d\tau}}{IDC(\tau)} \right) \right\} \quad (6)$$

Suppose now Y is an aggregate sequence of packet counts resulting from the superposition of M independent

²In log-log coordinates.

³The generality here is confined for those processes that are suitable in modeling network packet traffic.

packet-traffic sources, not necessarily identical. Then $N(t) = N_1(t) + \dots + N_M(t)$, where $N_i(t)$ denotes the number of packet arrivals in the interval $(0, t]$ from the i^{th} traffic source. Assuming again stationarity, then:

$$IDC(t) = \frac{\sum_{i=1}^M \text{Var}[N_i(t)]}{\sum_{i=1}^M \lambda_i t} = \sum_{i=1}^M \left(\frac{IDC_i(t)}{\Lambda_i} \right) \quad (7)$$

where λ_i is the mean packet arrival rate from the i^{th} source, and $\Lambda_i = \frac{\sum_{j=1}^M \lambda_j}{\lambda_i}$. In addition, $\frac{\log(IDC(t))}{\log(t)} = \frac{\log(\sum_{i=1}^M \text{Var}[N_i(t)])}{\log(t)} - \frac{\log(\sum_{i=1}^M \lambda_i t)}{\log(t)}$, and upon taking the derivative in respect to $\log(t)$ the Index of Variability for the aggregate traffic stream is computed to be:

$$\begin{aligned} H_v(\tau) &= 0.5\tau \left(\frac{\sum_{i=1}^M \frac{d\text{Var}[N_i(\tau)]}{d\tau}}{\sum_{i=1}^M \text{Var}[N_i(\tau)]} \right) \\ &= \frac{1}{2} \left\{ 1 + \tau \left(\frac{\sum_{i=1}^M \frac{d(IDC_i(\tau))}{d\tau} \left(\frac{1}{\Lambda_i} \right)}{\sum_{i=1}^M \left(\frac{IDC_i(\tau)}{\Lambda_i} \right)} \right) \right\}. \quad (8) \end{aligned}$$

As can be observed from (8), the variances or the indices of dispersion for counts of the M independent point-processes completely characterize the variability function of the aggregate packet-count sequence Y . If $\lim_{\tau \rightarrow \infty} IDC(\tau) = \lim_{\tau \rightarrow \infty} \left(\sum_{i=1}^M \left(\frac{IDC_i(\tau)}{\Lambda_i} \right) \right) = c < \infty$, then $\lim_{\tau \rightarrow \infty} H_v(\tau) = 0.5$. In the case that all M underlying point processes of making up Y are also identical, then (8) reduces to (6). If all M underlying point processes are Poisson, then $\frac{d(IDC_i(\tau))}{d\tau} = 0$ for all τ and i and hence $H_v(\tau) = 0.5$ for all τ .

III. ANALYSIS OF TRAFFIC MODELS IN TERMS OF THE INDEX OF VARIABILITY

In this section, we derive the Index of Variability functions for two traditional traffic models: two-state Markov Modulated Poisson Process (MMPP) and renewal process with hyperexponential interarrival time distributions of order two (RPH2). Two-state MMPP models have become popular for modeling the superposition of packet voice streams [31].

The work in [32] shows that long-tail distributions can be approximated by hyperexponential distributions. Thus, renewal processes with hyperexponential interarrival time distributions can be used for capturing the high variability of traffic over any range of (short or long) time scales. A major advantage of these models is their relative ease of analytically obtaining queueing performance predictions.

A. Two-state MMPP

Consider that the underlying point process of Y is an MMPP with two-state Markov chain where the mean

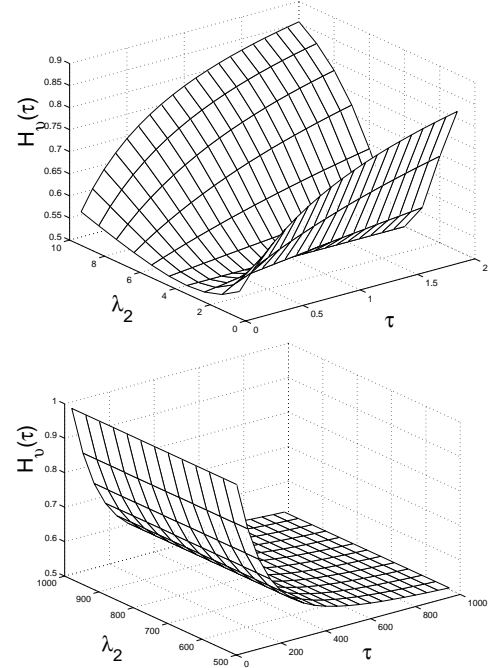


Fig. 2. Index of Variability for The Two-State MMPP: $\alpha^{-1} = \beta^{-1} = 100$ Seconds, $\lambda_1 = 4$ Packets/Second.

sojourn times in state 1 and 2 are α^{-1} and β^{-1} , respectively. When the chain is in state i ($i = 1, 2$) the point process is Poisson with rate λ_i . Letting $\eta = \alpha + \beta$ and $\nu = \lambda_1\beta + \lambda_2\alpha$, we have from [31] that $E[N(t)] = \frac{\nu t}{\eta}$ and $IDC(t) = 1 + \eta A - A \left(\frac{1 - e^{-\eta t}}{t} \right)$, where $A = \frac{2\alpha\beta(\lambda_1 - \lambda_2)^2}{\eta^3\nu}$. Clearly the $\lim_{t \rightarrow \infty} IDC(t) = 1 + \eta A$. Upon taking the derivative of $IDC(t)$, the Index of Variability of Y can be obtained:

$$H_v(\tau) = 0.5 \left\{ 1 + \frac{A [1 - (1 + \eta\tau) e^{-\eta\tau}]}{(1 + \eta A)\tau - A(1 - e^{-\eta\tau})} \right\}.$$

1) *Numerical Example:* Assume $\alpha^{-1} = \beta^{-1} = 100$ seconds, $\lambda_1 = 4$ packets/second and λ_2 to vary from 1 to 1000 packets/second. Fig. 2 shows the resulting index of variability curves as a function of time scale (τ) and state rates (λ_i). Notice that when $\lambda_2 = \lambda_1$, we have a pure Poisson process, and therefore zero variability; however, as the difference between λ_1 and λ_2 increases, so does the Index of Variability. From Fig. 2 it can be observed that the Index of Variability increases with λ_2 up to its maximum value, and any further increase in λ_2 does not have any affect on variability. It can be also observed that the Index of Variability increases with τ up to its maximum value and then decays exponentially. In addition, notice for values of λ_2 further from λ_1 , the packet-count process Y has substantial variability over a wide range of time scales that spans about 200 seconds.




Fig. 3. Squared Coefficient of Variation of the Interarrival Time vs. w_2 for the Case of Hyperexponential Distribution of Order Two: $a = 100$, and $b = 0.0001$.

B. RPH2

Assume here that the underlying point process of Y is a stationary renewal process with interarrival times hyperexponentially distributed. We call this model as the *hyperexponential model*. A hyperexponential distribution of order K , ($= 1, 2, 3, \dots$), is the weighted sum of K exponential distributions:

$$F_K(x) = Pr[X \leq x] = \sum_{i=1}^K w_i (1 - e^{-\alpha_i x}) \quad (9)$$

where $w_i > 0$ are the weights satisfying $\sum_{i=1}^K w_i = 1$, and $\alpha_i > 0$ are the rates of the exponential distributions [33]. It is shown in [34] that if $w_i = w^i$ and $\alpha_i = \frac{\mu}{\eta^i}$ for $0 < w < 1$, $\eta > 1$, and $\mu > 0$, then the tail of the hyperexponential distribution gets longer and longer with K . The major advantages of the hyperexponential distributions over heavy-tailed distributions, such as Pareto are two-fold: their Laplace transform exists, therefore they can be utilized in analytic models, and they have finite variance for all K .

In this paper, we only consider the case of $K = 2$ (i.e., the **RPH2** process). Letting $a = \alpha_1$ and $b = \alpha_2$, we get the pdf of the interarrival times to be:

$$f_2(x) = w_1 a e^{-ax} + w_2 b e^{-bx}. \quad (10)$$

The mean packet arrival rate is $\lambda = \frac{ab}{aw_2 + bw_1}$, and the squared coefficient of variation of the interarrival times is $C^2(X) = 2 \left[\frac{a^2 w_2 + b^2 w_1}{(aw_2 + bw_1)^2} \right] - 1$. Note that if $a = b$, then $\lambda = a = b$ and $C^2(X) = 1$ for all the values of w_1 and w_2 , and hence it is a Poisson process. In addition, $\lim_{w_2 \rightarrow 0} C^2(X) = 1$ and $\lim_{b \rightarrow 0} C^2(X) = \frac{2}{w_2} - 1$. As shown in Fig. 3, for constant values of a and b , $C^2(X)$ increases exponentially up to its maximum value and then decreases to one very abruptly. This indicates that the

hyperexponential distribution can be used to model the interarrival times distribution of highly bursty traffic.

From [35] we have that

$$Var[N(t)] = 2\lambda \int_0^t \Phi(u) du + \lambda t - \lambda^2 t^2 \quad (11)$$

where

$$\Phi(t) = L^{-1}[\Phi^*(s)] = L^{-1} \left[\frac{f_2^*(s)}{s\{1 - f_2^*(s)\}} \right]. \quad (12)$$

Note that the symbol L^{-1} denotes the inverse Laplace transform and

$$f_2^*(s) = L[f_2(x)] = w_1 \left(\frac{a}{s+a} \right) + w_2 \left(\frac{b}{s+b} \right)$$

is the Laplace transform of $f_2(x)$. Noting that:

$$\begin{aligned} \varphi(t) &= L^{-1} \left[\frac{f_2^*(s)}{1 - f_2^*(s)} \right] \\ &= \lambda - \frac{[(aw_1 + bw_2)^2 - (a^2 w_1 + b^2 w_2)] e^{-[aw_2 + bw_1]t}}{aw_2 + bw_1} \end{aligned} \quad (13)$$

$\Phi(t)$ is obtained:

$$\begin{aligned} \Phi(t) &= \int_0^t \varphi(u) du \\ &= \lambda t - \frac{[(aw_1 + bw_2)^2 - (a^2 w_1 + b^2 w_2)]}{(aw_2 + bw_1)^2} \\ &\quad \left(1 - e^{-[aw_2 + bw_1]t} \right). \end{aligned} \quad (14)$$

Performing the integration in (11) results:

$$\begin{aligned} Var[N(t)] &= \frac{2\lambda[(aw_1 + bw_2)^2 - (a^2 w_1 + b^2 w_2)]}{(aw_2 + bw_1)^3} \\ &\quad \left(1 - e^{-[aw_2 + bw_1]t} \right) + \lambda C^2(X)t, \end{aligned} \quad (15)$$

and hence

$$\frac{d}{dt} (Var[N(t)]) = \frac{2\lambda[(aw_1 + bw_2)^2 - (a^2 w_1 + b^2 w_2)]}{(aw_2 + bw_1)^2} e^{-[aw_2 + bw_1]t} + \lambda C^2(X), \quad (16)$$

and

$$\begin{aligned} IDC(t) &= \frac{2[(aw_1 + bw_2)^2 - (a^2 w_1 + b^2 w_2)]}{(aw_2 + bw_1)^3} \\ &\quad \left(\frac{1 - e^{-[aw_2 + bw_1]t}}{t} \right) + C^2(X). \end{aligned} \quad (17)$$

Observe that $\lim_{t \rightarrow \infty} IDC(t) = C^2(X)$, and if $a = b$ then $[(aw_1 + bw_2)^2 - (a^2 w_1 + b^2 w_2)] = 0$ and $C^2(X) = 1$ making $Var[N(t)] = \lambda t$ and $IDC(t) = 1$, i.e., we get a Poisson process. (15) or (17) can then be used in (5) or (6) to obtain the Index of Variability. It is obvious to see that $\lim_{\tau \rightarrow \infty} H_V(\tau) = 0.5$.

TABLE I

VALUES OF MEAN PACKET RATE (λ) AND SQUARED COEFFICIENT OF VARIATION OF INTERARRIVAL TIMES ($C^2(X)$) FOR THE NUMERICAL EXAMPLE OF THE CASE OF HYPEREXPONENTIAL DISTRIBUTION OF ORDER TWO: $a = 100$.

w_2	λ (packets/sec)		$C^2(X)$	
	$b = 0.01$	$b = 0.0001$	$b = 0.01$	$b = 0.0001$
10^{-3}	9.1000	0.0999	1.6522×10^3	1.9950×10^3
10^{-4}	50.0000	0.9901	5.0000×10^3	1.9605×10^4
10^{-5}	90.9000	9.0909	1.6536×10^3	1.6529×10^3
10^{-6}	99.0000	50.0000	197.0202	5.0000×10^3
10^{-7}	99.9000	90.9091	20.9561	1.6529×10^3
10^{-8}	99.9900	99.0099	2.9992	1.9607×10^4
10^{-9}	99.9990	99.9001	1.2000	1.9970×10^3
10^{-10}	99.9999	99.9900	1.0200	200.9596
10^{-11}	100.0000	99.9990	1.0020	20.9996
10^{-12}	100.0000	99.9999	1.0002	2.9999
10^{-13}	100.0000	100.0000	1.0000	1.2001

Deriving the symbolic expression of $\text{Var}[N(t)]$ for $K > 2$ is a difficult problem, mainly due to the difficulty in deriving $\phi(t)$, i.e., performing the following inverse Laplace transform:

$$L^{-1} \left[\frac{f_K^*(s)}{1 - f_K^*(s)} \right]$$

where $f_K^*(s)$ is the Laplace transform of the the K -order hyperexponential pdf of the interarrival times. However, it becomes trivial when the model parameters (e.g., w_i and α_i) are set to numerical values.

1) *Numerical Example:* Let $a = 100$. Table I lists the values of the mean packet rate (λ) and the squared coefficient of variation of the interarrival times ($C^2(X)$) for $b = 0.01$ and $b = 0.0001$ for different values of w_2 . Note that $w_1 + w_2 = 1$. Fig. 4 shows the resulted Index of Variability curves for different values of w_2 , and thus for different values of λ , and for $b = 0.0001$. Clearly, the value of w_2 has a direct impact on the variability of the resulting packet traffic. By observation, these curves can be divided into two groups, the ones shown at the left of Fig. 4 (group A) and the ones shown at the right (group B).

In group A, the curves have approximately the same rising behaviors but their falling behaviors, although having similar shapes, occur on different time scales. As $\lambda \rightarrow \frac{a}{2}$ (i.e., $w_2 \rightarrow 10^{-6}$), the range of time scales of high variability behavior and the maximum value of variability increases. On the other hand, in group B the curves have similar falling behaviors but different rising behaviors. In this case, the maximum value of variability as well as the range of time scales of substantial variability become smaller as $\lambda \rightarrow a$. If we let

$$\tau_{\text{on}} = \inf\{\text{range of time scales of substantial variability}\}$$

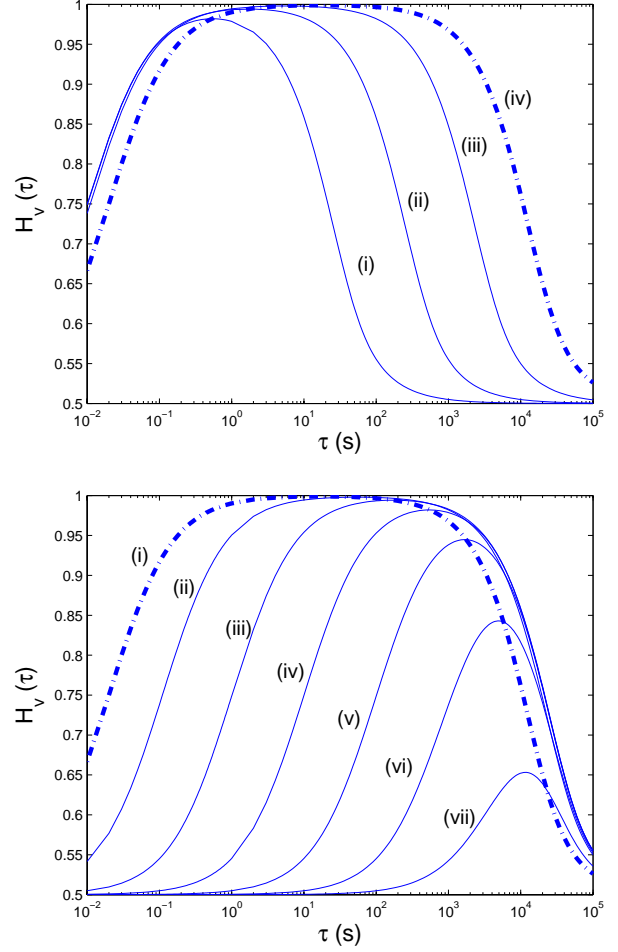


Fig. 4. Index of Variability vs. Time Scale for the Case of Hyperexponential Distribution of Order Two, $a = 100$, $b = 0.0001$: (left) $w_2 =$ (i) 10^{-3} (ii) 10^{-4} (iii) 10^{-5} (iv) 10^{-6} (right) $w_2 =$ (i) 10^{-6} (ii) 10^{-7} (iii) 10^{-8} (iv) 10^{-9} (v) 10^{-10} (vi) 10^{-11} (vii) 10^{-12} .

and

$$\tau_{\text{off}} = \sup\{\text{range of time scales of substantial variability}\},$$

then we can say that the process behaves like Poisson for all $\tau \notin [\tau_{\text{on}}, \tau_{\text{off}}]$.

As it can be observed in Fig. 4, the $w_2 = 10^{-6}$ -curve (left plot curve (iv) or right plot curve (i)) is the boundary curve between the two groups. At this value of w_2 , $\lambda = \frac{a}{2}$, the process attains the widest range of time scales (that spans 7 order of magnitude) of high variability, and in this range the index of variability reaches its maximum value ($\max H_v = 0.9988$). Interestingly, the maximum value of $C^2(X)$ occurs also at $\lambda = \frac{a}{2}$. Note that this widest range of time scales of high variability most likely covers all time scales that impact network performance evaluation [7].




Fig. 5. Normalized mean number of packets ($E[N]$) versus w_2 for various values of system load (ρ) assuming an RPH2/M/1 queueing system. The normalization is in terms of the M/M/1 queueing system.

Assuming an RPH2/M/1 queueing system⁴, we get the mean number of packets in the queue to be $E[N] = \rho/(1 - \sigma)$ where $\rho = \lambda/\mu$ is the load, μ the mean packet service rate, and

$$\sigma = \frac{(\mu + a + b - \sqrt{\mu^2 - 2(a+b)\mu + \frac{4ab}{\rho}} + (a-b)^2)}{2\mu}.$$

Fig. 5 depicts the normalized⁵ $E[N]$ as a function of w_2 and ρ (load). By studying these queueing performance curves, several interesting observations can be made: (1) even at very low loads, the mean number of packets in the queue (or the mean queueing delay) can get very large (see the curve for $\rho = 0.3$); (2) as the load increases, the normalized $E[N]$ peak increases (as expected); (3) as $\rho \rightarrow 1$, $w_2^{\text{peak}} \rightarrow 10^{-6}$, (i.e., $\lambda = \frac{a}{2}$) where w_2^{peak} is the value of w_2 where $E[N]$ attains its maximum value; (4) as w_2 decreases, $E[N]$ increases exponentially reaching its maximum value and then falls sharply to 1. Clearly, for those values of w_2 such that the normalized $E[N] \approx 1$, the queueing system behaves like M/M/1.

In addition, by careful examination of Figures 4 and 5, it can be determined that the performance of the RPH2/M/1 queue is correlated to the behavior of the Index of Variability curve at early time scales. For example, consider first the $w_2 = 10^{-5}$ -Index-of-Variability-curve. Although at $\tau \approx 10^4$ the value of H_v has fallen close to 0.5, $E[N]$ is relatively high, as shown in Fig. 5, even at

⁴RPH2/M/1 can be considered as special case of G/M/1 queue with RPH2 as the arrival process. See [37] for details about the G/M/1 queue.

⁵Normalized $E[N] = E[N]_{\text{RPH2/M/1}} / E[N]_{\text{M/M/1}}$.

a very low load. Next consider the $w_2 = 10^{-8}$ -Index-of-Variability-curve. The Index of Variability curve starts from close to 0.5 at low values of τ and reaches its maximum value around 100 seconds. Clearly, for $\tau < 1$ seconds, H_v is relatively low. Although H_v attains high variability at relatively high time scales for a wide range of time scales, the normalized $E[N]$ has values close to 1 for any load. It appears that high variability of the arrival process beyond a certain time scale does not affect the performance of the RPH2/M/1 queueing system, in agreement with [7].

Figure 6 depicts 3D Index of Variability curves generated using the hyperexponential model with $K = 2$. Clearly, both Figures 4 and 6 demonstrate that the hyperexponential model can yield a variety of Index of Variability curves. Hence, hyperexponential models can be used to model a wide range of network traffic types. Although hyperexponential models of order two (i.e., $K = 2$) are capable of generating a variety of Index of Variability curves, capturing the characteristics of traffic with multimodal Index of Variability curves would require using higher order ($K > 2$) hyperexponential models.

IV. ESTIMATING $H_v(\tau)$ FROM TRAFFIC TRACES

The estimation of the Index of Variability curve from a given traffic trace requires the estimation of the first derivative of $\text{Var}[N(\tau)]$ from discrete samples ($\text{Var}[N(\tau_i)]$, $i = 1, \dots, n$). To do this, we must first find an analytic function that best fits the discrete variance data. This in turn requires the use of an interpolation method such as polynomial-based interpolation, cubic spline and smoothing spline [39]-[43].

Since we use the sample variances as the estimates of $\text{Var}[N(\tau_i)]$, $i = 1, \dots, n$, we consider these estimates of the variances to be noisy samples. The smoothing spline interpolation methods are known to have optimal properties for estimating continuous functions and their derivatives from a finite number of noisy samples [40], [42], [43]. Note that nonsmoothing interpolation methods such as cubic spline have the characteristic that the estimated curve passes through all the given points. Hence, in case of noisy data, nonsmoothing interpolation methods yield rough curves, and therefore erroneously high first derivatives.

A. Smoothing Spline Interpolation Method

For a given data series (x_i, y_i) , $i = 1, 2, \dots, n$, the smooth function $f(x)$ is the solution of the minimization problem

$$\frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 + \xi \int_{x_1}^{x_n} (f^{(k)})^2 du, \quad (18)$$

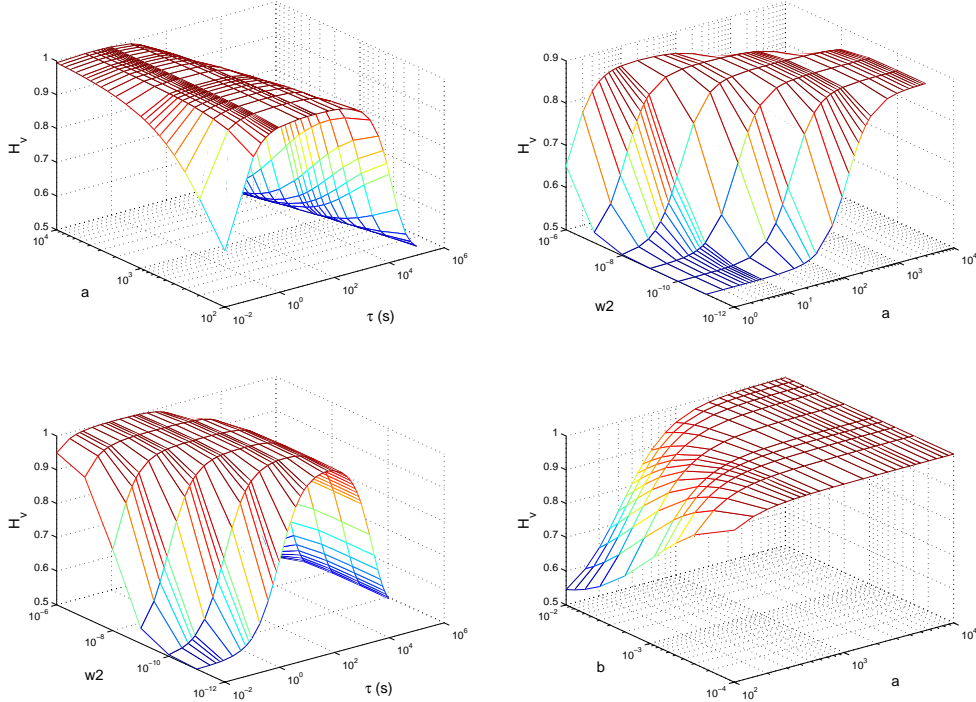


Fig. 6. Illustrations of 3D Index of Variability curves generated using the *hyperexponential* model with the following model parameter values: (top left) $b = 0.0001$, $w_2 = 10^{-6}$ (top right) $\tau = 1000$, $b = 0.001$ (bottom left) $a = 1000$, $b = 0.0001$ (bottom right) $\tau = 1$, $w_2 = 10^{-7}$.

where ξ is the smoothing parameter and $f^{(k)}$ is the k^{th} derivative of f . If $k = 2$, then f is a cubic smoothing spline.

The first term in (18) is the residual sum of squares, an indicator of the goodness-of-fit of the spline curve to the data. In other words, it measures the degree of fidelity of the smoothing spline function to the data. The second term measures the roughness of the resulting smoothing spline curve. The roughness of a function can be characterized by its curvature. For example, if a function is a straight line, then its second derivative (and therefore, roughness) is zero, that is, the second term is a penalty term measuring how close the function is to a straight line.

The smoothing parameter ξ plays an important role. It weights two aspects: smoothness and fit. Large values of ξ give a smoother curve, while small values of ξ result in a closer fit.

B. Steps for Estimating $H_v(\tau)$ from Traffic Traces

We now present a practical method for estimating the Index of Variability from traffic traces. Assuming that a given traffic trace is a realization of a second-order ergodic point process whose variance curve is continuous and differentiable. We can estimate $H_v(\tau)$ of the process as follows:

- Using the *Aggregated Variance* method [11] estimate the variance-time sequence: $\widehat{Var}[N(\tau_i)]$, $i = 1, \dots, n$.
- Using an appropriate smoothing spline implementation estimate the smoothing spline $\widehat{Var}[N(\tau)]$ from $\widehat{Var}[N(\tau_i)]$, $i = 1, \dots, n$.
- Using (5) estimate the Index of Variability $\widehat{H}_v(\tau)$.

Prior to experimentation, the accuracy and robustness of this procedure was validated by estimating and matching the Index of Variability curves shown in Fig. 4 from synthetically generated data using the *hyperexponential* traffic model.

C. Experimental Results

Using the steps outlined in the previous section, we estimated the Index of Variability curve ($H_v(\tau)$) from 12 NLANR network traffic long traces [38]. The dates at which each trace was collected and their durations are listed in Table II. For more information about these traffic traces, see [38]. Figures 7 and 8 show the estimated Index of Variability curves from the 12 long packet traces.

We used Matlab's spline toolbox to estimate all the smoothing splines. Its smoothing spline implementation is based on Reinsch's approach [40], [41]. Based on the

TABLE II

DATES AND DURATIONS OF THE NLANR NETWORK TRAFFIC TRACES. D:H:M \rightarrow DAYS:HOURS:MINUTES

Trace	Data Set	Date Collected	Duration (D:H:M)
19991129-134258-0	Auckland-II	9/29/1999	1:14:29
19991129-134258-1	Auckland-II	9/29/1999	1:14:29
19991201-192548-0	Auckland-II	12/1/1999	1:0:2
20010220-226-0	Auckland-IV	2/20/2001	6:4:58
20010220-226-1	Auckland-IV	2/20/2001	6:4:58
20010301-0310-0	Auckland-IV	3/1/2001	9:14:49
20010301-0310-1	Auckland-IV	3/1/2001	9:14:49
20010609-0613-0	Auckland-VI	6/9/2001	4:6:0
20010609-0613-1	Auckland-VI	6/9/2001	4:6:0
20010609-0613-e0	Auckland-VI	6/9/2001	4:6:0
20010609-0613-e1	Auckland-VI	6/9/2001	4:6:0
20020519-525	Bell-Lab-I	5/19/2002	7:0:0

input data, the algorithm computes the optimal smoothing parameter ξ such that the penalized residual sum of squares is less than a tolerance value $\varepsilon > 0$. In all cases we used the default value of $k (= 2)$ and $\varepsilon = 0.0001$.

V. DISCUSSION

As expected, the results displayed in Figures 7 and 8 show that the variability of real network traffic varies with time scales. An interesting observation is that the Index of Variability curves derived from the Auckland traffic traces exhibit similar monomodal behavior, while the Bell-Lab Index of Variability curve is multimodal. Hence, *hyperexponential* models of order two can be used to well approximate the Auckland traffic processes (see Figures 4 and 6), but they are not appropriate to be used to capture the characteristics of the Bell Lab traffic. Hence, the Index of Variability can be used to reveal significant differences between traffic traces.

We believe that the characteristics exhibited by the Bell Lab traffic can be captured either by *hyperexponential* models of order higher than two or by *mixprocess* models, that is the superposition of heterogeneous traffic processes. As an illustration, consider that an aggregated network traffic (packet) process is the result of the superposition of 10 renewal processes with hyperexponential interarrival time distribution of order two (RPH2), 20 two-state Markov Modulated Poisson processes (MMPP), 16 packetized voice streams, and 40 packet streams generated by ON/OFF traffic sources whose ON and OFF periods are both exponentially distributed⁶. The resulting Index of Variability curve for this *mixprocess* traffic model is shown in Fig. 9. Note that the dash-line curve is the H_v curve for the

case of aggregating only 10 identical RPH2 processes. Comparing the Index of Variability curves of Figures 7 and 9, it can be inferred that the Bell-Lab traffic can be well modeled using such a *mixprocess* traffic model.

Evidently, the Index of Variability is a valuable tool that can provide new insights and understanding of the dynamics of network traffic. H_v curves can be used to identify the characteristics of networks or applications that generate the various types of traffic patterns. An H_v curve can be regarded as a "signature" of a particular network or application. To further understand the difference between the Auckland and Bell-Lab Index of Variability curves, these two networks must be compared in terms of their architecture, protocols employed at each layer, physical links, applications, and other network specific characteristics.

VI. CONCLUSION

Knowledge of the traffic characteristics on multiple time scales can help to improve the efficiency of traffic control mechanisms. In addition, the design and provision of quality-of-service-guarantees over the Internet requires the understanding of traffic characteristics, such as variability. This paper presented an alternative measure of traffic burstiness called the *Index of Variability* ($H_v(\tau)$), that describes the degree of variability of a typical network traffic process at each time scale and is analytically tractable for many traffic models. Analytical derivations of $H_v(\tau)$ were given for two conventional traffic models. Several generated 2D and 3D Index of Variability curves predicate that the Index of Variability can help in determining the complexities of the network traffic variability over time scales that are relevant to network performance.

This paper also presented a practical method for estimating the Index of Variability curves from traffic traces. Using this method, the Index of Variability curves for 12 NLANR network traffic long traces were estimated. The results indicate that the variability of real network traffic varies with time scales. In summary, the Index of Variability offers the potential to gain insights into the dynamics of network traffic that existing tools do not offer. Future work involves developing methods of fitting analytically obtained Index of Variability functions to empirically obtained Index of Variability curves. To accomplish this, nonlinear optimization techniques will be utilized. In addition, studies to find relations that associate $H_v(\tau)$ with queuing performance metrics, such as packet loss rate and delay, will be performed.

⁶See [36] for more details.

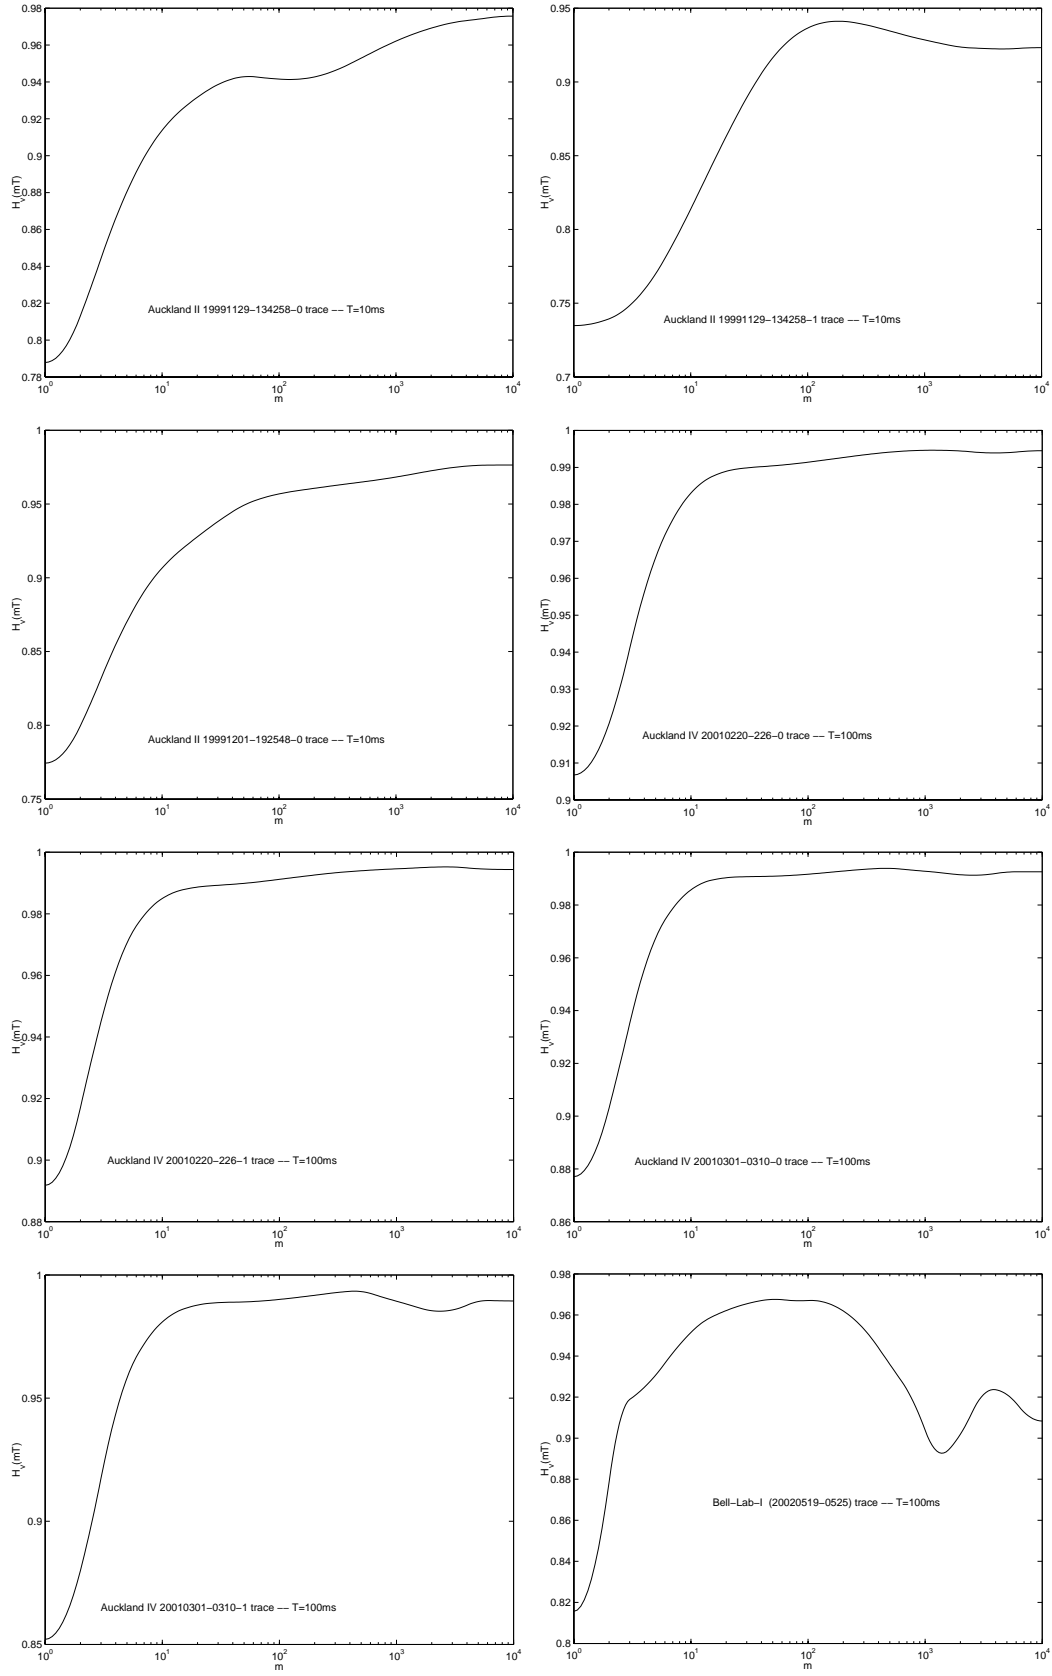


Fig. 7. Estimated Index of Variability Curves for Auckland II (top first three), Aucland IV (next four), and Bell-Lab-I (last) traces

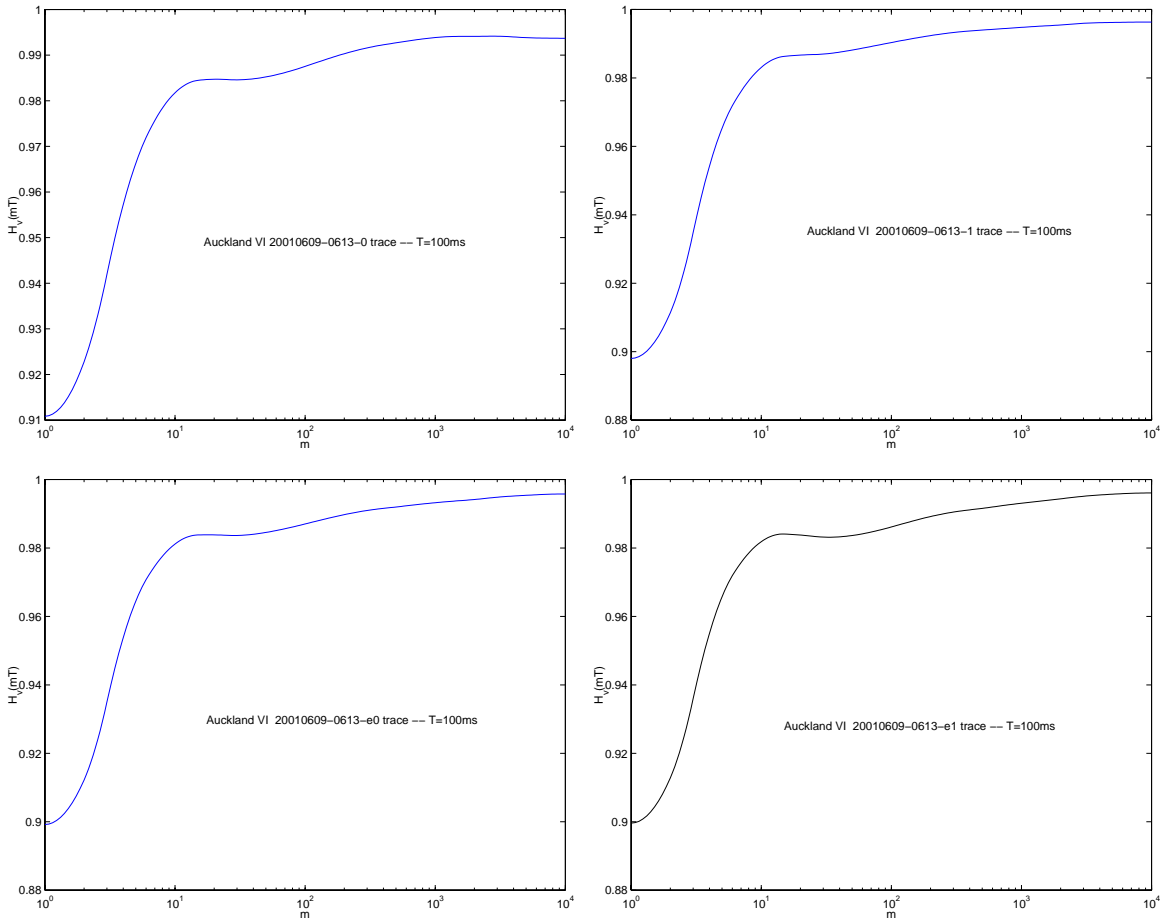


Fig. 8. Estimated Index of Variability Curves for Auckland VI traces





Fig. 9. Analytically derived Index of Variability curve for the case of *mixprocess* traffic model example.

REFERENCES

- [1] G. Y. Lazarou, X. Xiangdong and V. S. Frost, "Internet traffic modeling using the index of variability," in *Proc. IASTED International Conference on Modeling and Simulation*, 2003, pp. 31–37.
- [2] D. Duffy, A. McIntosh, M. Rosenstein and D. Wilson, "Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks," *IEEE JSAC*, vol. 12, no. 3, pp. 544–551, Apr. 1994.
- [3] H. J. Fowler and W. Leland, "Local area network traffic characteristics, with implication for network congestion management," *IEEE JSAC*, vol. 9, no. 7, pp. 1139–1149, Sept. 1991.
- [4] W. Leland, M. Taqqu, W. Willinger and D. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Trans. on Networking*, vol. 2, no. 1, pp. 1–15, Feb. 1994.
- [5] V. Paxson and S. Floyd, "Wide-area traffic: The failure of Poisson modeling," *IEEE/ACM Trans. on Networking*, vol. 3, no. 3, pp. 226–244, Jun. 1995.
- [6] A. Erramilli, O. Narayan and W. Willinger, "Experimental queueing analysis with long-range dependent packet traffic," *IEEE/ACM Trans. on Networking*, vol. 4, no. 2, pp. 209–223, Apr. 1996.
- [7] M. Grossglauser and J.-C. Bolot, "On the relevance of long-range dependence in network traffic," *IEEE/ACM Trans. on Networking*, vol. 7, no. 5, pp. 629–640, Oct. 1999.
- [8] K. R. Krishnan, A. L. Neidhardt and A. Erramilli, "Scaling analysis in traffic management of self-similar processes," in *Proc. 15th ITC*, 1997, pp. 1087–1096.
- [9] A. L. Neidhardt and J. L. Wang, "The concept of relevant time scales and its application to queueing analysis of self-similar traffic (or is Hurst naughty or nice?)," in *Proc. ACM SIGMETRICS*, 1998, pp. 222–232.
- [10] K. Park and W. Willinger, Editors, *Self-Similar Network Traffic And Performance Evaluation*, New York: John Wiley & Sons, 2000.
- [11] J. Beran, *Statistics for Long-Memory Processes. Monographs on Statistics and Applied Probability*, New York: Chapman and Hall, 1994.
- [12] M. Taqqu, V. Teverovsky and W. Willinger, "Estimators for long-range dependence: An empirical study," *Fractals*, vol. 3, no. 4, pp. 785–798, 1995.
- [13] D. Veitch and P. Abry, "A wavelet based joint estimator for the parameters of LRD," Special issue on Multiscale Statistical Signal Analysis and its Applications, *IEEE Trans. Information Theory*, vol. 45, no. 3, Apr. 1999.
- [14] A. Baiocchi, N. Melazzi, M. Listanti, A. Roveri and R. Winkler, "Modeling issues on an ATM multiplexer within a bursty traffic environment," in *Proc. IEEE INFOCOM*, 1991, pp. 83–91.
- [15] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.
- [16] W. Leland and D. Wilson, "High time-resolution measurement and analysis of LAN traffic: Implication for LAN interconnection," in *Proc. IEEE INFOCOM*, 1991, pp. 1360–1366.
- [17] N. G. Duffield, J. T. Lewis, N. O'Connell, R. Russell and F. Toomey, "Statistical issues raised by the Bellcore data," in *Proc. of the 11th Teletraffic Symposium on Performance Engineering in Telecommunications Networks*, 1994, pp. 1–8.
- [18] N. G. Duffield, J. T. Lewis, N. O'Connell, R. Russell and F. Toomey, "Predicting quality of service for traffic with long-range fluctuations," in *Proc. IEEE ICC*, 1995, pp. 473–477.
- [19] K. R. Krishnan, "A new class of performance results for a fractional brownian traffic model," *Queueing Systems*, vol. 22, pp. 277–285, 1996.
- [20] Q. Li and D. Mills, "Investigating the scaling behavior, crossover and anti-persistence of Internet packet delay dynamics," in *Proc. IEEE Globecom*, 1999, pp. 1843–1852.
- [21] V. J. Ribeiro, R. Riedi and R. G. Baraniuk, "Wavelets and multifractals for network traffic modeling and inference," in *Proc. IEEE ICASSP*, 2001, pp. 3429–3432.
- [22] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi and D. Veitch, "Multiscale nature of network traffic," *IEEE Signal Processing Magazine*, vol. 19, no. 3, pp. 28–46, May 2002.
- [23] K. Kant, "On aggregate traffic generation with multifractal properties," in *Proc. IEEE Globecom*, 1999, pp. 801–804.
- [24] R. Riedi, M. S. Crouse, V. J. Ribeiro and R. G. Baraniuk, "A multifractal wavelet model with application to network traffic," *IEEE Trans. on Information Theory*, vol. 45, no. 3, pp. 992–1018, Apr. 1999.
- [25] J. L. Veהל and B. Sikdar, "A multiplicative multifractal for TCP traffic," in *Proc. of the 6th IEEE Symposium on Computers and Communications*, 2001, pp. 714–719.
- [26] A. C. Gilbert, W. Willinger and A. Feldmann, "Scaling analysis of conservative cascades, with applications to network traffic," *IEEE Trans. on Information Theory*, vol. 45, no. 3, pp. 971–991, Apr. 1999.
- [27] M. Taqqu, V. Teverovsky and W. Willinger, "Is network traffic self-similar or multifractal?" *Fractals*, vol. 5, pp. 63–73, 1997.
- [28] A. Horvath and M. Telek, "A markovian point process exhibiting multifractal behaviour and its application to traffic modeling," in *Proc. 4th International Conference Matrix Analytic Methods in Stochastic Models*, 2002.
- [29] D. Veitch, P. Abry, P. Flandrin and P. Chainais, "Infinitely divisible cascade analysis of network traffic," in *Proc. IEEE ICASSP*, 2000, pp. 245–248.
- [30] W. Willinger, V. Paxson, R. Riedi and M. Taqqu, "Long-range dependence and data network traffic," in *Long-range Dependence: Theory and Applications*, P. Doukhan, G. Oppenheim, M. Taqqu, Eds., Birkhauser, 2002.
- [31] H. Heffes and D. M. Lucantoni, "A markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance," *IEEE JSAC*, vol. SAC-4, no. 6, pp. 856–868, Sept. 1986.
- [32] A. Feldman and W. Whitt, "Fitting mixtures of exponentials to long-tail distributions to analyze network performance models," *Performance Evaluation*, vol. 31, no. 8, pp. 963–976, Aug. 1998.
- [33] R. Gusella, "Characterizing the variability of arrival processes with indexes of dispersion," *IEEE JSAC*, vol. 9, no. 2, pp. 203–211, Feb. 1991.
- [34] M. Greiner, M. Jobmann and L. Lipsky, "The importance of power-tail distributions for telecommunications traffic modeling," Institute für Informatik, Technische Universität München, München, Germany, Tech. Rep., 1995.
- [35] D. R. Cox, *Renewal Theory*, London: J. Wiley & Sons, 1962.
- [36] G. Y. Lazarou, *On the Variability of Internet Traffic*, Ph.D. Dissertation, Electrical Engineering and Computer Science, University of Kansas, 2000.
- [37] L. Kleinrock, *Queueing Systems, Volume I: Theory*, New York: J. Wiley & Sons, 1975.
- [38] J. Micheel, NLANR PMA: Special Traces Archive, [Online]. Available: <http://pma.nlanr.net/Traces/long/>
- [39] R. L. Burden and J. D. Faires, *Numerical Analysis 7th ed.*, Brooks/Cole, 2001.
- [40] C. M. Reinsch, "Smoothing by spline functions," *Numerische Mathematik*, vol. 10, pp. 177–183, 1967.
- [41] R. L. Eubank, *Spline Smoothing and Nonparametric Regression*, New York: Marcel Dekker, 1988.
- [42] J. Fox, *Nonparametric Simple Regression: Smoothing Scatterplots*, Californian: Sage Publications, Thousand Oaks, 2000.
- [43] B. Shahraray and D. J. Anderson, "Optimal estimation of contour properties by cross-Validated regularization," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 11, no. 6, pp. 600–610, Jun. 1989.



Georgios Y. Lazarou is currently an Assistant Professor in the Electrical and Computer Engineering Department at Mississippi State University. He is a founding member of the Telecommunication and Information Technology Laboratory (TITL) at Mississippi State University. Currently, he is a member of the Intelligent Electronic Systems (IES) program at CAVS. His current research interests lie in

the area of next generation smart networking technologies, wireless networking for intelligent transportation systems, and energy efficient protocols for wireless sensor networks. He has been involved in research projects concerning the design of adaptive and reconfigurable networks, modeling and evaluating the performance of high-speed wide-area asynchronous transfer mode (ATM) networks and Internet, and characterization of network traffic. Results from his work have been published in numerous journal and conference articles, and in technical reports.



Joseph B. Evans is the Deane E. Ackers Distinguished Professor of Electrical Engineering & Computer Science and Director of Research Information Technology at the University of Kansas. He recently served as a Program Director in the Division of Computer & Network Systems in the Directorate for Computer & Information Science & Engineering at the National Science Foundation. His research inter-

ests include mobile and wireless networking, pervasive computing systems, high speed networks, and adaptive computing systems. He has been involved in major national high performance networking testbeds and broadband wireless mobile networking efforts, and has published over 100 journal and conference works. Dr. Evans received his PhD degree from Princeton University in 1989, is a senior member of the IEEE, and a member of the ACM.



Julie Baca is currently a Research Assistant Professor at the Center for Advanced Vehicular Systems at the Mississippi State University. Prior to this, Dr. Baca was a member of the computer science research group of the Information Technology Laboratory at U.S. Army Corps of Engineers. Currently, Dr. Baca directs the mobile computing usability and the in-vehicle dialogue interface projects with an

emphasis on voice over IP. She has over a decade of experience in designing complex software system interfaces for voice and multimedia applications over the Internet. Her research interests include mobile computing, wireless networking, operating systems, and user interface design in mobile environments. Dr. Baca received her Ph.D. degree from Mississippi State University in 1998, and is a member of ACM.



Victor S. Frost is currently the Dan F. Servey Distinguished Professor of Electrical Engineering and Computer Science and Director of the University of Kansas Telecommunications and Information Technology Center (ITTC). He is a Fellow of the IEEE and received a Presidential Young Investigator Award from the National Science Foundation in 1984. His current research interest is in the areas of inter-

net quality of service, traffic management, and integrated broadband communication networks. He has been involved in research on several national scale high speed wide area testbeds. Government agencies, including, NSF, DARPA, Rome Labs, and NASA have sponsored his research. Professor Frost has been involved in research for numerous corporations, including Harris, Sprint, NCR, BNR, Telesat Canada, AT&T, McDonnell Douglas, DEC, and COMDISCO Systems. He has published over 100 journal articles and conference papers. Dr. Frost received the BS, MS, and Ph.D. degrees from the University of Kansas, Lawrence in 1977, 1978, and 1982, respectively. In 1982 he joined the faculty of the University of Kansas