INTERNATIONAL JOURNAL OF COMMUNICATION SYSTEMS
Int. J. Commun. Syst. 2001; **14**:783–801 (DOI: 10.1002/dac.509)

Multiplicative multifractal modelling of long-range-dependent network traffic

Jianbo Gao and Izhak Rubin*

Department of Electrical Engineering, University of California, Los Angeles (UCLA), Los Angeles, California 90095, U.S.A.

SUMMARY

We present a multiplicative multifractal process to model traffic which exhibits long-range dependence. Using traffic trace data captured by Bellcore from operations across local and wide area networks, we examine the interarrival time series and the packet length sequences. We also model the frame size sequences of VBR video traffic process. We prove a number of properties of multiplicative multifractal processes that are most relevant to their use as traffic models. In particular, we show these processes to characterize effectively the long-range dependence properties of the measured processes. Furthermore, we consider a single server queueing system which is loaded, on one hand, by the measured processes, and, on the other hand, by our multifractal processes (the latter forming a $MF_e/MF_g/1$ queueing system model). In comparing the performance of both systems, we demonstrate our models to effectively track the behaviour exhibited by the system driven by the actual traffic processes. We show the multiplicative multifractal process to be easy to construct. Through parametric dependence on one or two parameters, this model can be calibrated to fit the measured data. We also show that in simulating the packet loss probability, our multifractal traffic model provides a better fit than that obtained by using a fractional Brownian motion model. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS: performance modelling; multiplicative multifractal; network traffic; analytic analysis

1. INTRODUCTION

Recent analysis of high-quality traffic measurements have revealed the prevalence of long-range-dependent (LRD) (or self-similar) features in traffic processes loading packet switching communications networks. Included are local area networks (LANs) [1], wide area networks (WANs) [2], variable-bit-rate (VBR) video traffic [3,4], and world wide web (WWW) traffic [5]. These findings have greatly challenged the commonly assumed models for network traffic; e.g. Poisson- or Markov-based processes. Since a LRD traffic process exhibits bursts over many or all

* Correspondence to: Izhak Rubin, Department of Electrical Engineering, University of California, Los Angeles (UCLA), Los Angeles, California 90095, U.S.A.

Contract/grant sponsor: UC MICRO/SBC Pacific Bell; contract/grant number: 98-131
Contract/grant sponsor: ARO; contract/grant number: DAAGIJ-98-1-0338

Received 1 February 2000
Revised 28 March 2001
Accepted 1 May 2001

Published online 25 June 2001

time scales, while a Poisson or Markovian process, which displays short-range dependence, exhibit burstiness over much shorter time scales. As a result, Poisson or Markovian models tend to yield overly optimistic performance predictions.

Recent modelling works have therefore focused on obtaining parsimonious models capable of capturing the basic LRD property of traffic processes. Such approaches include chaotic maps [6], a LRD ON/OFF model [7], Cox's $M/G/\infty$ type models [8–11], the fractional Brownian motion (FBm) model [12,13], fractional autoregressive integrated moving-average (FARIMA) models [1,14], point processes [15], and pseudo models [16,17]. An issue of much interest is whether and how multifractal can be employed to model LRD traffic. Recently, Taqqu *et al.* [18] have analysed aggregated network traffic processes using the multifractal concept. They conclude that when self-similar traffic models can be applied, multifractal models may not be needed.

The concept of multifractal is mostly developed in understanding the intermittent features of turbulence [19]. Intermittency, when paraphrased to a meaning appropriate for network traffic, is a combination of burstiness of traffic and the variation of the burstiness over time. Hence, it is conceivable that a simple multifractal model would suffice to capture the basic (and possibly time varying) bursty features of LRD traffic. In this paper, we show that, based on a multiplicative process structure, two simple multifractals, one used to model interarrival time series, and the other used for packet size sequences, each characterized by only one or two parameters, can be readily constructed based on analysis of the measured traffic trace data. We show that these models provide excellent descriptions for LAN, WAN, and VBR video traffic processes.

We have chosen Bellcore's Ethernet traffic data and VBR video traffic data[†] for use in this study. Four Ethernet data sets, denoted as pAug.TL, pOct.TL, OctExt.TL, and OctExt4.TL, and a VBR video data, denoted as MPEG.data, have been made available. Each Ethernet data set contains 1 million points representing measured values for arrival time stamps and packet sizes. Two Ethernet data sets (pAug.TL and pOct.TL) were measured on the 'purple cable', involving LAN traffic. Another two Ethernet sets (OctExt.TL and OctExt4.TL) were collected on Bellcore's link to the outside world, and have been classified as WAN traffic [18]. These sets cover time spans of 3142.8, 1759.6, 122797.8, and 75943.1 s, respectively. The video data consists of 174136 integers, representing the number of bits per video frame (24 frames/s for approximately 2 h). For the study presented in this paper, we select a LAN traffic set (pAug.TL), a WAN traffic set (OctExt.TL), and the video traffic data (MPEG.data). We examine the interarrival time series (which is deterministic for the video traffic involving frame units) and the packet length sequences derived from these data sets. We show the multiplicative multifractal processes to effectively characterize the long-range dependence properties of these processes.

The remaining of the paper is organized as follows. In Section 2, after overviewing the definition of self-similar stochastic processes, we show that the interarrival time series and the packet length sequences derived from Bellcore's LAN and WAN traffic trace data are self-similar. We then describe in Section 3 a class of multifractals, namely, random multiplicative processes. To ease construction of multifractal models from measured traffic trace data, we prove a number of properties of multiplicative processes, and demonstrate that multiplicative processes can possess self-similar properties. We proceed in Section 4 with a detailed analysis of Bellcore's LAN, WAN, and video traffic data. We show that the interarrival time series and the packet length sequences of the Ethernet traffic, as well as the frame size sequences of the video traffic, exhibit

[†] This is available at ftp.bellcore.com under the directory/pub/world/wel/lan_traffic and /pub/vbr.video.trace, respectively.

stochastic features which are consistent with the stochastic behaviour of random multiplicative processes. In Section 5, we consider a single server queueing system which is loaded, on one hand, by the measured processes, and, on the other hand, by our multifractal processes. In comparing the performance of both systems, we demonstrate our model to effectively track the behaviour exhibited by the system driven by the actual traffic processes. We furthermore show in Section 6 that in modelling the above-mentioned processes, the multifractal traffic model yields better queue size tail probability fit than that attained by a FBm model. We summarize our findings in Section 7.

2. SELF-SIMILAR INTERARRIVAL TIME SERIES AND PACKET LENGTH SEQUENCES OF NETWORK TRAFFIC

In this section, we overview the definition of second order self-similar processes, for the purpose of showing that the interarrival time series and the packet length sequences derived from Bellcore's LAN and WAN traffic trace data are self-similar. These descriptions also serve us in a later section where we show multifractals to exhibit self-similar properties. The following definitions follow those made by Leland *et al.* [1] and Beran *et al.* [3].

Let $X = \{X_i: i = 0, 1, 2, \ldots\}$ be a covariance stationary stochastic process with mean μ, variance σ^2 and autocorrelation function $r(k)$, $k \geqslant 0$. Assume $r(k)$ to be of the form

$$r(k) \sim k^{-\beta}, \quad \text{as } k \to \infty \tag{1}$$

where $0 < \beta < 1$. Note that $\sum_k r(k) = \infty$. This is referred to as the LRD property.

In characterizing the self-similarity of such processes, the most important parameter identified is the Hurst parameter H. It is equivalent to β, $1/2 < H = 1 - \beta/2 < 1$. The value of H measures the degree of persistence of the correlation: the larger the H value, the more persistent the correlation is.

For each $m = 1, 2, 3, \ldots$, let $X^{(m)} = \{X_i^{(m)}: i = 1, 2, 3, \ldots\}$ denote the new covariance stationary time series obtained by averaging the original series X over non-overlapping blocks of size m, i.e.

$$X_i^{(m)} = (X_{im-m+1} + \cdots + X_{im})/m, \quad i \geqslant 1 \tag{2}$$

Self-similarity for X means that the process $X^{(m)}$ exhibits (exactly or asymptotically) the same second-order statistics as those characterizing the process X.

An efficient method used to detect the self-similarity character of a stochastic process employs the variance-time relation [1]: $\mathrm{Var}(X^{(m)}) \sim m^{-\beta}$. Thus, in plotting the variance of $X^{(m)}$ vs the aggregating block size m on a log–log scale, we expect self-similar processes to yield a curve which tends to become linear for large m, having a slope larger than -1. By contrast, for a short-range-dependent process, we have $\mathrm{Var}(X^{(m)}) \sim m^{-1}$.

Two random processes, the interarrival time series, $\{T_i\}$, where T_i denote the ith interarrival time between two successive packet arrivals, and the packet length sequences, $\{B_i\}$, where B_i represents the length of the ith packet, are used to define a traffic process. Two other random counting processes can be derived from these two time series. The latter processes represent the number of packets and the number of bytes arriving during successive periods each of length Δt seconds (the frame size sequences of the video traffic described above is such an example).

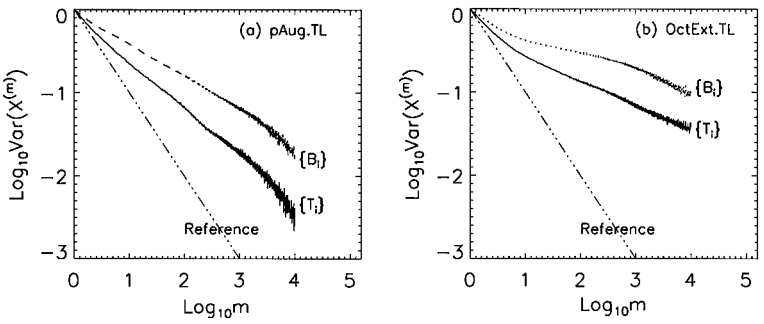

Figure 1. Variance–time plots (in logarithmic scale) for interarrival time series $\{T_i\}$ and packet size sequences $\{B_i\}$ for traffic traces (a) pAug.TL, and (b) OctExt. The line denoted as 'Reference' has slope -1.

Following Leland *et al.* [1], self-similarity has been detected focusing on the analysis of these counting processes. In contrast, we study in this paper the processes $\{T_i\}$ and $\{B_i\}$. Using the traffic measurements mentioned in Section 1, we show these processes to be self-similar. This is demonstrated by the results shown in Figure 1. We will furthermore show in this paper that simple multifractal models for the $\{T_i\}$ and $\{B_i\}$ sequences can be readily constructed. Note that $T_i = 1/24$ s, for the video traffic.

3. MULTIPLICATIVE MULTIFRACTALS

In this section, we first overview the definition of multifractals. Multifractals are typically constructed through multiplicative processes. In the following, we define multiplicative processes and present examples of such models. We then prove a number of properties of such processes that are most relevant to their use as network traffic models.

3.1. Definition

Consider a unit interval. Associate it with a unit mass. Partition the unit interval into a series of small intervals, each of linear length ε. Also partition the unit mass into a series of weights $\{w_i\}$, and associate w_i with the ith interval. Now consider the moments

$$M_q(\varepsilon) = \sum_i w_i^q \tag{3}$$

where q is real. Note the convention that whenever w_j is zero, the term w_j^q is dropped. We also note that a positive q value emphasizes large weights, while a negative q value emphasizes small weights. If we have, for a real function $\tau(q)$ of q,

$$M_q(\varepsilon) \sim \varepsilon^{\tau(q)} \quad \text{as } \varepsilon \to 0 \tag{4}$$

for every q, and the weights $\{w_i\}$ are non-uniform, then the weights $w_i(\varepsilon)$ are said to form a multifractal measure. Note that the normalization $\sum_i w_i = 1$ implies that $\tau(1) = 0$.

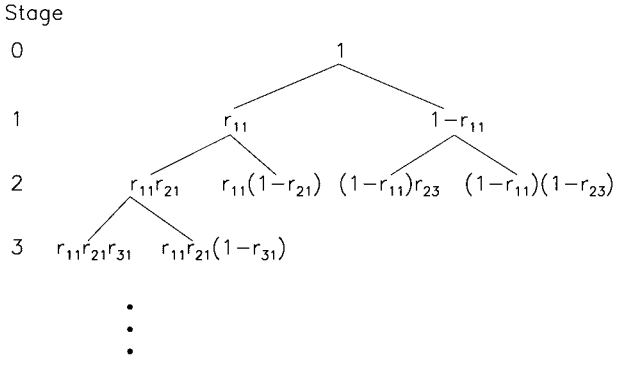

Figure 2. A schematic illustrating the construction rule of a multiplicative multifractal.

Note that if $\{w_i\}$ are uniform, then $\tau(q)$ is linear in q. When $\{w_i\}$ are weakly non-uniform, visually $\tau(q)$ may still be approximately linear in q. The nonuniformity in $\{w_i\}$ is better characterized by the so-called generalized dimensions D_q defined as [20]:

$$D_q = \frac{\tau(q)}{q-1} \tag{5}$$

D_q is a monotonically decreasing function of q [21]. It exhibits a non-trivial dependence on q when the weights $\{w_i\}$ are non-uniform.

3.2. *Construction of multiplicative multifractals*

Consider a unit interval. Associate it with a unit mass. Divide the unit interval into two segments of equal length. Also partition the associated mass into two fractions, r and $1-r$, and assign them to the left and right segments, respectively. The parameter r is in general a random variable, governed by a probability density function (pdf) $P(r)$, $0 \leqslant r \leqslant 1$. The fraction r is called the multiplier. Each new subinterval and its associated weight are further divided into two parts following the same rule. The procedure is schematically shown in Figure 2, where the multiplier r is written as r_{ij}, with i indicating the stage number. Note the scale (i.e. the interval length) associated with stage i is 2^{-i}. We assume that $P(r)$ is symmetric about $r = 1/2$, and has successive moments $\mu_1, \mu_2, \ldots$. Hence r_{ij} and $1 - r_{ij}$ both have marginal distribution $P(r)$. The weights at the stage N, $\{w_n, n = 1, \ldots, 2^N\}$, can be expressed as $w_n = u_1 u_2 \cdots u_N$, where u_l, $l = 1, \ldots, N$, are either r_{ij} or $1 - r_{ij}$. Thus, $\{u_i, i \geqslant 1\}$ are independent identically distributed random variables having pdf $P(r)$. In the following, we illustrate this process by selecting a specific pdf $P(r)$.

3.2.1. Deterministic binomial multiplicative process. In this case, the pdf is set to be equal to $P(r) = \delta(r - p)$, where $\delta(x)$ is the Kronecker delta function. Thus, $r = p$ with probability 1, where $0 < p < 1$ is a fixed number. The weights obtained for the first several stages are schematically shown in Figure 3.

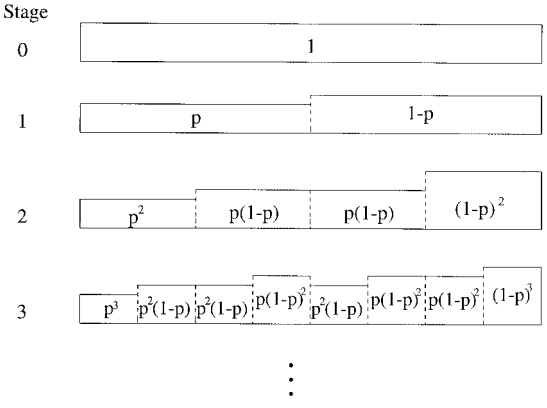

Figure 3. A schematic showing the weights at the first several stages of the binomial multiplicative process.

For this process, at stage n, we have

$$M_q(\varepsilon) = \sum_{i=0}^{n} C_n^i p^{qi}(1-p)^{q(n-i)} = (p^q + (1-p)^q)^n \tag{6}$$

Since at stage n, $\varepsilon = 2^{-n}$, we obtain

$$\tau(q) = -\ln(p^q + (1-p)^q)/\ln 2 \tag{7}$$

which is independent of n (or ε). Hence, this weight process constitutes a multifractal.

3.2.2. Random binomial multiplicative process. To make the weight series random, we modify $P(r)$ to become

$$P(r) = (\delta(r-p) + \delta(r-(1-p)))/2, \tag{8}$$

so that $P(r=p) = P(r=1-p) = 1/2$. Hence, $P(r)$ is symmetric about $r = 1/2$. A realization of the weights at stage 12 (with $p = 0.3$) is shown in Figure 4(a).

3.2.3. Random multiplicative process. The function $P(r)$ can be selected to follow any functional form [21]. The following piece-wise linear $P(r)$ function is used to generate the weight realization (at stage 12) shown in Figure 4(b):

$$P(r) = \begin{cases} 2r + 0.5 & \text{if } 0 \leqslant r \leqslant 0.5 \\ -2r + 2.5 & \text{if } 0.5 \leqslant r \leqslant 1 \end{cases}$$

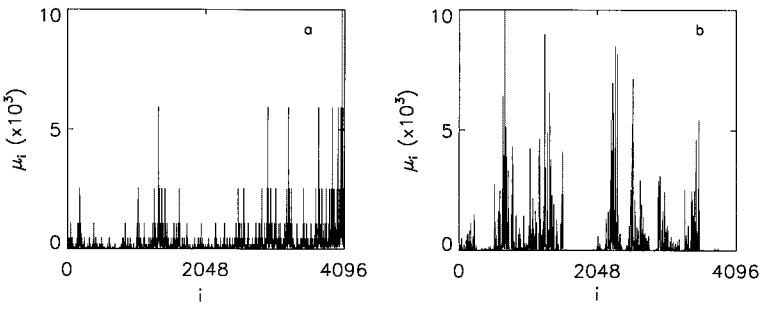

Figure 4. Weight series at stage 12 for (a) the modified binomial multiplicative process ($p = 0.3$), and (b) the random multiplicative process with the multiplier pdf given in the text.

3.3. *Properties of multiplicative multifractals*

For the weights at stage N, we prove the following properties to hold.

(i) $M_q(\varepsilon) \sim \varepsilon^{\tau(q)}$, with $\varepsilon = 2^{-N}$, $\tau(q) = -\ln(2\mu_q)/\ln 2$. This follows the observation that at stage N, $M_q(\varepsilon) = E(\sum_{n=1}^{2^N}(w_n)^q) = 2^N E(w^q) = 2^N E((u_1u_2\cdots u_N)^q) = 2^N\mu_q^N$. This property indicates that a multiplicative process is a multifractal, and relates the $\tau(q)$ spectrum to the moments of the multiplier distribution.
Note that for $P(r) = [\delta(r-p) + \delta(r-(1-p))]/2$, we have $\mu_q = [p^q + (1-p)^q]/2$. Hence, $\tau(q) = -\ln[p^q + (1-p)^q]/\ln 2$. We thus note that the function $\tau(q)$ assumed by a random binomial process is the same as that exhibited by a deterministic binomial process (Equation (7)). This is expected to be the case since the value assumed by $M_q(\varepsilon)$, and thus also $\tau(q)$, is independent of the specific ordering of the weights.

(ii) $E(w) = E(w_n) = E(u_1u_2\cdots u_N) = 2^{-N}$, $n = 1, \ldots, 2^N$.

(iii) $\mathrm{Var}(w) = \mathrm{Var}(w_n) = \mu_2^N - 2^{-2N}$, $n = 1, \ldots, 2^N$. To prove this relation, we note that $E(w^2) = E(w_n^2) = E((u_1u_2\cdots u_N)^2) = \mu_2^N$.

(iv) When $N \gg 1$, the weights at stage N have log-normal distribution. This is deduced directly by taking the logarithm of $w_n = u_1u_2\cdots u_N$.

(v) $E[(w_n - E(w))(w_{n+m} - E(w))] = (1/2 - \mu_2)\mu_2^{N-1}(4\mu_2)^{-k} - 2^{-2N}$, for $m = 2^k$, where k is an integer. Hence, the covariance function decays with time lag m in a power-law manner.
To prove assertion (v), we consider two weights w_{n_1} and w_{n_2} at stage N. Assume they share the same ancestor weight x at stage $N-k$, i.e. $w_{n_1} = xr\prod_{l=1}^{k-1} r_{1l}$, $w_{n_2} = x(1-r)\prod_{l=1}^{k-1} r_{2l}$, where r and $\{r_{il}, i = 1, 2, l = 1, \ldots, k-1\}$ are independent random variables with distribution $P(r)$. Then $E[(w_{n_1} - 2^{-N})(w_{n_2} - 2^{-N})] = E(x^2)E[r(1-r)]E[\prod_{l=1}^{k-1} r_{1l}r_{2l}] - 2^{-2N} = 2^{-2(k-1)}\mu_2^{N-k}(1/2 - \mu_2) - 2^{-2N}$. For $m = 2^k$, all pairs of $\{w_n, w_{n+m}\}$, for $n \geqslant 1$, share an ancestor at stage $N-k-1$. Hence, $E[(w_n - E(w))(w_{n+m} - E(w))] = 2^{-2k}\mu_2^{N-k-1}(1/2 - \mu_2) - 2^{-2N} = (1/2 - \mu_2)\,\mu_2^{N-1}(4\mu_2)^{-k} - 2^{-2N}$.

(vi) $\mathrm{Var}(W^{(m)}) = \mu_2^N(4\mu_2)^{-k} - 2^{-2N}$, where $W^{(m)} = (w_{im-m+1} + \cdots + w_{im})/m$, $m = 2^k$, $k = 1, 2, \ldots$, and $i \geqslant 1$. This is proven by expressing $W^{(m)} = 2^{-k}x$, where x is a weight at stage $N-k$.

The equation in (vi) expresses a variance-time relation. For LRD traffic [16], $\mathrm{Var}(W^{(m)}) \sim m^{2H-2}$, where $1/2 < H < 1$ is the Hurst parameter. For multiplicative multifractal

processes, when N is large and $\mu_2 > 0$, the term $\mu_2^N(4\mu_2)^{-k}$ dominates. Hence, the functional variation of $\log \mathrm{Var}(W^{(m)})$ vs $\log m$ is linear, with the resulting slope, $-\log(4\mu_2)/\log 2$, providing an estimate of $2H - 2$. The linear property is demonstrated by Figure 5 which shows the variance-time plots for the times series of Figure 4.

Also note that the dependence of $\mathrm{Var}(W^{(m)})$ on m is the same as that of $E[(w_n - E(w))(w_{n+m} - E(w))]$ on m. Below we show that, by analysing measured traffic processes, one can effectively model LRD traffic by using multiplicative multifractals by properly selecting the multiplier distribution $P(r)$.

4. MULTIPLICATIVE MULTIFRACTAL ANALYSIS OF LAN, WAN, AND VBR VIDEO TRAFFIC

Our purpose of multifractal analysis of network traffic is to check whether the interarrival time series $\{T_i\}$ and packet length sequences $\{B_i\}$ of the traffic data can be viewed as realizations of multiplicative processes. If they are only approximately multifractals, an equivalent multifractal model may still be constructed. In this section, we show that the interarrival time series and packet length sequences of Bellcore's LAN and WAN traffic data, and the frame size sequences of the video traffic data, exhibit stochastic features which are consistent with the stochastic behaviour of random multiplicative processes.

There are two ways to check whether $\{T_i\}$ and $\{B_i\}$ are realizations of certain multiplicative processes. One method is to compute the moments $M_q(\varepsilon)$ at different stages, and check whether Equation (4) is valid for certain ε ranges. Another method is to compute the multiplier distributions at different stages, and check whether they are stage independent. We note the latter method to be typically more useful when constructing a multiplicative process for $\{T_i\}$ or $\{B_i\}$. In the following, we describe in detail a general procedure for obtaining weight sequences at different stages needed for computing the moments $M_q(\varepsilon)$ and the multiplier distributions.

Assume there are 2^N arrivals. For ease of illustration, we denote $\{T_i\}$ or $\{B_i\}$ by $\{X_i\}$. We view $\{X_i, i = 1, \ldots, 2^N\}$ as the weight series of a certain multiplicative process at stage N. Note that the total weight $\sum_{i=1}^{2^N} X_i$ is set equal to 1 unit. Also note the scale associated with stage N is $\varepsilon = 2^{-N}$. This is the smallest time scale resolvable by the measured traffic data.

Given the weight sequence at stage N (which represents the measured data), the weights at stage $N - 1$, $\{X_i^{(2^1)}, i = 1, \ldots, 2^{N-1}\}$, is obtained by simply adding the consecutive weights at stage N over non-overlapping blocks of size 2, i.e. $X_i^{(2^1)} = X_{2i-1} + X_{2i}$, for $i = 1, \ldots, 2^{N-1}$, where the superscript 2^1 for $X_i^{(2^1)}$ is used to indicate that the block size used for the involved summation at stage $N - 1$ is 2^1. Associated with this stage is the scale $\varepsilon = 2^{-(N-1)}$. This procedure is carried out recursively. That is, given the weights at stage $j + 1$, $\{X_i^{(2^{N-j-1})}, i = 1, \ldots, 2^{j+1}\}$, we obtain the weights at stage j, $\{X_i^{(2^{N-j})}, i = 1, \ldots, 2^j\}$, by adding consecutive weights at stage $j + 1$ over non-overlapping blocks of size 2, i.e.

$$X_i^{(2^{N-j})} = X_{2i-1}^{(2^{N-j-1})} + X_{2i}^{(2^{N-j-1})} \tag{9}$$

for $i = 1, \ldots, 2^j$. Here the superscript 2^{N-j} for $X_i^{(2^{N-j})}$ is used to indicate that the weights at stage j can be equivalently obtained by adding consecutive weights at stage N over non-overlapping blocks of size 2^{N-j}. Associated with stage j is the scale $\varepsilon = 2^{-j}$. This procedure stops at stage 0,

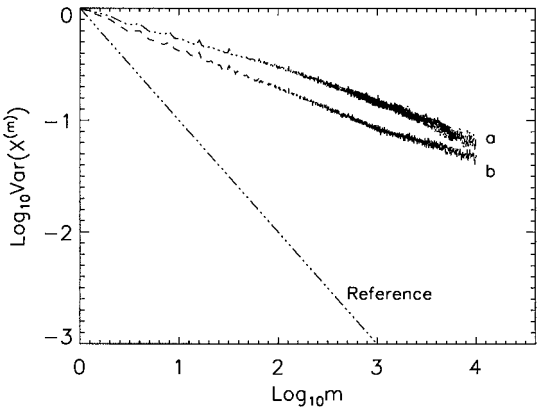

Figure 5. Variance–time plots for the time series of Figure 3, (a) the modified binomial multiplicative process, and (b) the random multiplicative process. The length for both time series is 2^{19}. The line denoted as 'Reference' has slope -1.

where we have a single unit weight, $\sum_{i=1}^{2^N} X_i$, and $\varepsilon = 2^0$. The latter is the largest time scale associated with the measured traffic data. Figure 6 schematically shows this procedure.

After we have obtained all the weights from stages 0 to N, we compute the moments $M_q(\varepsilon)$ according to Equation (3) for different values of q. We then plot $\log M_q(\varepsilon)$ vs $\log \varepsilon$ for different values of q. If these curves are linear over wide ranges of ε, then these weights are consistent with a multifractal measure. Note that, according to Equation (4), the slopes of the linear part of $\log M_q(\varepsilon)$ vs $\log \varepsilon$ curves provide an estimate of $\tau(q)$, for different values of q.

Next, we explain how to compute the multiplier distributions at different stages. From stage j to $j+1$, the multipliers are defined by the following equation, based on Equation (9):

$$r_i^{(j)} = \frac{X_{2i-1}^{(2^{N-j-1})}}{X_i^{(2^{N-j})}} \tag{10}$$

for $i = 1, \ldots, 2^j$. We view $\{r_i^{(j)}, i = 1, \ldots, 2^j\}$ as sampling points of the multiplier distribution at stage j, $P_j = \{P_j(r), 0 \leqslant r \leqslant 1\}$. Hence P_j can be determined from its histrogram based on $\{r_i^{(j)}, i = 1, \ldots, 2^j\}$. We then plot $P_j(r)$ vs r for different stages (j) together. If these curves collapse

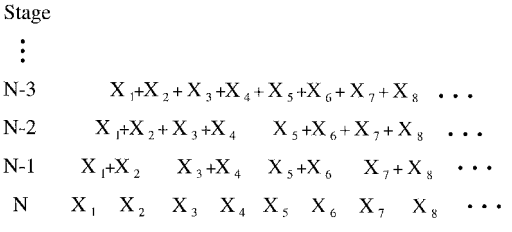

Figure 6. A schematic showing the weights at the last several stages for the analysis procedure described in the text.

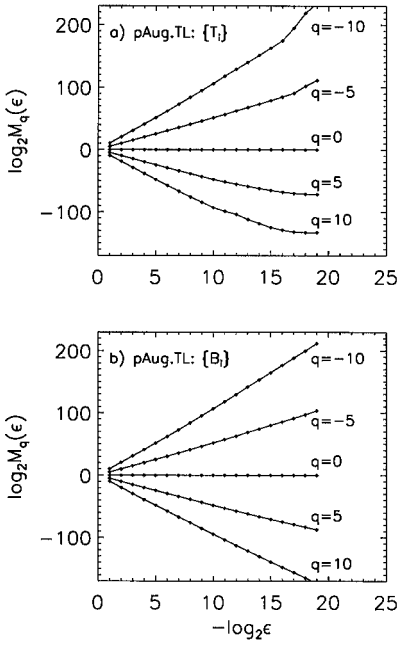

Figure 7. $\log_2 M_q(\varepsilon)$ vs $-\log_2 \varepsilon$ for $\{T_i\}$ and $\{B_i\}$ of pAug.TL for several different q's.

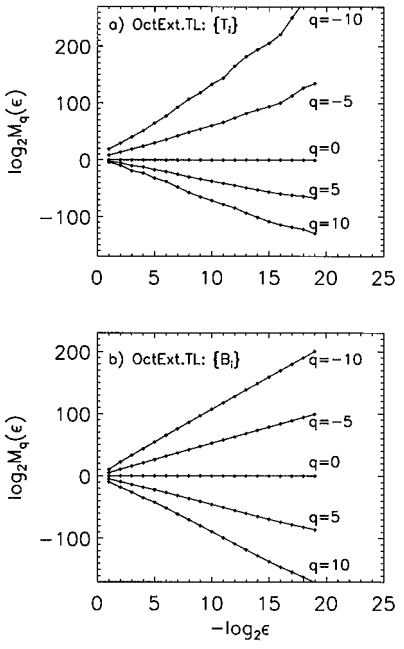

Figure 8. $\log_2 M_q(\varepsilon)$ vs $-\log_2 \varepsilon$ for $\{T_i\}$ and $\{B_i\}$ of OctExt.TL for several different q's.

together so that $P_j \sim P$, then the multiplier distributions are stage independent, and the weights form a multifractal measure P.

We illustrate the above procedure by analysing Bellcore's traffic trace data pAug.TL, OctExt.TL, and MPEG.data. We use the first 2^{19} arrivals of the Ethernet trace data, and the first 2^{17} video frame size data for this analysis. Figures 7, 8 and 10(a) show $\log_2 M_q(\varepsilon)$ vs $-\log_2 \varepsilon$ for processes $\{T_i\}$ and $\{B_i\}$ of pAug.TL and OctExt.TL, and the frame size sequences of the video traffic data, respectively. We observe that the scaling between $M_q(\varepsilon)$ and ε (i.e. the degree of linearity between $\log_2 M_q(\varepsilon)$ and $-\log_2 \varepsilon$) for $\{B_i\}$ of the Ethernet data sets are excellent. While the scaling between $M_q(\varepsilon)$ and ε for $\{T_i\}$ of the Ethernet data and the frame size sequences of the video traffic is slightly worse, yet it is still quite good. To further check whether these data sets are truly multifractals, we compute D_q. These results are shown in Figure 9, for the Ethernet data, and Figure 10(b), for the frame size sequences of the video traffic. Indeed we observe that in all cases D_q has a nontrivial dependence on q. Therefore, we conclude that these time series are consistent with multifractals.

Next we compute the multiplier distributions at different stages for $\{T_i\}$ and $\{B_i\}$ of pAug.TL and OctExt.TL. Figure 11 shows, for pAug.TL, the multiplier distributions P_j for $\{T_i\}$ and $\{B_i\}$ at different stages j. Plotted in Figure 11(a) is an asteroid curve $P(r) \sim \varepsilon^{-\alpha_e|r-1/2|}$, with $\alpha_e = 8$, where subscript e designates the double exponential profile of $P(r)$. Collapsed on it are five P_j curves with $j = 11, \ldots, 15$. Also shown is P_j with $j = 9$ as a diamond curve, which is observed to deviate from the asteroid curve. Note that if we fit it with a double exponential curve, then α_e has a value larger than 8. In fact, α_e increases monotonically with the decrease of the stage number j when $j \leqslant 11$. As

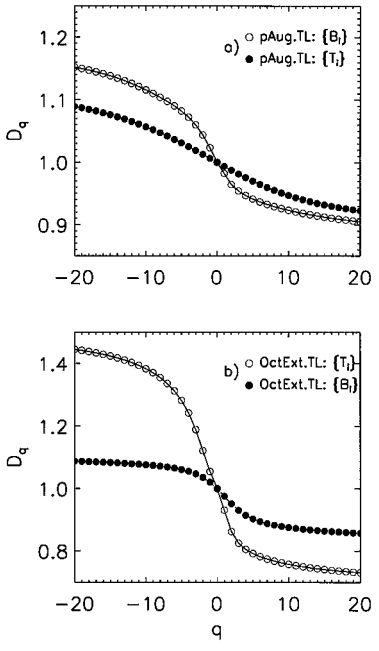

Figure 9. The generalized dimension spectrum for $\{T_i\}$ and $\{B_i\}$ of (a) pAug.TL, and (b) OctExt.TL.

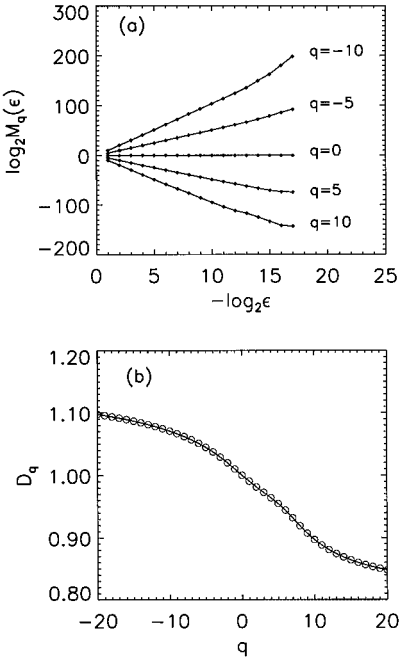

Figure 10. (a) $\log_2 M_q(\varepsilon)$ vs $-\log_2 \varepsilon$, and (b) the generalized dimension spectrum for the frame size sequence data MPEG.data.

will be discussed in the next section, taking this into proper consideration is of key importance to a successful modelling.

Next, we consider the $\{B_i\}$ process for pAug.TL. This is shown in Figure 11(b). The asteroid curve is generated from $P(r) \sim e^{-\alpha_g |r-1/2|^2}$, with $\alpha_g = 80$, where subscript g designates the Gaussian shape of $P(r)$. Collapsed on it are three P_j curves with $j = 9$, 10 and 11. Also shown is a diamond curve for P_j with $j = 8$. Again, if we fit a Gaussian-shaped curve to this P_j, then the value for α_g is larger than 80. Note that α_g also monotonically increases with the decrease of the stage number j when $j \leqslant 9$. Thus we conclude that for pAug.TL, both $\{B_i\}$ and $\{T_i\}$ processes are multifractals in a certain time scale range.

The results for OctExt.TL are shown in Figure 12, where asteroid curves are generated from $P(r) \sim e^{-\alpha_e |r-1/2|}$, with $\alpha_e = 7$ and $P(r) \sim e^{-\alpha_g |r-1/2|^2}$ with $\alpha_g = 270$ for $\{T_i\}$ (Figure 12(a)) and $\{B_i\}$ (Figure 12(b)), respectively. Collapsed on them are P_j curves with $j = 11$, 12 and 13 for $\{T_i\}$, and $j = 9$, 10 and 11 for $\{B_i\}$. Also plotted as diamond curves are P_j with $j = 8$ for $\{T_i\}$, and P_j with $j = 7$ for $\{B_i\}$. Note that if we fit a double exponential curve to P_j with $j = 8$ for $\{T_i\}$, then α_e is larger than 7. However, now α_e no longer increases monotonically with the decrease of the stage number j. We thus note a distinctive difference between the data collected for LAN and WAN traffic. This will be further discussed in the next section. For the $\{B_i\}$ sequence of OctExt.TL, if we fit a Gaussian-shaped curve to P_j with $j = 7$ for $\{B_i\}$, then α_g is much smaller than 270, and α_g decreases monotonically with the decrease of the stage number j when $j \leqslant 9$. Recall that for the

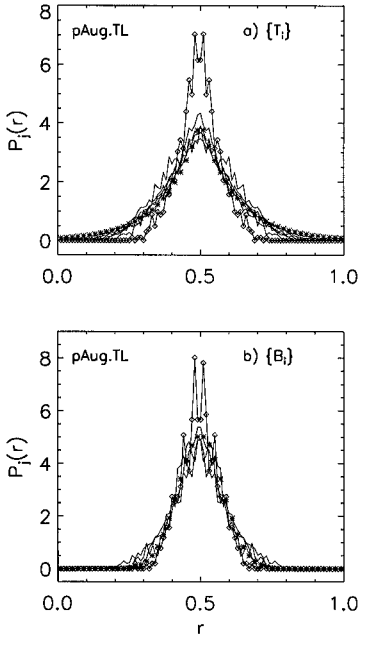

Figure 11. Multiplier distributions P_j for (a) $\{T_i\}$, and (b) $\{B_i\}$ of pAug.TL. See the text for more detail.

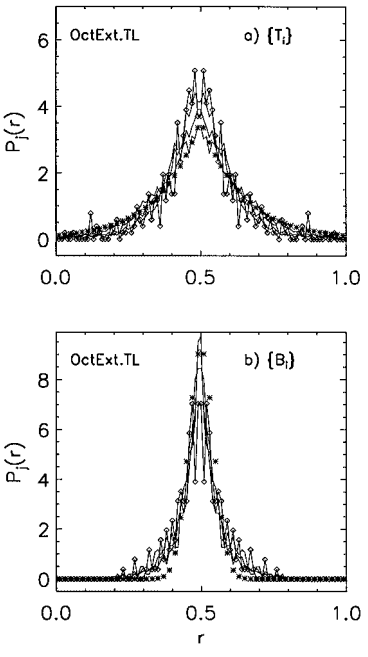

Figure 12. Multiplier distributions P_j for (a) $\{T_i\}$, and (b) $\{B_i\}$ of OctExt.TL. See the text for more detail.

$\{B_i\}$ process of pAug.TL, α_g increases monotonically with the decrease of the stage number j. The consequence of this difference will be further discussed in the next section.

The result for MPEG.data is shown in Figure 13, where the asteroid curve is generated from $P(r) \sim e^{-\alpha_g |r-1/2|^2}$, with $\alpha_g = 200$. Collapsed on it are P_j curves with $j = 8, \ldots, 11$. Again we see that the frame size sequences of the video traffic form a multifractal process over certain time scales.

We have observed an interesting feature that the weights at those particular stages where they form a multiplicative process follow approximately log-normal distribution. This is consistent with property (iv) of Section 3.3.

In summary, we have shown that the $\{B_i\}$ and $\{T_i\}$ sequences for the pAug.TL and OctExt.TL data sets, and the frame size sequences of MPEG.data, behave as multifractals over certain time scale ranges. The behavior of the process outside the latter time scale is not governed by a multiplicative multifractal process. Nevertheless, we can approximate its behaviour by still using a multifractal model, as illustrated in the next section.

5. $MF_e/MF_g/1$ QUEUEING SYSTEM

In this section, we consider a $MF_e/MF_g/1$ queueing system, where MF_e and MF_g denote multifractal interarrival time series and packet size sequences with appropriate parameters. We show that the system size tail distribution behaviour exhibited by the $MF_e/MF_g/1$ model closely

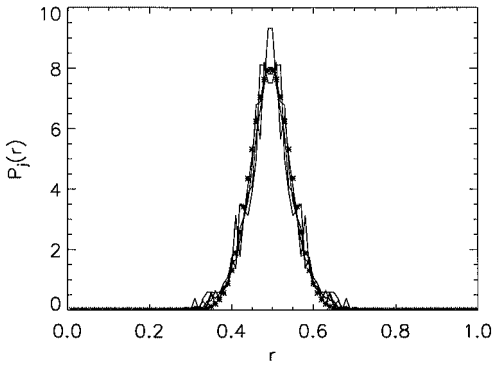

Figure 13. Multiplier distributions P_j for the frame size sequence data MPEG.data. See the text for more detail.

matches that obtained when the queueing system is driven by the actual traffic processes. Furthermore, this has been noted to be the case for all examined LAN, WAN, as well as VBR video traffic processes. Let us briefly recapitulate the procedure for constructing a multiplicative process.

Assume that over the observation period of the traffic process, there are 2^N arrivals, with mean interarrival time and mean packet length given by T and b, respectively. The observation time is thus equal to $\sum_{i=1}^{2^N} T_i = 2^N T$, while the total length of the packets is $\sum_{i=1}^{2^N} B_i = 2^N b$, where T_i and B_i denote the ith interarrival time and packet length, respectively. We view $2^N T$ and $2^N b$ as the weights at stage zero of the interarrival and packet-length multiplicative processes, respectively. Assume the multiplier distribution is already chosen (this will be discussed shortly). Then, we can follow the standard procedure described in Section 3.2, to construct the 2^N samples $\{T_i'\}$ and $\{B_i'\}$ at stage N for the two modelled processes.

Assume a single server queueing system using a FIFO service discipline and an infinite buffer, with interarrival times and packet sizes modelled by random multiplicative processes, as described above. Denote such a queueing system by $\mathrm{MF_e/MF_g/1}$, where subscripts e and g designate the multiplier distributions $P(r)$ to be double exponential ($P(r) \sim \mathrm{e}^{-\alpha_e|r-1/2|}$) and Gaussian ($P(r) \sim \mathrm{e}^{-\alpha_g|r-1/2|^2}$) characterized by parameters α_e and α_g, respectively. Thus our multifractal traffic model contains four parameters: α_e, α_g, the mean interarrival time T, and the mean packet length b. Our purpose is to find proper values for α_e and α_g to model the actual traffic trace data. For this purpose, we simulate a $\mathrm{MF_e/MF_g/1}$ queueing system and compare its behaviour (through observation of the system size tail distribution) to a queueing system driven by the actual traffic trace. We then select values for α_e and α_g so that the latter two queueing systems exhibit similar performance under several loading ratio conditions.

Here again, we use the first 2^{19} arrivals of the Ethernet traffic trace data. For modelling the video traffic, we use the whole sequence (of length 174136) of the frame size data. Since $2^{17} < 174136 < 2^{18}$, we generate a multiplicative process till stage 18 and retain the first 174136 weights at stage 18. We then rescale them so that their average is equal to the mean of the observed frame sizes.

Assuming $P(r) \sim \mathrm{e}^{-\alpha_e|r-1/2|}$ or $\sim \mathrm{e}^{-\alpha_g|r-1/2|^2}$, then $P(r) \to \delta(r-1/2)$ when α_e or $\alpha_g \to \infty$. The latter distribution leads to sequences of fixed variables, and thus represent non-bursty traffic

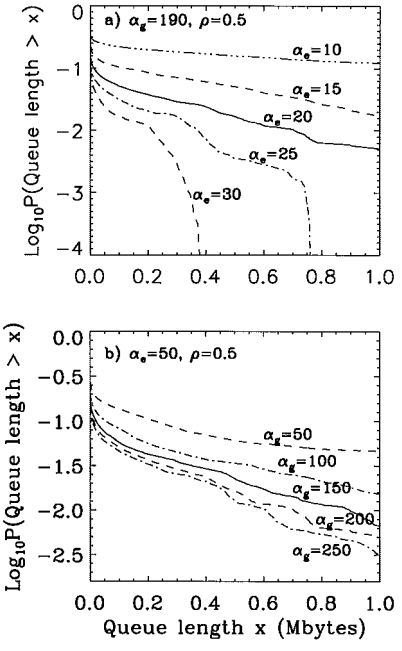

Figure 14. Complementary queue length distributions of $MF_e/MF_g/1$ queueing systems with (a) $\rho = 0.5$, $\alpha_g = 190$, and 5 different α_e's; and (b) $\rho = 0.5$, $\alpha_e = 50$, and 5 different α_g's. The mean interarrival time and mean packet length are the same as those for pAug.TL.

processes. This leads us to expect that given the mean interarrival time and the mean packet length, the burstiness of the traffic monotonically increases with the decrease of α_e (or α_g). Figure 14(a) shows the complementary queue length distributions when α_g is fixed to be 190 and the utilization is set equal to $\rho = 0.5$. The five illustrated curves, from top to bottom, correspond to $\alpha_e = 10$, 15, 20, 25 and 30. The results for a fixed value of $\alpha_e = 50$ and $\rho = 0.5$ is shown in Figure 14(b), where five curves, from top to bottom, correspond to $\alpha_g = 50$, 100, 150, 200 and 250, respectively. These results clearly confirm the above-mentioned property, illustrating the monotonic increase in the burstiness of the traffic processes resulting as we decrease the α parameters.

We note an interesting feature exhibited by Figure 14(b). The burstiness of the traffic decreases non-uniformly with the increase of α_g: it decreases quite fast when α_g is small, and very little when α_g is already large. We note that when α_g is already large, the burstiness contributed by α_g is much smaller than that contributed by α_e. Further increase of α_g will not result in a significant reduction in the burstiness of the traffic. When $\alpha_g \to \infty$, the burstiness of the traffic is solely determined by α_e. This observation is also valid when α_g is fixed and α_e is varied.

The observed burstiness behaviour of the $MF_e/MF_g/1$ system is consistent with the properties (iii), (v) and (vi) of Section 3.3. Larger values for α_e and/or α_g correspond to smaller second momentum, hence, smaller variation.

We have simulated the queue-size behaviour of the $MF_e/MF_g/1$ system by selecting the α_e and α_g values associated with the asteroid curves in Figures 11 and 12. We have found this model to yield a system-size tail distribution which is longer than that obtained when pAug.TL is used to drive the queueing system, while it is shorter than that obtained when OctExt.TL is used to drive

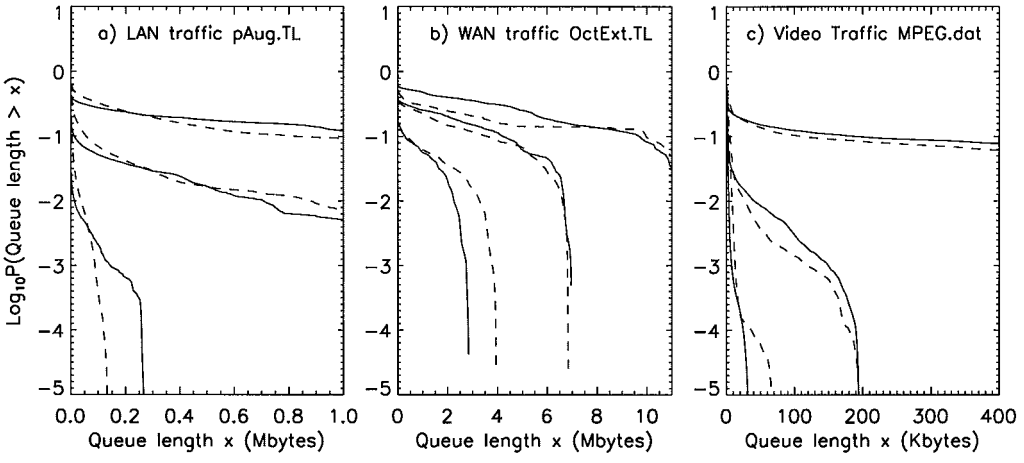

Figure 15. Comparison of complementary queue length distributions of single server FIFO queueing systems driven by measured data (a) pAug.TL, (b) OctExt.TL, and (c) MPEG.data (dashed lines), and of corresponding $MF_e/MF_g/1$ queueing systems (solid lines). The parameters used for the $MF_e/MF_g/1$ queueing systems are (a) $(\alpha_e, \alpha_g) = (20, 190)$, (b) $(\alpha_e, \alpha_g) = (6, 50)$, and (c) $(\alpha_e, \alpha_g) = (\infty, 150)$. Three curves, from top to bottom, correspond to $\rho = 0.7$, 0.5 and 0.3, respectively.

the queueing system. This is because the asteroid curves correspond to the estimated multiplier distributions P_j at quite large values of j, corresponding to relatively short time scales. As pointed out in the last section, for pAug.TL, the α_e and α_g values fitted from P_j at smaller values of j (corresponding to longer time scales) are larger than those associated with the asteroid curves in Figure 11, while for OctExt.TL, the α_e and α_g values fitted from P_j at smaller values of j are smaller than those associated with the asteroid curves in Figure 12. The results exhibited by Figure 14 indicate that the LAN traffic process (pAug.TL) is less bursty for longer time scales than for shorter time scales, while the WAN traffic process (OctExt.TL) is more bursty for longer time scales than for shorter time scales. By selecting the α_e and α_g values to be those associated with the asteroid curves in Figures 11 and 12, the modelled traffic processes exhibit the same degree of burstiness at all time scales. This is the underlying reason that the $MF_e/MF_g/1$ system yields a system-size performance behaviour which is not very close to that obtained when the measured traffic data, pAug.TL and OctExt.TL, are used to drive the queueing system.

Consequently, better values for α_e and α_g are selected by fitting the means of P_j over a multitude of stages. To obtain a good fit of the system size tail distribution of a queueing system driven by the corresponding measured traffic processes, we find (by trial and error) that for pAug.TL, $(\alpha_e, \alpha_g) = (20, 190)$, and for Oct.Ext.TL, $(\alpha_e, \alpha_g) = (6, 50)$. For the video frame size sequence data, the frame durations are fixed so that $\alpha_e = \infty$. We then find (also by trial and error) that $\alpha_g = 150$. Note that the value for α_g for the VBR video data is quite close to that associated with the asteroid curve in Figure 13 ($\alpha_g = 180$).

Figures 15(a)–(c) show a comparison of the complementary queue length distributions for a queueing system loaded by the measured traffic, pAug.TL, OctExt.TL, and MPEG. data, and the corresponding $MF_e/MF_g/1$ queueing system models based on the above selected α parameters. The solid curves display the results obtained from $MF_e/MF_g/1$ simulation, and the

dashed curves exhibit the results obtained from a queueing system driven by the measured traffic trace. The three curves in each figure, from top to bottom, correspond to three different utilization levels, $\rho = 0.7$, 0.5 and 0.3. Using the α_e and α_g parameter values mentioned above, the $MF_e/MF_g/1$ model proves to yield excellent fit of the complementary queue size distributions of a queueing system loaded by the measured LAN, WAN, and VBR video traffic.

6. COMPARISON BETWEEN THE MULTIFRACTAL TRAFFIC MODEL AND THE FBm MODEL FOR THE OBSERVED TRAFFIC PROCESSES

In this section, we make a comparison between the $MF_e/MF_g/1$ queueing system model and the FBm model, as it relates to fitting the queue size tail distribution for a single server queueing system loaded by the traffic processes mentioned above.

FBm is the most intensively studied model exhibiting the long-range dependence property in traffic engineering. After Norros [12, 13] introduced FBm as a traffic model, Erramilli *et al.* [22] have checked the complementary queue length distribution formula of Norros [12,13], and found excellent agreement with simulations for a single server queueing system operated at utilization $\rho = 0.5$. By breaking their Ethernet traffic into source–destination pairs, Willinger *et al.* [7] have shown that the long-range-dependent ON/OFF model is consistent with the data, and further proven that the ON/OFF model asymptotically approaches a FBm model. It has also been shown that analytic results similar to those exhibited by a single server queueing system driven by a FBm traffic process can be obtained for a queue fed by a traffic process modelled by LRD ON/OFF sources [23] and by a FARIMA model [14]. These works demonstrate the effectiveness of the FBm model. In the following, we demonstrate, by considering measured LAN, WAN, and VBR video traffic data, that a multiplicative multifractal traffic model can provide an even better fit to the underlying complementary queue length distribution.

For this purpose, for the FBm model, we use the Norros' lower bound complementary queue length distribution formula [12,13]:

$$P(V > x) \sim \exp[-\alpha x^{2-2H}] \tag{11}$$

with

$$\alpha = \frac{m^{2H-1}((1-\rho)/\rho)^{2H}}{2a}\left[\left(\frac{1-H}{H}\right)^{H} + \left(\frac{H}{1-H}\right)^{1-H}\right]^2 \tag{12}$$

where the random variable V represents the (steady-state) unfinished load, H is the Hurst parameter, $m > 0$ is the mean input rate, ρ is the utilization level, and a is a variance coefficient. Using the measured traffic trace data, we first estimate the parameters a [24] and H [1,3]; then we use Equation (11) to compute the complementary queue length distributions. We compare the computed queue length distributions with those obtained for a queueing system driven by the measured traffic trace data. Figure 16(a) shows the complementary queue length distributions for pAug.TL for three different utilization levels, $\rho = 0.7$, 0.5 and 0.3 (from top to bottom). The dashed lines are from a queueing simulation with pAug.TL as input traffic. The solid lines are computed from Equation (11). Note that if we plot the middle curves ($\rho = 0.5$) for $x \leqslant 100$ kbytes,

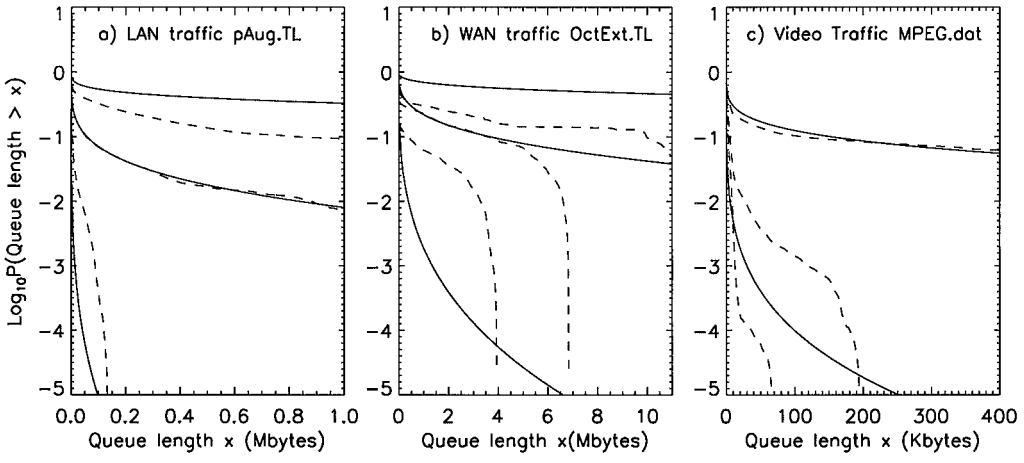

Figure 16. Comparison of complementary queue length distributions of single server FIFO queueing systems driven by measured data (a) pAug.TL, (b) OctExt.TL, and (c) MPEG.data (dashed lines), and corresponding FBm traffic processes (solid lines). Three curves, from too to bottom, correspond to $\rho = 0.7$, 0.5 and 0.3, respectively. Note for MPEG.data, the solid curve for $\rho = 0.3$ is too close to the y-axis to be seen.

then the solid line and the dashed line almost coincide, which is consistent with the result presented by Erramilli *et al.* [22]. A salient feature exhibited by Figure 16(a) is that in terms of the complementary queue length distribution, given a certain utilization level (say, for example $\rho = 0.5$), by choosing the H value carefully, the FBm model can provide excellent fit. However, for lighter and heavier loading conditions (smaller and larger ρ values), when the same H value is used, the model does not provide a good fit.

As is evident from Figure 16(c), this problem is also associated with the modelling of VBR video traffic data.

The aforementioned problem can become even more pronounced when FBm is used to model WAN traffic. This is shown in Figure 16(b), where the solid curves are generated from Equation (11), and the dashed lines are obtained from a queueing simulation with OctExt.TL as input traffic. The exhibited three curves, from top to bottom, correspond to $\rho = 0.7$, 0.5 and 0.3.

In ATM networks, packet loss probability levels as small as 10^{-12} are of interest. Such low probabilities lead to the almost vertically dropping segments of Figures 15 and 16. This feature is not predicted by Equation (11).

By comparing Figures 15 and 16, it is clear that the $MF_e/MF_g/1$ model has overcome both problems. Hence, for the traffic processes studied here, it yields more accurate results than the FBm model in predicting the behaviour of the complementary queue length distribution of the system (or equivalently, the packet loss probability).

7. CONCLUSIONS

We introduce a multiplicative multifractal process for the modelling of long-range-dependent network traffic processes. We prove a number of properties of such a process that are most

relevant to its use as network traffic models. Through analysis of Bellcore's LAN, WAN and VBR video traffic trace data, we show this process to provide excellent fit with measured traffic data. The model employs two multifractals to characterize the self-similar interarrival time series and the packet length sequence of the measured traffic trace data, respectively. The multiplicative model involves one to two basic parameters. Once the parameters are chosen, it is easy to construct. To calibrate the model, we have considered a single server queueing system which is loaded, on one hand, by the measured processes, and, on the other hand, by our multiplicative multifractal processes. In comparing the performance of both systems, we have demonstrated our model to effectively track the behaviour exhibited by the system driven by the actual traffic processes. We have also shown that, for the measured traffic processes studied here, our multiplicative multifractal traffic model provides more accurate results concerning the behaviour of the packet loss probability than those obtained using a FBm model.

ACKNOWLEDGEMENTS

Our thanks to the Bellcore researchers (Drs Leland and Garret) for making available their Ethernet traffic trace data and the VBR video data. Without these data, this work would not be possible. J.B. Gao would also like to thank his friend Johnny Lin for teaching him LATEX. This work is supported by UC MICRO/SBC Pacific Bell research grant 98-131 and by ARO grant DAAGIJ-98-1-0338.

REFERENCES

1. Leland WE, Taqqu MS, Willinger W, Wilson DV. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking* 1994; **2**:1–15.
2. Paxson V, Floyd S. Wide area traffic—the failure of Poisson modeling. *IEEE/ACM Transactions on Networking* 1995; **3**:226–244.
3. Beran J, Sherman R, Taqqu MS, Willinger W. Long-range-dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications* 1995; **43**:1566–1579.
4. Garret MW, Willinger W. Analysis, modeling and generation of self-similar VBR video traffic. *Proceedings of ACM SIGCOMM*, London, England, 1994.
5. Crovella ME, Bestavros A. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking* 1997; **5**:835–846.
6. Erramilli A, Singh PR, Pruthi P. An application of deterministic chaotic maps to model packet traffic. *Queueing Systems* 1995; **20**:171–206.
7. Willinger W, Taqqu MS, Sherman MS, Wilson DV. Self-similarity through high-variability: Statistical analysis of ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking* 1997; **5**:71–86.
8. Cox DR. Long-range dependence: a review. In *Statistics*: *An Appraisal*, David HA, Davis HT (eds). The Iowa State University Press: Ames, IA, 1984; 55–74.
9. Likhanov N, Tsybakov B, Georganas ND. Analysis of an ATM buffer with self-similar ('fractal') input traffic. *Proceedings of IEEE InfoCom*, Boston, MA, 1995.
10. Parulekar M, Markowski AM. Buffer overflow probabilities for a multiplexer with self-similar input. *Proceedings of IEEE InfoCom*, San Francisco, CA, 1996.
11. Tsybakov B, Georganas ND. On self-similar traffic in ATM queues: definitions, overflow probability bound, and cell delay distribution. *IEEE/ACM Transactions on Networking* 1997; **5**:397–409.
12. Norros I. A storage model with self-similar input. *Queueing Systems* 1994; **16**:387–396.
13. Norros I. On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE JSAC* 1995; **13**:953–962.
14. Su YT. Personal communication, 1998.
15. Ryu B, Loven SB. Point process approaches for modeling and analysis of self-similar traffic: Part II—applications. *Proceedings of the Fifth International Conference on Telecommunications Systems, Modeling, and Analysis*, Nashville, TN, March 1997.
16. Addie RG, Zukerman M, Neame T. Fractal traffic: measurements, modeling and performance evaluation. *Proceedings of IEEE InfoCom*, Boston, MA, 1995.

17. Robert S, Boudec J-Y. New models for pseudo self-similar traffic. *Performance Evaluation* 1997; **30**:57–68.
18. Taqqu MS, Teverovsky V, Willinger W. Is network traffic self-similar or multifractal? *Fractals* 1997; **5**:63–73.
19. Frisch U. *Turbulence—The legacy of A. N. Kolmogorov*. Cambridge University Press: Cambridge, 1995.
20. Hentschel HGE, Procaccia I. The infinite number of generalized dimensions of fractals and strange attractors. *Physica D* 1983; **8**:435–444.
21. Gouyet JF. *Physics and Fractal Structures*. Springer: New York, 1995.
22. Erramilli A, Narayan O, Willinger W. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking* 1996; **4**:209–223.
23. Brichet F, Roberts J, Simonian A, Veitch D. Heavy traffic analysis of a storage model with long range dependent on/off sources. *Queueing Systems* 1996; **23**:197–215.
24. Krishnan KR. A new class of performance results for a fractional Brownian traffic model. *Queueing Systems* 1996; **22**:277–285.

AUTHORS' BIOGRAPHIES

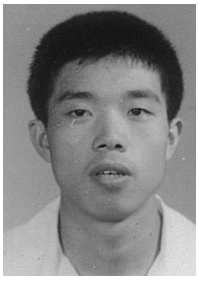

Jianbo Gao received his BSc in Electrical Engineering from Zhejang University, China, his MSc in Fluid Mechanics from the Institute of Mechanics, Chinese Academy of Sciences, and his PhD degree in Electrical Engineering from UCLA. He is now a research engineer in the EE Dept at UCLA. He has been actively working on network traffic modelling and network performance evaluation in recent years. Dr Gao has very broad interests. He is an internationally recognized expert in the study of chaos. Recently, he has also been working on bioinformatics using information theory, chaos theory, and fractal geometry.

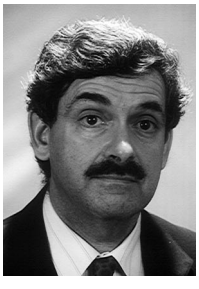

Izhak Rubin received the BSc and MSc from the Technion—Israel Institute of Technology, Haifa, Israel, and the PhD degree from Princeton University, Princeton, NJ, all in Electrical Engineering. Since 1970, he has been a faculty of the UCLA School of Engineering and Applied Science where he is currently a Professor in the Electrical Engineering Department. Dr Rubin has had extensive research, publications, consulting, and industrial experience in the design and analysis of commercial and military computer communications and telecommunications systems and networks. At UCLA, he is leading a large research group. He also serves as President of IRI Computer Communications Corporation, a leading team of computer communications and telecommunications experts engaged in software development and consulting services. Professor Rubin served as co-chairman of the 1981 IEEE International Symposium on Information Theory; as programme chairman of the 1984 NSF-UCLA workshop on Personal Communications; as program chairman for the 1987 IEEE INFOCOM conference; and as program co-chair of the IEEE 1993 workshop on Local and Metropolitan Area networks. Dr Rubin has been elected as a Fellow of the IEEE for his contributions to the analysis and design of computer communications networks. He has served as editor of the *IEEE Transactions on Communications*, on the ACM/Baltzer journal *Wireless Networks*, on the Kluwer journal *Photonic Network Communications* and on the Wiley journal *International Journal of Communication Systems*.

