On the Relevance of Long-Range Dependence in Network Traffic *

Matthias Grossglauser Jean-Chrysostome Bolot

INRIA
BP 93
06902 Sophia Antipolis Cedex
France
{Matthias.Grossglauser, Jean.Bolot}@inria.fr

Abstract

There is mounting experimental evidence that network traffic processes exhibit ubiquitous properties of self-similarity and long range dependence (LRD), i.e. of correlations over a wide range of time scales. However, there is still considerable debate about how to model such processes and about their impact on network and application performance. In this paper, we argue that much recent modeling work has failed to consider the impact of two important parameters, namely the finite range of time scales of interest in performance evaluation and prediction problems, and the first-order statistics such as the marginal distribution of the process.

We introduce and evaluate a model in which these parameters can be easily controlled. Specifically, our model is a modulated fluid traffic model in which the correlation function of the fluid rate is asymptotically second-order self-similar with given Hurst parameter, then drops to zero at a cutoff time lag. We develop a very efficient numerical procedure to evaluate the performance of the single server queue fed with the above fluid input process. We use this procedure to examine the fluid loss rate for a wide range of marginal distributions, Hurst parameters, cutoff lags, and buffer sizes.

Our main results are as follows. First, we find that the amount of correlation that needs to be taken into account for performance evaluation depends not only on the correlation structure of the source traffic, but also on time scales specific to the system under study. For example, the time scale associated to a queueing system is a function of the maximum buffer size. Thus for finite buffer queues, we find that the impact on loss of the correlation in the arrival process becomes nil beyond a time scale we refer to as the correlation horizon. Second, we find that loss depends in a crucial way on the marginal distribution of the fluid rate process. Third, our results suggest that reducing loss by buffering is hard. We advocate the use of source traffic control and statistical multiplexing instead.

*This work was supported in part by a grant from France Telecom/CNET.

1 Introduction

Experimental data obtained from the observation of systems is typically considered for modeling purposes as a realization, or sample path, of an underlying stochastic process. In practice, statistical analysis of the data proceeds with the additional hypotheses that the process is stationary and ergodic. Such analysis has shown that many systems of interest in the physical world exhibit a property of correlation over many different time scales, often referred to as long range dependence (LRD) or long memory. Some of the better known examples of such systems are found in hydrology [16]. However, the phenomenon of LRD occurs in many other systems including chemical, astronomical, and biological systems (see [3] for references).

In spite of much statistical evidence, the existence of LRD has often been met with resistance or at least puzzlement. This was caused in large part by the absence of physical explanations for the observed phenomenon. Hydrologists for example wondered "By what sort of physical mechanism can the influence of, say, the mean temperature of this year at a particular geographic location be transmitted over decades and centuries?" [19]. Two approaches then are possible. One approach is to argue that the LRD observed in the measurement data is a consequence of inadequate hypotheses, in particular the stationarity hypothesis, made about the underlying process that (it is assumed) did generate the data. For example, the superposition of a process with short range dependence (SRD) and an appropriately chosen on/off trend [19] or a hyperbolically decreasing trend [5] is difficult to distinguish from a stationary process with LRD. Another approach is to not worry about a physical explanation and to develop and use models that do exhibit LRD, on the grounds that a model is good not because it explains a phenomenon correctly, but rather because it provides good prediction ability and it is numerically and/or analytically tractable.

Unfortunately, it is not possible to tell with certainty whether or not a realization is stationary. Therefore, the jury is still out on which of the above two approaches is "the right one". Clearly, it is better for a model to match more properties of the data. However, a model is a tool for decision making. Thus, its quality depends on the quality of the decisions it leads to rather than on its closeness to physical reality.

The situation in the area of communication systems in general, and computer networks in particular, is no exception to that described above. Careful statistical analysis of data collected over a wide variety of networks has provided ample evidence that network traffic processes exhibit properties of self-similarity and LRD [23, 21, 9, 20, 27, 4]. However, there is still considerable debate about how to model such processes. Different approaches have been taken that parallel those taken in other areas and described earlier. One approach has been to argue that the observed LRD may be due to non-stationarity in the data caused by the superposition of level shifts [8] or Dirac pulses [13] with short range dependent (SRD) stationary processes. Another approach has been to use stochastic models (such as fractional Brownian motion [24], zero-rate renewal processes [30] and various other point processes [28]) or deterministic models (such as chaotic maps [12]) that exhibit the LRD observed in the experimental data. However, these models are analytically difficult to handle. Furthermore, they do not provide much insight into why they are meaningful on physical grounds. This explains in part that much modeling work still relies on more traditional multi-state Markovian models (e.g. [22, 2]).

However, recent work has shown that the superposition of many on/off sources with heavy-tailed on- and off-periods results in aggregate traffic with LRD [31, 6]. Furthermore, there is widespread evidence that human as well as computer sources of traffic do tend to behave as heavy tailed on/off sources [7, 20, 31]. Thus, LRD in network traffic can be explained simply in terms of the nature of the traffic generated by individual sources.

This intuitively appealing explanation suggests that LRD will remain a salient feature of network traffic even as network characteristics such as bandwidth and topology evolve over time. Thus, the fundamental question for both current and future networks is that of the practical impact of LRD on network and application performance. Not surprisingly, much effort has focused on trying to answer this question. The main result is that the performance of queueing systems with LRD in the input or service processes can be radically different from the performance of usual Markovian systems [11, 6].

However, we believe that not enough work has considered the impact of parameters other than the correlation structure of the traffic. Consider for example the asymptotic behavior of an infinite queue fed with three different arrival processes that all exhibit the LRD property. First, if the arrival process is a fractional Brownian motion, then the queue length distribution is Weibullian. Second, if the arrival process is a single on/off source with heavy-tailed on and off periods, then the queue length distribution is hyperbolic. Third, if the arrival process is a single on/off source in which the off periods only are heavy-tailed, then the queue length distribution decays exponentially [6, 26]. Thus, processes with the same correlation structure can generate vastly different queueing behavior. Therefore, it is important to consider parameters other than the correlation of the input process for accurate performance prediction.

Two such parameters stand out, namely the marginal distribution of the arrival process, and the finite range of time scales of interest in performance evaluation and prediction problems. The main goal of this paper is to evaluate the impact of these parameters, as well as the correlation structure of traffic sources, on network and application performance.

To achieve this goal, we develop a model in which all three parameters can be easily controlled. Specifically, our model is a modulated fluid traffic model in which the correlation function of the fluid rate is asymptotically second-order self-similar with given Hurst parameter, then drops to zero at a cutoff time lag. We then consider the behavior of a finite-buffer queue fed with the above fluid input process. We cannot describe this behavior with closed-form analytic expressions. However, we develop a very efficient numerical procedure to evaluate various performance measures. In this paper, the measure of interest is the fluid loss rate, i.e. the amount of work lost because of buffer overflow to the amount of work arriving at the queue.

Our main results are as follows. First, we find that the amount of correlation that needs to be taken into account for performance evaluation depends not only on the correlation structure of the source traffic, but also on time scales specific to the system under study. For example, the time scale associated to a queueing system is a function of the maximum buffer size. Thus for finite buffer queues, we find that the impact on performance of the correlation in the arrival process becomes nil beyond a time scale we refer to as the correlation horizon.

Second, we find that the loss rate depends in a crucial way on the buffer size, but more importantly on the marginal distribution. An obvious consequence is that the marginal distribution must be taken into account for accurate loss prediction. Another consequence is that controlling the loss rate by increasing the buffer size is inefficient. Instead, statistical multiplexing and source traffic control are two much more efficient ways to achieve high utilization while keeping loss low. The rest of the paper is organized as follows. In Section 2, we describe the model and the numerical solution procedure. In Section 3, we describe the behavior of the loss rate as as function of system parameters. In Section 4, we discuss the implications of our results. Section 5 concludes the paper.

2 Model Description

In this section, we describe our modulated fluid traffic model and the numerical procedure we developed to evaluate the behavior of a finite buffer queue fed with this input traffic.

Recall that the goal of the model is to examine the impact on the performance measure of interest of parameters such as time scales and the marginal distribution of the traffic process. Thus, we need a traffic model in which these parameters can be controlled easily.

Specifically, the source traffic model is described by a random process $\{X_t\}$ which represents the fluid rate at time t. We assume that X_t takes on a finite set of possible rates $\{\lambda_1, \ldots, \lambda_M\}$. Furthermore, we assume that the fluid rate process is piecewise constant. Thus, the rate remains constant over intervals the lengths of which are determined by arrivals of a stationary point process $\{\tau_n\}$. We denote $X_t = \lambda(n)$ for $\tau_n \leq t < \tau_{n+1}$. The interarrival times $T_n = \tau_{n+1} - \tau_n$ are i.i.d. with survival function $\Pr\{T_n > t\} = F_T(t)$. Furthermore, the constant fluid rate $\lambda(n)$ is i.i.d. with distribution $\Pr\{\lambda(n) = \lambda_i\} = \pi_i$. For i.i.d. random variables, we drop the subscript if this does not lead to confusion. Note that this model can be specialized into the familiar on/off source model with identically distributed on and off periods.

The rest of this section proceeds in three steps. In the first step, we derive the covariance function of the fluid process

$\{X_t\}$ in terms of the interarrival time distribution $F_T(t)$, the rate matrix $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_M)$, and the marginal distribution of the fluid rate $\Pi = (\pi_1, \ldots, \pi_M)$. In the second step, we derive the occupancy distribution at time τ_n of a queue fed with $\{X_t\}$. In the third step, we use this distribution to derive the performance measure of interest here, namely the stationary fluid loss rate.

The covariance function of $X(t)$ is defined by

$$\phi(t) = \mathrm{E}[X_0 X_t] - (\mathrm{E}[X_0])^2 \tag{1}$$

where

$$\begin{aligned} \mathrm{E}[X_0 X_t] &= \sum_{i=1}^{M}\sum_{j=1}^{M} \lambda_i \lambda_j \Pr\{X_0 = \lambda_i, X_t = \lambda_j\} \\ &= \sum_{i=1}^{M} \lambda_i^2 \Pr\{X_0 = \lambda_i, X_t = \lambda_i\} \\ &\quad + 2\sum_{i=1}^{M-1}\sum_{j=i+1}^{M} \lambda_i \lambda_j \Pr\{X_0 = \lambda_i, X_t = \lambda_j\}. \end{aligned} \tag{2}$$

The first term on the right-hand side of (2) corresponds to the case when the source rates at time 0 and at time t are identical. To compute this term, we have to account for the case when 0 and t lie in different intervals and the case when they lie in the same interval. Let $p(t)$ denote the probability that 0 and t lie in the same interval, or equivalently, that there is no arrival in $[0, t]$. Then,

$$\sum_{i=1}^{M} \lambda_i^2 \Pr\{X_0 = \lambda_i, X_t = \lambda_i\} = \\ = [1 - p(t)] \sum_{i=1}^{M} \pi_i^2 \lambda_i^2 + p(t) \sum_{i=1}^{M} \pi_i \lambda_i^2.$$

The second term of (2) corresponds to the case when the source rates at time 0 and at time t are different. Thus, there must have been a rate change (i.e. an arrival) in this interval, which implies that X_t and X_0 are independent. Therefore, $\Pr\{X_0 = \lambda_i, X_t = \lambda_j\} = \pi_i \pi_j$. Also,

$$(\mathrm{E}[X_0])^2 = (\sum_{i=1}^{M} \pi_i \lambda_i)^2 = \sum_{i=1}^{M} \pi_i^2 \lambda_i^2 + 2\sum_{i=1}^{M-1}\sum_{j=i+1}^{M} \pi_i \pi_j \lambda_i \lambda_j. \tag{3}$$

Combining the above results yields

$$\begin{aligned} \phi(t) &= p(t)\sum_{i=1}^{M} \pi_i \lambda_i^2 + [1 - p(t)] \sum_{i=1}^{M} \pi_i^2 \lambda_i^2 - \sum_{i=1}^{M} \pi_i^2 \lambda_i^2 \\ &= p(t)\Pi\Lambda^2 \mathbf{1}^T + [1 - p(t)]\Pi\Lambda^2\Pi^T - \Pi\Lambda^2\Pi^T \\ &= p(t)\Pi\Lambda^2(\mathbf{1} - \Pi)^T. \end{aligned} \tag{4}$$

Furthermore, it follows from renewal theory that the probability $p(t)$ is equal to the probability that the residual life τ_{res} of the interarrival time T_n exceeds t, which is given by [18, p. 172]

$$\Pr\{\tau_{res} \geq t\} = \int_t^\infty \frac{F_T(x)}{E[T_n]} dx. \tag{5}$$

We now consider the special case when the interarrival time distribution $F_T(t)$ is a truncated Pareto distribution defined by

$$F_T(t) = \begin{cases} \left(\frac{t+\theta}{\theta}\right)^{-\alpha} & \text{if } t < T_c \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

where $1 < \alpha < 2$. We refer to this distribution as truncated Pareto because $F_T(t)$ is a Pareto distribution with parameters θ and α for $0 < t < T_c$. We refer to the parameter T_c as the *cutoff lag*.

Since the length T_n of an interval cannot exceed T_c, and since the rates in consecutive intervals are independent, it follows that there is no correlation in the fluid rate process beyond lag T_c.

We now compute the covariance function $\phi(t)$ for the above distribution. We have

$$\Pr\{\tau_{res} \geq t\} = \begin{cases} \frac{(t+\theta)^{-\alpha+1} - (T_c+\theta)^{-\alpha+1}}{\theta^{-\alpha+1} - (T_c+\theta)^{-\alpha+1}} & \text{if } t < T_c \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

Therefore

$$\phi(t) = \Pi\Lambda^2(\mathbf{1}-\Pi)^T \times \begin{cases} \frac{(t+\theta)^{-\alpha+1} - (T_c+\theta)^{-\alpha+1}}{\theta^{-\alpha+1} - (T_c+\theta)^{-\alpha+1}} & \text{if } t < T_c \\ 0 & \text{otherwise} \end{cases}. \tag{8}$$

We observe that $\phi(t)$ behaves asymptotically as $t^{-\alpha+1}$ when $T_c \to \infty$. This implies that $\{X_t\}$ is asymptotically second-order self-similar with Hurst parameter H such that $-\alpha + 1 = -(2 - 2H)$, i.e. $H = \frac{3-\alpha}{2}$ [21]. When T_c is finite, the model behaves like an asymptotic second-order self-similar model for $t < T_c$, but the correlation drops to zero at lag T_c.

In summary, our source model allows us to control the marginal fluid distribution Π, the time-scale T_c over which the correlation structure matches that of an asymptotic second-order self-similar model, and the Hurst parameter H.

We next examine the performance of a queue with constant service rate c and a finite buffer B fed with the above source model. This queueing model does not appear to be analytically tractable. However, we have developed a very efficient numerical procedure to determine the queue occupancy at the arrival instants τ_n. Although this queue occupancy is not equal to the queue occupancy at a random point in time, it is sufficient to derive our chosen performance metric, namely the long-term loss rate.

Let Q_n be the continuous random variable describing the queue occupancy at arrival instant τ_n, and let $W_n = T_n \times (\lambda(n) - c)$. W_n represents the difference between arriving and departing work in interarrival interval n. Note that the $\{W_n\}$ are i.i.d. because T_n and $\lambda(n)$ are i.i.d. and jointly independent. The queue occupancy is recursively given by

$$Q_{n+1} = \max(0, \min(B, Q_n + W_n)). \tag{9}$$

(X_n, Q_n) is a two-dimensional Markov process because X_n and X_{n-1} are independent. Assuming ergodicity, (X_n, Q_n) converges to (X_∞, Q_∞) as $n \to \infty$. Our goal is to derive upper and lower bounds of the stationary density $f_Q(x)$ of the queue occupancy. For convenience of notation, we assume that $\lambda_1 < \lambda_2 < \ldots < \lambda_L < c < \lambda_{L+1} < \lambda_{L+2} < \ldots < \lambda_M$. We exclude the trivial case when one of the fluid rates is equal to c, in which case the queue occupancy does not change.

The probability density of W_n is given by

$$f_W(w) = \sum_{i=1}^{M} \pi_i dF_T\left(\frac{w}{\lambda_i - c}\right) \tag{10}$$

From (9), we can compute the density of Q_{n+1} given the densities of Q_n and W. Let $U_n = Q_n + W_n$. Then

$$f_{U_n}(x) = f_{Q_n}(x) * f_W(x) \tag{11}$$

where $*$ denotes convolution. Then

$$f_{Q_{n+1}}(x) = \begin{cases} f_{U_n}(x) + [\int_{-\infty}^{0} f_{U_n}(y)dy]\delta(x) + & \\ +[\int_{B}^{\infty} f_{U_n}(y)dy]\delta(x-B) & \text{if } x \in [0,B] \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

where $\delta(.)$ is a Dirac impulse.

We know that $f_{Q_{n+1}}(x) = f_{Q_n}(x)$ in steady state. Furthermore, Q_n converges to this distribution as $n \to \infty$. Our numerical solution relies on computing an approximation of the stationary queue occupancy distribution by discretizing (12).

We define two discrete random variables $Q_{l,n}$ and $Q_{h,n}$ in the following way. Let $d = B/N$ be the quantization interval size, where B is the buffer size and N a positive integer. Let

$$\begin{aligned} Q_{l,n} &= d\lfloor Q_n/d \rfloor \\ Q_{h,n} &= d\lceil Q_n/d \rceil \end{aligned} \tag{13}$$

such that $\Pr\{Q_{l,n} \geq x\} \leq \Pr\{Q_n \geq x\} \leq \Pr\{Q_{h,n} \geq x\}$. We denote their densities by

$$\begin{aligned} q_{l,n}(i) &= \Pr\{Q_{l,n} = id\}, i = 0, \ldots, N \\ q_{h,n}(i) &= \Pr\{Q_{h,n} = id\}, i = 0, \ldots, N. \end{aligned} \tag{14}$$

The random variable U_n is quantized in the same way. Furthermore, we define the following vector

$$w(i) = \Pr\{W \in [id, (i+1)d)\} \tag{15}$$

Then (11) can be written as a discrete convolution for both the upper and the lower bound

$$\begin{aligned} u_{l,n}(i) &= \sum_{j=0}^{N} q_{l,n}(j) w(i-j) \\ u_{h,n}(i) &= \sum_{j=0}^{N} q_{h,n}(j) w(i-j-1). \end{aligned} \tag{16}$$

Deriving the discrete version of (12) is then straightforward.

A simple sample path argument shows that the queue occupancy at τ_n for an initially full queue is larger than or equal to the queue occupancy at τ_n for an initially empty queue. Therefore, initializing $q_{l,0}$ and $q_{h,0}$ to an empty and full queue, respectively, and iterating (16) will always yield a lower and an upper bound on the limiting queue occupancy. We have found that in this approach, the upper and lower bounds converge rapidly to the limiting occupancy (a graphical illustration of this is shown below in Figure 1). Nevertheless, it is possible to improve the efficiency of the procedure by using a fast Fourier transform (FFT) with appropriate zero-padding [25].

We can now derive a numerical procedure to compute the long-term stationary loss rate, defined as the ratio of work lost to work arriving, from the stationary queue occupancy at arrival instants. This is done as follows. The random variable $W_l = (W - (B - Q))^+$ represents the amount of work lost in an interarrival interval. The loss rate l is then

$$l = \frac{\mathrm{E}[W_l]}{\overline{\lambda}\mathrm{E}[T]} \tag{17}$$

with

$$\mathrm{E}[W_l] = \int_0^B f_Q(x)\mathrm{E}[(W - (B-Q))^+ | Q = x]dx \tag{18}$$

and

$$\begin{aligned} \mathrm{E}[W_l | Q = x] &= \int_0^\infty F_{W_l|Q=x}(y)dy = \\ &= \int_0^\infty \sum_{\{i: T_c(\lambda_i - c) - B + x > 0\}} \pi_i F_T\left(\frac{y + B - x}{\lambda_i - c}\right) dy \\ &= \frac{\theta}{\alpha - 1} \sum_{\{i: T_c(\lambda_i - c) - B + x > 0\}} \pi_i(\lambda_i - c) \times \\ &\times \left[\left(\frac{B - x}{\theta(\lambda_i - c)} + 1\right)^{1-\alpha} - \left(\frac{T_c}{\theta} + 1\right)^{1-\alpha}\right] \end{aligned}$$

We can compute an upper and lower bound on l using the upper and lower bounds on the queue occupancy derived above. We obtain

$$\frac{\sum_{i=0}^{N} q_l(i)\mathrm{E}[W_l|Q = id]}{\overline{\lambda}\mathrm{E}[T]} \leq l \leq \frac{\sum_{i=0}^{N} q_h(i)\mathrm{E}[W_l|Q = id]}{\overline{\lambda}\mathrm{E}[T]} \tag{19}$$

We have found that the above iterative procedure is quite efficient in terms of computational complexity and that it yields tight lower and upper bounds with a reasonable number of iterations. This is illustrated in Figure 1 which shows the upper and lower bounds for the queue occupancy after 5, 10, and 30 iterations in a typical case when the queue has been discretized using 100 bins (i.e., $d = B/100$).

The efficiency can be explained in part by the fact that in each iteration of (16), we "jump" from one arrival epoch to the next one without having to consider the details of the queueing behavior during the entire interval.

3 Numerical Results

In this section, we present results of numerical experiments we have conducted using the model and the analysis technique described above. The goal of these experiments is to evaluate the impact on performance of various parameters of the model.

The performance metric we consider is the loss rate, i.e. the ratio of work lost due to buffer overflow to total work arriving at the server. Other performance metrics would be of interest as well. However, the loss rate is a very natural metric for finite buffer queues. Furthermore, it relates to the tail distribution of the queue occupancy, which is the metric considered in most of the analyses pertaining to infinite

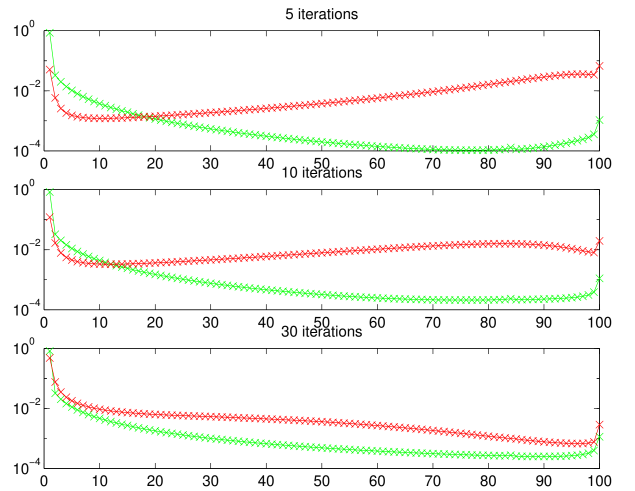

Figure 1: The upper and lower bounds on the arrival instant queue occupancy after 5, 10, and 30 iterations.

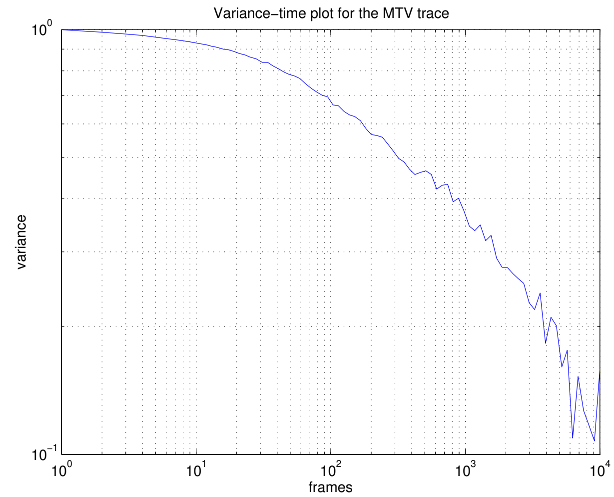

Figure 2: The variance-time plot of the MTV trace.

buffer queues[1]. We also note that the loss rate has tremendous impact on both application performance (e.g. on the quality of the audio or video delivered from a source to a destination) and network performance (e.g. on the number of retransmissions required to achieve reliable communication).

The parameters of interest here are the buffer size B (or rather the normalized buffer size, which is equal to the actual buffer size divided by the service rate c), the cutoff lag T_c, and the marginal fluid rate distribution Π. We let these parameters vary within ranges consistent with practical networking situations. We use normalized buffer sizes of up to a few seconds. These values are typical of currently available switches. For example, the Fore ATM 200BX/1000 switch has a per-port buffer of 13312 cells. Since the slowest available link on this switch is a 1.5 Mb/s T1 link, the maximum delay in the buffer is equal to 3.3 s. Higher link speeds would yield correspondingly smaller delays.

We use marginal distributions of fluid rates obtained from traces of various traffic sources. In this paper, we consider two traces. The first trace has been generated by JPEG-encoding an NTSC TV channel (MTV) for one hour. The trace has been recorded on June 11, 1995, 14:59 EST. It includes 107892 frames, with a mean rate of 9.5222 Mb/s. The second trace is based on the August 1989 "purple-cable" Ethernet trace collected at Bellcore [21]. Each trace element is a rate averaged over a 10 ms interval.

Both traces have been found to exhibit long-range dependence (refer to [21] for a detailed analysis of the Bellcore trace). This is visible in Figure 2, which shows the variance-time plot of the MTV trace. Estimations of the Hurst parameter from the variance-time plot (which has been shown to be a relatively poor estimator [1]) yield $H_{MTV} \approx 0.83$. The Hurst parameter for the Bellcore trace is 0.9, compatible with the findings in [21].

There still remains to match the marginal rate distribution of the traces to the fluid rate vector Π. Figure 3 shows the marginal fluid rate distributions for both traces. Recall

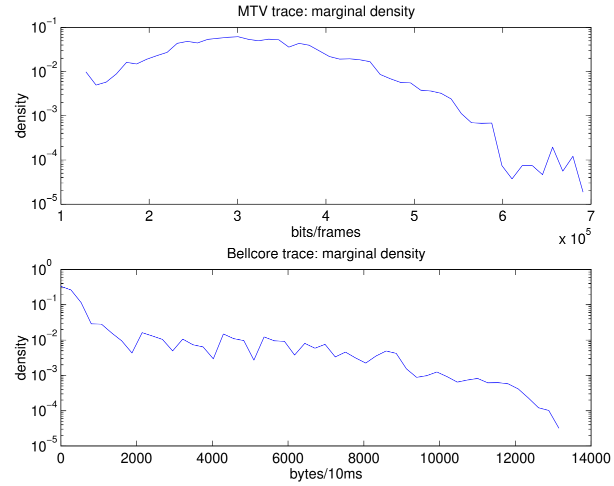

Figure 3: The marginal distributions of the MTV and the Bellcore traces.

that the traces represent the amount of work arriving within constant-length time intervals (33ms for the MTV trace, 10ms for the Bellcore trace). Thus, the marginal distribution vectors Π and the rate matrices Λ are simply obtained from a constant bin-size histogram of the traces. We set the number of bins to 50 in all experiments. We determine θ in (6) as follows. We first compute the average number of consecutive samples in the trace that fall within the same histogram bin. We then set θ such that the mean interval duration is equal to this value. We find from the trace data that the mean epoch durations are quite short, specifically about 80ms for the MTV trace and 15ms for the Bellcore traces.

We have now completed the description of the setup for our

[1] The overflow probability for a infinite buffer queue is an upper bound to the loss rate of the corresponding finite buffer queue.

numerical experiments. We next describe the experiments proper.

We have carried out three sets of experiments. In each set, we examine our chosen performance metric (i.e. the loss rate) as a function of two of the four parameters B, Π, H, and T_c. In the first set, we consider the impact of the buffer size and the cutoff lag on the loss rate. In the second set, we consider the impact of the Hurst parameter and the marginal distribution on the loss rate. In the third set, we consider the impact of the buffer size and the marginal distribution on the loss rate.

In the *first set of experiments*, we examine the impact of buffer size (or rather normalized buffer size) and cutoff lag on the loss rate. Figure 4 shows the loss rate for the MTV trace with a utilization equal to 0.8. Figure 5 shows the loss rate for the Bellcore trace with a utilization equal to 0.4. For clarity of presentation, the figures only show the upper bound for loss computed from (19). However, the discretization interval has been chosen small enough that the upper and lower bounds are essentially indistinguishable except for loss rates smaller than 10^{-10}.

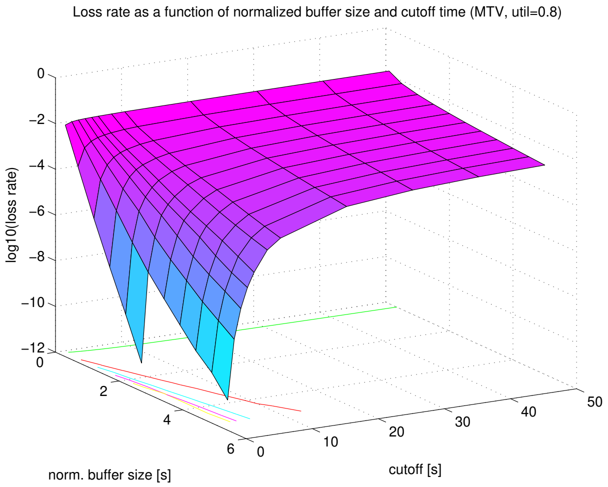

Figure 4: The loss rate predicted by the model for the MTV trace as a function of normalized buffer size and cutoff time, at utilization 0.8.

The figures bring out two important results. First, we observe that there exists a *correlation horizon (CH)* for each buffer size, such that the loss rate is not affected if the cutoff lag increases beyond CH. An important consequence of this is that it is sufficient for a model to take into account correlation up to this correlation horizon to accurately predict loss.

Second, we observe that the rate at which the loss rate decreases as the buffer size increases depends on the value of the cutoff lag. For small cutoff lags, the decrease is approximately exponential. However, as the cutoff lag increases, the rate of decrease actually decreases. This is an illustration of the "buffer ineffectiveness" phenomenon also reported elsewhere (e.g. [17]), whereby increasing buffer sizes beyond a certain value only slightly decreases loss rates. Note that this phenomenon is not unexpected, since an input process with non negligible correlation over long lags generates occasional bursts of traffic that cannot be absorbed even by

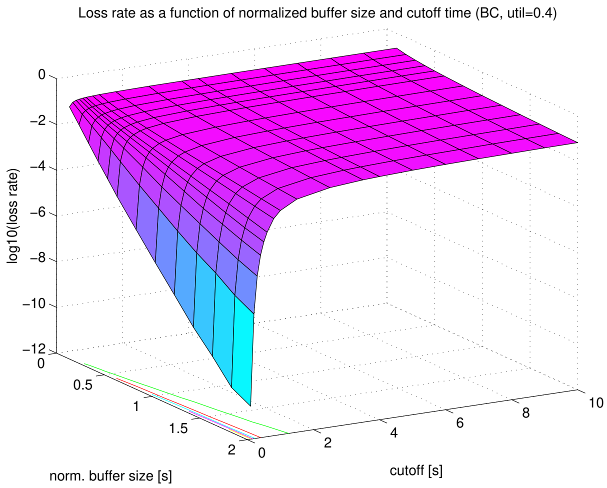

Figure 5: The loss rate predicted by the model for the Bellcore trace as a function of normalized buffer size and cutoff time, at utilization 0.4.

very large buffers. In such a case, decreasing the loss rate should be done by lowering the utilization or by statistically multiplexing several streams [14].

Recall that the cutoff lag T_c eliminates correlation in the input fluid process beyond a lag equal to T_c. Thus, its impact is similar to that of the "external shuffling" procedure described in [11]. In this procedure, a time series representing a realization of a process is divided up into blocks and the blocks are shuffled. However, the structure of the time series inside a block remains unchanged. Thus, external shuffling removes correlation from the series beyond a lag equal to the length of a block.

It is natural to compare the numerical results to those obtained by external shuffling of the input traces. We have carried out this comparison for both the MTV and the Bellcore traces. Figures 6 and 7 show the loss rate as a function of buffer size for different values of the shuffle block size (which is referred to as "cutoff" in the figures). We observe that the loss predicted by the model is very close to that obtained with shuffling and simulation for the MTV trace. The agreement is not so good with the Bellcore trace. The discrepancy is probably due to the fact that the residence time distribution for the fluid rates is not matched well enough in our model for this trace. The main results, however, namely the existence of a correlation horizon and the buffer ineffectiveness phenomenon for large values of the cutoff lag, still hold in both figures.

In the *second set of experiments*, we examine the impact of the marginal distribution and the Hurst parameter on the loss rate. To do this, we examine two transformations of the marginal distribution, namely a scaling and a convolution transformation. The first transformation scales the density of the marginal distribution by a constant factor α, while keeping the mean $\overline{\lambda} = \Pi\Lambda\mathbf{1}^T$ constant. Thus, we simply replace λ_i with $\lambda'_i = \overline{\lambda} + \alpha(\lambda_i - \overline{\lambda})$. Figure 8 shows the loss rate for H in the range $(0.5, 1.0)$ and α in the range $(0.5, 1.5)$. The normalized buffer size is set to 1 s. The cutoff lag is set to 100 s. Referring to Figure 4, we note that this cutoff lag

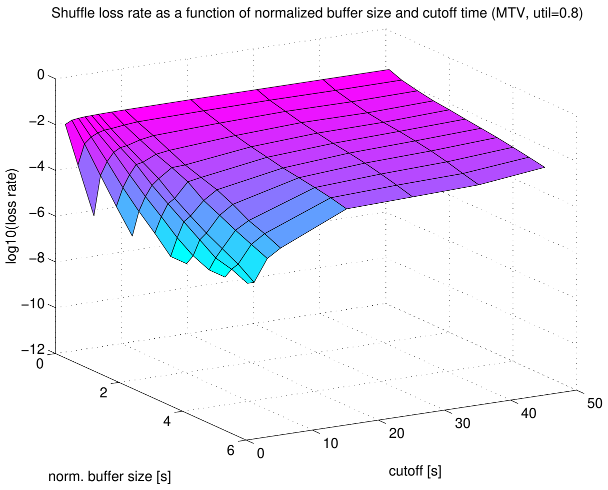

Figure 6: The loss rate obtained with shuffling for the MTV trace as a function of normalized buffer size and cutoff lag, at utilization 0.8.

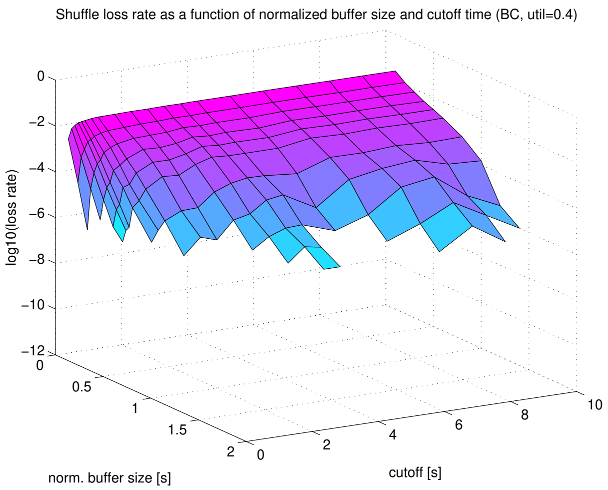

Figure 7: The loss rate obtained with shuffling for the Bellcore trace as a function of normalized buffer size and cutoff lag, at utilization 0.4.

is considerably larger than the CH.

The second transformation convolves the original distribution n times and renormalizes it to the original mean. Note that this transformation amounts to considering the superposition of n of the original streams, where the buffer size and the service rate per stream are kept constant. Figure 9 shows the loss rate for H in the range $[0.55, 0.95]$ and n in the range $1, \ldots, 5$.

The figures bring out a very interesting result, namely that the impact of the transformations on the loss rate is much greater than that of the Hurst parameter. For example, we observe in Figures 8 and 9 that changing α from 1.0 to 0.5 or superposing 3 streams decreases the loss rate by about 3 orders of magnitude. In contrast, changing the value of H

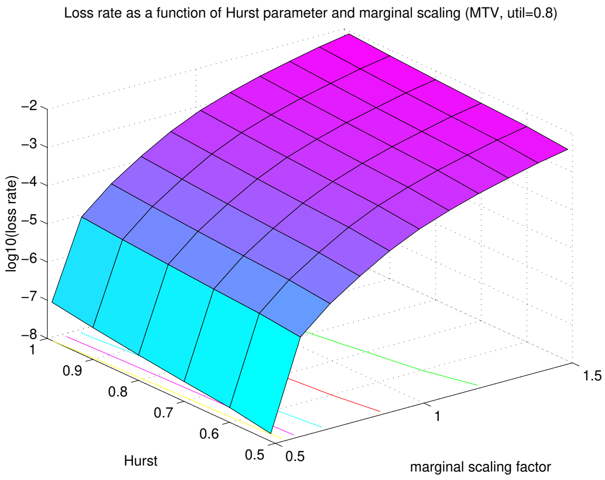

Figure 8: The loss rate predicted by the model for the MTV trace as a function of the Hurst parameter and the marginal distribution scaling, at utilization 0.8.

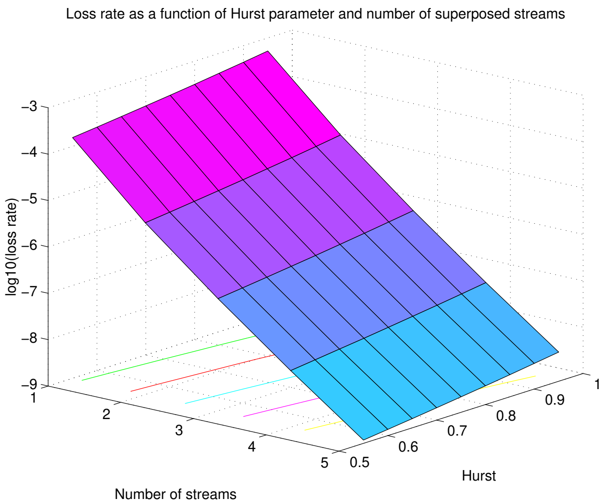

Figure 9: The loss rate predicted by the model for the MTV trace as a function of the Hurst parameter and the marginal distribution obtained by n convolutions of the original marginal distribution, at utilization 0.8.

has much less of an impact on the loss rate.

This result is confirmed in our *third set of experiments*, in which we examine the impact of the marginal distribution and the normalized buffer size on the loss rate. Figures 10 and 11 show that small changes in the marginal scaling factor again yield dramatic changes in the loss rate, more so than changes in the buffer size would do. For example, reducing the width of the marginal distribution by a factor of two (from $\alpha = 1$ to $\alpha = 0.5$) decreases the loss rate more than increasing the buffer size even up to 5 s (which is an extremely large value in practice).

The consequences of this are threefold.

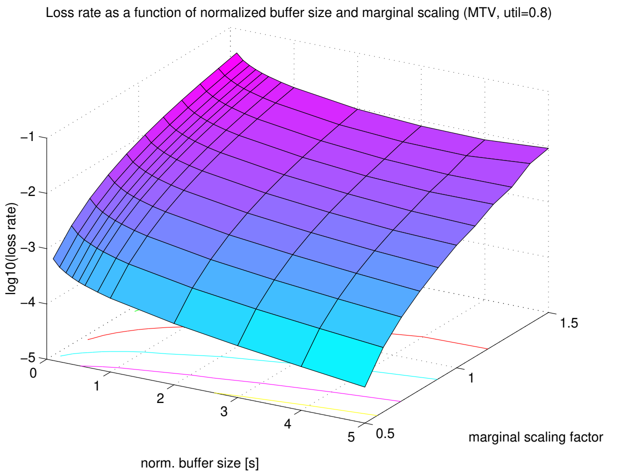

Figure 10: The loss rate for the MTV trace as a function of normalized buffer size and marginal scaling factor for a utilization of 0.8.

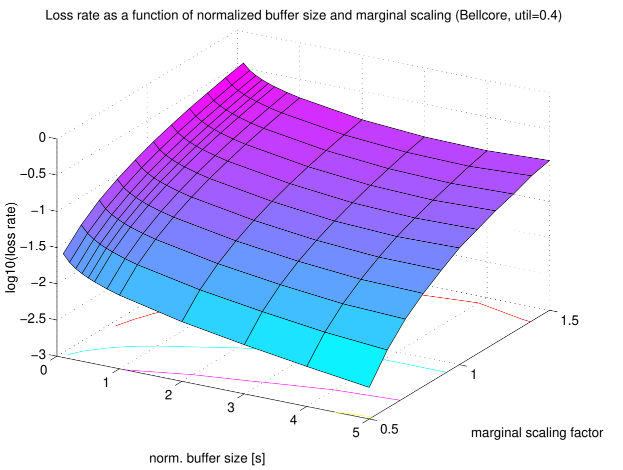

Figure 11: The loss rate for the Bellcore trace as a function of normalized buffer size and marginal scaling factor for a utilization of 0.4.

The first consequence relates to modeling. Clearly, the marginal distribution is a crucial parameter and it must be taken into account for accurate loss prediction. This is in agreement with results obtained by others using analytic approaches regarding the impact of the marginal distribution on the tail of the queue occupancy in infinite buffers (e.g [26]).

The second consequence relates to multiplexing. Our result above, namely that superposing even a moderate number of streams sharply decreases the loss rate, indicates that statistical multiplexing is an efficient mechanism (more so than buffering) to achieve high utilization while keeping loss low.

The third consequence relates to traffic control. Indeed, source traffic control mechanisms can be thought of as modifying the scaling of the marginal distribution. Our result then indicates that such control mechanisms (which we expect would typically be feedback-based) are effective for loss control with LRD sources.

4 Discussion

We start this section by summarizing the key findings from Section 3. These are

- There exists a correlation horizon (CH) such that the loss rate is not affected if the cutoff lag increases beyond CH. Thus, CH separates relevant and irrelevant correlation with respect to the loss rate.
- Large buffers are helpful to significantly reduce the loss rate only for short-range dependent traffic; for long-range dependent traffic, increasing the buffer size has little impact.
- The marginal scaling factor has considerable impact on the loss rate, even more so than the Hurst parameter or the buffer size.
- Adjusting the marginal scaling factor by statistical multiplexing several streams or by using source traffic control mechanisms is a very efficient way of reducing loss while keeping utilization high.

Two seemingly contradictory conclusions have usually been drawn in the literature from experiments with long-range dependent traffic. On the one hand, mathematical analysis of queueing systems with LRD input shows that the queue occupancy exhibits an asymptotic behavior very much different from that observed with Markov sources [24, 8]. This behavior has also been confirmed through simulation [11]. On the other hand, the literature on Markov modeling reports good performance prediction for finite buffer systems even when input traffic streams are correlated over many time-scales [10, 15, 29]. However, the contradiction is only apparent. Indeed, the existence of a finite buffer queue sets a limit on the memory of the system. This is because the buffer "forgets" about the past as soon as it is either empty *or* full (this is referred to as the *resetting effect* in [15]). We refer to the length of the time interval having an impact on the current system state as the *implicit time scale of the system*. Note that the implicit time scale of an infinite buffer is infinite, because the buffer can accumulate an arbitrary amount of information about the past. It becomes clear that analyzing finite or infinite queues in the presence of long-memory input processes leads to the seemingly contradictory results mentioned above. While in the infinite queue, correlation on all time-scales has an impact on performance, only the correlation up to the implicit system time scale has an effect in the finite buffer queue. This fits exactly our observation in Section 3 of the existence of the correlation horizon CH.

We next describe a simple way to estimate CH in terms of the various system parameters. The estimation procedure is based on the resetting argument mentioned earlier, i.e. we assume that when the buffer becomes empty or full, information about the past is lost. Then, we take the correlation horizon estimate T_{CH} to be the time interval for which the probability that the buffer empties or overflows at least once is close to one.

We make the assumption that the correlation horizon is much longer than the average interarrival time μ. If the converse were true, then the buffer would either empty or overflow with very high probability within a single interarrival interval. This would mean that the utilization is close to zero, or that the loss rate is extremely high.

Let $\Gamma_n = \sum_{i=1}^{n} W_n$ denote the sum of the excess work in n consecutive intervals, and $\Theta = \sum_{i=1}^{n} T_n$ denote the sum of the lengths of the n intervals. Now assume we are looking at the queue at an arrival instant τ_m and that the queue occupancy at that moment is equal to Q. Note that the probability of either emptying or overflowing the buffer at some point during the next n intervals is bounded below by $\Pr\{Q + \Gamma_n \leq 0\} + \Pr\{Q + \Gamma \geq B\}$. We wish to find n such that the probability $p = 1 - \Pr\{Q+\Gamma_n \leq 0\} - \Pr\{Q+\Gamma \geq B\}$ of no reset occurring during n intervals is very small.

Given our assumption that $T_{CH} \gg \mu$, we look for n large. Then, the central limit theorem says that for large values of n, $Q + \Gamma_n$ is approximately normally distributed with mean $Q + n\mu(\overline{\lambda} - c)$ and variance $n\sigma_T^2\sigma_\lambda^2$, where μ is the average interval length, $\overline{\lambda}$ is the average rate, σ_T^2 is the variance of the interarrival time, and σ_λ^2 is the variance of the marginal distribution. Furthermore, Θ is approximately normally distributed as well, with mean $n\mu$ and variance $n\sigma_T^2$.

As we do not explicitly know the distribution of Q, we find an upper bound on p by setting the mean of the normal distribution to $B/2$, which corresponds to selecting the value of Q that is least likely to result in a reset during the next n intervals. With this assumption, the upper bound on p becomes $p \leq \text{erf}(\frac{B}{2\sqrt{2}n\sigma_T\sigma_\lambda})$, where erf(.) denotes the error function. Let us now choose T_{CH} such that the probability that n intervals fit into T_{CH} is high, i.e. $T_{CH} = n(\mu + \beta\sigma_T)$, where β is a small constant (e.g. $\beta = 2$ results in this probability being ≈ 0.9.) Then we may write T_{CH} as

$$T_{CH} = \frac{B}{2\sqrt{2}\sigma_T\sigma_\lambda \text{erf}^{-1}(p)}(\mu + \beta\sigma_T) \tag{20}$$

The above estimate provides valuable insight into the impact on CH of the various parameters of the model. For example, we note that T_{CH} scales linearly with the buffer size, and that it is proportional to the inverse of the standard deviation of the marginal distribution. This is intuitively not surprising. Indeed, increasing the buffer size by a factor α implies that (the input rate being constant) it takes α times longer to overflow or empty the buffer, and thus to reach a resetting point. On the other hand, scaling the marginal distribution by a factor α has exactly the opposite impact (refer to Figure 8 and the discussion there). Finally, we observe that, as σ_T increases, T_{CH} converges to a limiting value that depends only on the buffer size and on the width of the marginal distribution. This is in agreement with the "flattening out" of the loss rate with increasing cutoff lag observed in Figures 4 to 7. This behavior of T_{CH} illustrates the above-mentioned implicit time-scale of the system.

5 Conclusion

We have focussed in this paper on a particular instance of performance prediction relying on traffic modeling, where the performance metric has been defined as the long-term loss rate. We have then shown that the relevant correlation is limited to a time-scale smaller than the correlation horizon. This might seem to support the claim that only models with limited memory (up to CH) make sense, and hence that self-similar models have no place in cases like these. This means instead that for this type of performance problem, we may choose any model among all the available models as long as it captures the correlation structure up to CH. The choice can be based on analytic tractability, on ease of parameter identification from traces, etc. Markov models are one possible choice since they can capture correlations up to some given value CH. And indeed, as we mentioned earlier, several studies have used Markov models for performance prediction with finite buffers and found them to work well.

However, we would like to stress that the amount of correlation that must be taken into account in a model also depends on the performance metric of interest in the model. We next illustrate this with a simple example. Assume that we would like to compare the performance of closed-loop (ARQ) and open-loop (FEC) error control schemes for reliable point-to-point communications. Let us try to guess the relevant time-scales of this problem. ARQ schemes perform well when losses are bursty because they can accumulate information about the loss of the whole burst and request retransmission of the whole block in one go. FEC schemes perform well when losses are spread out over time because they can correct errors of the type "among n packets, $k \leq k_{\max} < n$ have been lost". The probability that k exceeds $k_{\max}$ is smaller for independent losses than for correlated losses. This suggests that extending the time-scale of the correlation structure of the packet arrival process in a model of error control schemes amounts to increasing the advantage of ARQ over FEC. Therefore, it seems necessary for this problem to make an effort to correctly model the arrival and loss processes over a wide range of time-scales. Thus, a parsimonious self-similar model would be appropriate in this case.

Parsimonious models of self-similar processes capture the behavior of these processes over *all* time-scales using a small number of parameters. They appear to be particularly well suited in a number of different cases. First, in cases such as that described above for the ARQ/FEC evaluation when the system under study does not have a clearly bounded time-scale. Second, in cases when we wish to generate traces that match LRD behavior observed in actual networks. Since we do not want to restrict the use of these traces to a specific modeling task, we cannot make any assumptions about the relevance of the model parameters. Thus, traces should be derived from a self-similar traffic model. Of course, our results in Section 3 show that trace generation based on self-similar model (or any other model for that matter) should also take into account statistics such as the marginal distribution if we want to achieve accurate performance prediction.

Acknowledgments

Extensive discussions with Francois Baccelli, Philippe Nain and Darryl Veitch (INRIA Sophia Antipolis) have been a valuable contribution to this work. We have greatly appreciated useful comments from Debasis Mitra (Bell Labs) and Alain Jean-Marie (INRIA Sophia Antipolis). The MTV trace was recorded by one of the authors while a visitor at AT&T Bell Laboratories; thanks to James Pawlyk and Robert Safranek for their help with the recording equipment.

Finally, our thanks to the anonymous reviewers for the many helpful suggestions.

References

[1] P. Abry and D. Veitch. Wavelet Analysis of Long-Range Dependent Traffic. INRIA Technical Report, in preparation.

[2] R. Addie, M. Zukerman, and T. Neame. Fractal Traffic: Measurements, Modelling and Performance Evaluation. In *Proc. IEEE Infocom '95*, pages 977–984, Boston, MA, April 1995.

[3] J. Beran. Statistical Methods for Data with Long Range Dependence. *Statistical Science*, 7(4):404–427, 1992.

[4] J. Beran, R. Sherman, and W. Willinger. Long Range Dependence in Variable Bit Rate Video Traffic. *IEEE Trans. on Communications*, 43(3):1566–1579, February 1995.

[5] R. N. Bhattacharya, V. K. Gupta, and E. Waymire. The Hurst Effect under Trends. *Journal of Applied Probability*, 20:649–662, 1983.

[6] F. Brichet, J. Roberts, A. Simonian, and D. Veitch. Heavy Traffic Analysis of a Storage Model with Long-Range Dependent On/Off Sources. Technical Report, CNET, 1995.

[7] M. Crovella and A. Bestavros. Explaining World Wide Web Self-Similarity. Technical Report 95-015, Boston University, August 1995.

[8] N. G. Duffield, J. T. Lewis, N. O'Connell, R. Russell, and F. Toomey. Predicting Quality of Service for Traffic with Long-Range Dependence. In *Proc. IEEE ICC'95*, pages 473–477, Seattle, WA, September 1995.

[9] D. Duffy, A. McIntosh, M. Rosenstein, and D. Wilson. Statistical Analysis of CCSN/SS7 Traffic Data from Working CCS Subnetworks. *JSAC*, 12(3):544–551, 1994.

[10] A. Elwalid, D. Heyman, T. V. Lakshman, D. Mitra, and A. Weiss. Fundamental Bounds and Approximations for ATM Multiplexers with Applications to Video Teleconferencing. *IEEE J. Sel. Areas in Communication*, August 1995.

[11] A. Erramilli, O. Narayan, and W. Willinger. Experimental Queueing Analysis with Long-Range Dependent Packet Traffic. to appear in *IEEE/ACM Trans. on Networking*, April 1996.

[12] A. Erramilli, R. P. Singh, and P. Pruthi. An Application of Deterministic Chaotic Maps to Model Packet Traffic. *Queueing Systems: Theory and Applications*, to appear, 1995.

[13] M. Grasse, M. R. Frater, and J. F. Arnold. On the Non-Stationarity of MPEG-2 Video Traffic. Technical Report COST 242, University of New South Wales, 1995.

[14] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic. In *Proc. ACM Sigcomm '95*, pages 219–230, Boston, Mass., August 1995.

[15] D. P. Heyman and T. V. Lakshman. What are the Implications of Long-Range Dependence for VBR Video Traffic Engineering? Technical Report, Bellcore, 1995.

[16] H. E. Hurst. Long-Term Storage Capacity of Reservoirs. *Proc. American Society of Civil Eng.*, 76(11), 1950.

[17] C. L. Hwang and S. Q. Li. On Input State Space Reduction and Buffer Noneffective Region. In *Proc. IEEE Infocom '94*, Toronto, Canada, April 1994.

[18] Leonard Kleinrock. *Queueing Systems, Volume 1: Theory.* John Wiley & Sons, New York, NY, 1975.

[19] V. Klemes. The Hurst Phenomenon: A Puzzle? *Water Resources Research*, 10(4):675–688, 1974.

[20] S. Klivansky, A. Mukherjee, and C. Song. On Long-Range Dependence in NSFNET Traffic. Technical Report, Georgia Institute of Technology, 1995.

[21] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Trans. on Networking*, 2(1):1–15, February 1994.

[22] S. Q. Li and C. L. Hwang. Queue Response to Input Correlation Function: Continuous Spectral Analysis. *IEEE/ACM Trans. Networking*, 1(3):678–692, 1993.

[23] K. Meier-Hellstern, P. Wirth, Y-L. Yang, and D. Hoeflin. Traffic Models for ISDN Data Users: Office Automation Application. In *Proc. ITC'13*, pages 167–172, Copenhagen, June 1991.

[24] I. Norros. A Buffer with Self-Similar Input. *Queueing Systems: Theory and Applications*, 16(2):382–396, 1994.

[25] Alan V. Oppenheim and Ronald W. Schafer. *Discrete-Time Signal Processing.* Prentice-Hall, Englewood Cliffs, 1989.

[26] M. Parulekar and A. Makowski. Buffer Overflow Probabilities for a Multiplexer with Self-Similar Traffic. In *Proc. IEEE Infocom '96*, San Fransisco, CA, March 1996.

[27] V. Paxson and S. Floyd. Wide Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Trans. on Networking*, 3(3):226–244, June 1995.

[28] B. K. Ryu and S. Lowen. Point Process Approaches to Modeling and Analysis of Self-Similar Traffic. Part I: Model Construction. In *Proc. IEEE Infocom '96*, San Fransisco, CA, March 1996.

[29] P. Skelly, M. Schwartz, and S. Dixit. A Histogram-Based Model for Video Traffic Behavior in an ATM Multiplexer. *IEEE/ACM Trans. on Networking*, 1(4), August 1993.

[30] D. Veitch. Novel Models of Broadband Traffic. In *Proc. Globecom '93*, pages 362–368, Houston, TX, December 1993.

[31] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson. Self-Similarity through High Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level. In *Proc. ACM Sigcomm '95*, pages 100–113, Cambridge, MA, August 1995.