On the Relevance of Long-Range Dependence in Network Traffic

Matthias Grossglauser, *Member, IEEE*, Jean-Chrysostome Bolot

Abstract— **There is much 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 controlled. Specifically, our model is a modulated fluid traffic model in which the correlation function of the fluid rate matches that of an asymptotically second-order self-similar process with given Hurst parameter up to an arbitrary cutoff time lag, then drops to zero. We develop a very efficient numerical procedure to evaluate the performance of a 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 with 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. This means in particular that for performance modeling purposes, we may choose any model among the panoply of available models (including Markovian and self-similar models) as long as the chosen model captures the correlation structure of the source traffic up to the correlation horizon. Second, we find that loss can depend in a crucial way on the marginal distribution of the fluid rate process. Third, our results suggest that reducing loss by buffering is hard for traffic with correlation over many time scales. We advocate the use of source traffic control and statistical multiplexing instead.

I. 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 [19]. However, the phenomenon of LRD occurs in many other systems including chemical, astronomical, and biological systems (see [4] 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?" [22]. 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 [22] or a hyperbolically decreasing trend [6] 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 from its observation. 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 [25], [23], [10], [29], [5]. 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 [9] or Dirac pulses [15] with short range dependent (SRD) stationary processes. Another approach has been to use stochastic models (such as fractional Brownian motion [26], zero-rate renewal processes [35] and various other point processes [32]) or deterministic models (such as chaotic maps [13]) 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 why much modeling work still relies on more traditional multi-state Markovian models (e.g., [24], [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 [36], [7]. Furthermore, there is widespread evidence that human as well as computer sources of traffic do tend to behave as heavy tailed on/off sources [8], [36]. 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

This work was done while both authors were at INRIA Sophia Antipolis, France. This work was supported in part by a grant from France Telecom/CNET. An earlier version of this work appeared in ACM SIGCOMM '96.

Matthias Grossglauser is with the Network Mathematics Research Department, AT&T Labs - Research, Florham Park NJ 07932, USA (e-mail: mgross@research.att.com).

Jean Bolot is with Ensim Corporation, 1215 Terra Bella Avenue, Mountain View, CA 94043, USA (e-mail: bolot@ensim.com).

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 infinite buffer size and LRD in the input or service processes can be radically different from the performance of usual Markovian systems [26], [12], [28].

However, consider for example the asymptotic behavior of an infinite queue fed with three different arrival processes that all exhibit the LRD property: i) if the arrival process is a fractional Brownian motion, then the queue length distribution is Weibullian, ii) if the arrival process is a single on/off source with heavy-tailed on and off periods then the queue length distribution is hyperbolic, iii) 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 [7], [28]. 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, in particular low order statistics such as the marginal distribution of the arrival process.

Another parameter stands out, namely 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 controlled. Specifically, our model is a modulated fluid traffic model in which the correlation function of the fluid rate matches that of an asymptotically second-order self-similar process with given Hurst parameter up to an arbitrary cutoff time lag, then drops to zero. 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 ratio of 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 with 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 marginal distribution of the fluid arrival process. An obvious but important 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 much less efficient than controlling the loss rate by adjusting the marginal distribution. Statistical multiplexing and source traffic control are two efficient ways to do this, and to achieve high utilization while keeping loss low.

The rest of the paper is organized as follows. In Section II, we describe the model and the numerical solution procedure. In Section III, we describe the behavior of the loss rate as as function of system and traffic parameters. In Section IV, we discuss the implications of our results. Section V concludes the paper.

II. 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 whose lengths 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 ccdf $\Pr\{T_n > t\} = F_T(t)$. 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 autocovariance function of $X(t)$ is defined by

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

where

$$\mu = \Pi\Lambda\mathbf{1}^T \tag{2}$$

is the mean rate, and where $\mathbf{1} = [1, 1, \ldots, 1]$.

Let A_t denote the event of an arrival occurring within $[0, t]$. Then, noting that conditional on A_t, X_0 and X_t are independent, and conditional on A_t^c (the complement of event A_t), $X_0 = X_t$, we obtain

$$\begin{aligned}\phi(t) &= \mathrm{E}[(X_0 - \mu)(X_t - \mu)|A_t] + \mathrm{E}[(X_0 - \mu)(X_t - \mu)|A_t^c] \\ &= \sigma^2 p(t) \end{aligned}\tag{3}$$

where

$$\sigma^2 = \mathrm{E}[(X_0 - \mu)^2] = \Pi\Lambda^2\mathbf{1}^T - (\Pi\Lambda\mathbf{1}^T)^2 \tag{4}$$

is the variance of X_t.

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 [21, p. 172]

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

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

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

We observe that $\phi(t)$ behaves asymptotically as $t^{-\alpha+1}$ when $T_c \to \infty$. Thus, when $T_c = \infty$, $\{X_t\}$ is asymptotically second-order self-similar with Hurst parameter H such that $-\alpha+1 = -(2-2H)$, i.e., $H = \frac{3-\alpha}{2}$ [23]. When T_c is finite, the correlation drops to zero at lag T_c and $\{X_t\}$ has no LRD.

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. However, it is difficult in this model to control separately the correlation structure of the source over short and long time scales. Indeed, both parameters θ and α (or H) have an impact on the correlation structure at time scales shorter than T_c, and they cannot be cleanly separated into a "short term parameter" and a "long term parameter". This means the model is not well-suited for sources with separate structures for the short term and long term correlation, for example VBR video sources typically characterized by an exponential decrease in the short term followed by an hyperbolic decrease in the long term [14], [33].

We next examine the performance of a queue with constant service rate c and a finite buffer B fed with our fluid source model. The 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 that 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 $W(n) = T_n \times (\lambda(n) - c)$. The continuous random variable $W(n)$ is 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

$$Q(n+1) = \max(0, \min(B, Q(n) + W(n))). \tag{9}$$

$(W(n), Q(n))$ is a two-dimensional Markov process because $W(n)$ and $W(n-1)$ are independent. Assuming ergodicity, $(W(n), Q(n))$ converges to $(W(\infty), Q(\infty))$ as $n \to \infty$. Our goal is to derive 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 $W(n)$. Let $U(n) = Q(n) + W(n)$. Then

$$f_{U(n)}(x) = f_{Q(n)}(x) * f_W(x) \tag{11}$$

where $*$ denotes convolution. It follows that

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

where $\delta(.)$ is a Dirac impulse. In steady state, $f_{Q_{n+1}}(x) = f_{Q(n)}(x) = f_Q(x)$.

The random variable $W_l = (W - (B-Q))^+$ represents the amount of lost work in an interarrival interval. The performance measure of interest in this paper, namely the long-term stationary loss rate defined as the ratio of work lost to work arriving, is then

$$l = \frac{\mathrm{E}[W_l]}{\overline{\lambda}\mathrm{E}[T]} \tag{13}$$

with

$$\mathrm{E}[W_l] = \mathrm{E}[\mathrm{E}[W_l|Q]] = \int_0^B f_Q(x)\mathrm{E}[W_l|Q=x]dx \tag{14}$$

and

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

We now derive a numerical procedure to compute l. This procedure relies on computing an approximation of the stationary queue occupancy distribution by discretizing (12). We define two discretization operators ϕ_L^M and ϕ_H^M in the following way. Let $d = B/M$ be the quantization interval size, where B is the buffer size and M a positive integer. We define

$$\begin{aligned} \phi_L^M x &= d\lfloor x/d \rfloor \\ \phi_H^M x &= d\lceil x/d \rceil \end{aligned} \tag{15}$$

so that $\phi_L^M x \leq x \leq \phi_H^M x$.

We then define two discretized queueing processes $Q_L(n)$ and $Q_H(n)$ as follows

$$
\begin{aligned}
Q_L^M(n+1) &= \phi_L^M \max(0, \min(B, Q_L^M(n) + W(n))) \\
Q_H^M(n+1) &= \phi_H^M \max(0, \min(B, Q_H^M(n) + W(n)))
\end{aligned} \quad (16)
$$

with

$$
\begin{aligned}
Q_L^M(0) &= 0 \\
Q_H^M(0) &= B.
\end{aligned} \quad (17)
$$

The goal of the processes $Q_L^M(n)$ and $Q_H^M(n)$ is to give lower and upper bounds on the loss rate l. We formalize the properties of these two processes below. Assuming $\phi_{L,H}^M Q_{L,H}^M(n) = Q_{L,H}^M(n)$ (i.e., $Q_{L,H}^M(n)$ are discrete) it follows that $Q_{L,H}^M(n+1)$ are discrete as well, such that

$$
\begin{aligned}
Q_L^M(n+1) &= \max(0, \min(B, Q_L^M(n) + \phi_L^M W(n))) \\
Q_H^M(n+1) &= \max(0, \min(B, Q_H^M(n) + \phi_H^M W(n)))
\end{aligned} \quad (18)
$$

The above recursion, which is the discrete equivalent of (9), can be computed using a discrete convolution:

$$
u_{L,H}^n(i) = \sum_{j=0}^{M} q_{L,H}^n(j) w_{L,H}(i-j) \quad (19)
$$

where $q_{L,H}^n(j) = \Pr\{Q_{L,H}^M(n) = jB/M\}$. The discrete version of (12) becomes

$$
q_{L,H}^{n+1}(j) = \begin{cases} \sum_{i=-\infty}^{0} u_{L,H}^n(i) & j = 0 \\ u_{L,H}^n(j) & 0 < j < M \\ \sum_{i=M}^{\infty} u_{L,H}^n(i) & j = M \end{cases} \quad (20)
$$

As the probability mass of $u_{L,H}^n$ in $(-\infty, 0]$ and $[B, \infty)$ is concentrated in 0 and μ of $Q_{L,H}(n+1)$ (corresponding to the two terms involving $\delta(x)$ and $\delta(x-B)$ in (12)), we can actually constrain the length of $w_{L,H}$ and define

$$
w_L(i) = \begin{cases} \Pr\{W \in (-\infty, (i+1)d)\} & i = -M \\ \Pr\{W \in [id, (i+1)d)\} & -M < i < M \\ \Pr\{W \in [id, \infty)\} & i = M \end{cases} \quad (21)
$$

$$
w_H(i) = \begin{cases} \Pr\{W \in (-\infty, id]\} & i = -M \\ \Pr\{W \in ((i-1)d, id]\} & -M < i < M \\ \Pr\{W \in ((i-1)d, \infty)\} & i = M \end{cases} \quad (22)
$$

Then, the convolution (19) is between a vector of length $M+1$ and a vector of length $2M+1$.

We now examine how the discrete processes $Q_{L,H}^M(n)$ relate to the queueing process $Q(n)$, and how they depend on M and n. Recall that we introduced the discretized processes $Q_L^M(n)$ and $Q_H^M(n)$ as lower and upper bounds of the process $Q(n)$. Denote by $l(Q_L^M(n))$ and $l(Q_H^M(n))$ the loss rates computed using the discretized recursions (19) and (20) with $Q_L^M(n)$ and $Q_H^M(n)$, respectively. We next show that i) the value we are looking for, namely l, is bounded below and above by $l(Q_L^M(n))$ and $l(Q_H^M(n))$, respectively, and ii) $l(Q_L^M(n))$ is increasing in n and M and $l(Q_H^M(n))$ is decreasing in n and M. Thus, in practice, we recursively compute the loss bounds $l(Q_L^M(n))$ and $l(Q_H^M(n))$ using (19) and (20); if the bounds do not get tight enough, we multiply

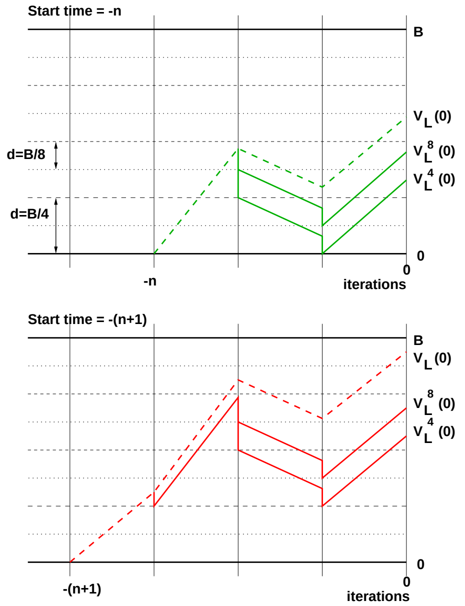

Fig. 1: An illustration of the sample path arguments to show the stochastic ordering relationships between $Q_L(n)$ and $Q_L^M(n)$ (resp. their shifted version $V_L^M(n)$). The situation for the upper bound is analogous.

M by some integer value m and resume computation until the desired accuracy is reached.

The main result is formalized in the following proposition.

Proposition II.1: The loss rate l is bounded by $l(Q_L^M(n))$ and $l(Q_H^M(n))$, i.e., $l(Q_L^M(n) \leq l \leq l(Q_H^M(n))$. Furthermore, $l(Q_L^M(n))$ is increasing in n and M, and $l(Q_H^M(n)$ is decreasing in n and M.

Proof: We sketch the main points of the proof. Refer to [3] and [31] for more details on sample path and Loynes type arguments. For clarity, we break down the proof into five steps. We denote stochastic order relations using the superscript "st" as in [31]. We focus on the lower bound in the proof; the corresponding results for the upper bound are derived in the same way.

(i) $Q_1 \overset{\text{st}}{\leq} Q_2$ implies $l_1 \leq l_2$: $\mathrm{E}[W_l|Q]$ is by definition an increasing function of Q. The result then follows from (14) and Proposition 9.1.2 in [31].

(ii) $Q_L(n) \overset{\text{st}}{\leq} Q \overset{\text{st}}{\leq} Q_H(n)$: A Loynes' type of argument shows that $Q_L(\infty) = Q = Q_H(\infty)$ (a.s.). We do not go through the details here (refer to [3] for details) but the argument goes as follows. The distribution of $Q_L(n)$ is equal to the distribution $V_L(0)$ of a process $V_L(.)$ which is a shifted version of $Q_L(n)$ starting at $-n$, i.e., $V_L(-n) = 0$. A straightforward sample path argument illustrated in Figure 1 shows that the value of every sample path at time 0 cannot decrease if the start time of process V_L is changed from $-n$ to $-(n+1)$. Therefore, $Q_L(n)$ is stochastically increasing in n. The fact that Q is the limiting distribution of $Q_L(n)$ as $n \to \infty$ shows the result.

(iii) $Q_L^M(n) \overset{\text{st}}{\leq} Q_L(n)$, $Q_H^M(n) \overset{\text{st}}{\geq} Q_H(n)$: From (9) and (18), it follows that if $Q_L^M(n) \leq Q_L(n)$, then $Q_L^M(n+1) \leq Q_L(n+1)$. The result then follows by induction.

(iv) $Q_L^M(n) \overset{\text{st}}{\leq} Q_L^M(n+1)$, $Q_H^M(n) \overset{\text{st}}{\geq} Q_H^M(n+1)$: We show this again using Loynes' argument. The distribution of $Q_L^M(n)$ is equal to the distribution $V_L^M(0)$ of a process $V_L^M(.)$

which is a shifted version of $Q_L^M(n)$ starting at $-n$, i.e., $V_L^M(-n) = 0$. Then, changing the start time of process V_L^M from $-n$ to $-(n+1)$, it follows from (9) that each sample path can only increase or remain constant.

(v) $Q_L^{mM}(n) \overset{st}{\geq} Q_L^M(n)$, $Q_H^{mM}(n) \overset{st}{\leq} Q_H^M(n)$, where m is an integer: A discretization with interval length d/m divides the interval $[0, B]$ in such a way that each interval of length d/M is entirely contained in a interval of length d corresponding to a coarser discretization. Therefore, $\phi_L^M x \leq \phi_L^{mM} x$, and $\phi_H^M x \geq \phi_H^{mM} x$. Applying this argument to the sample paths of the shifted processes $V_L^M(.)$ and $V_L^{mM}(.)$ shows the result (cf. Fig. 1).

Now, (iv) and (v) imply that $Q_L^M(n)$ is stochastically increasing in n and in M, and because of (i), $l(Q_L^M(n))$ is increasing in n and M. Similarly, $Q_H^M(n)$ is stochastically decreasing in n and in M, and therefore $l(Q_H^M(n))$ is decreasing in n and M. Finally, (i), (ii) and (iii) show that these are lower and upper bounds of $l = l(Q)$, respectively. This completes the proof. ∎

Thus, we obtain a lower and an upper bound on the limiting queue occupancy by initializing $q_L^0 = [1, 0, 0, \ldots, 0, 0]$ and $q_H^0 = [0, 0, \ldots, 0, 0, 1]$ which corresponds to starting the recursion with an initially empty and with an initially full queue, respectively, and by iterating (19). The lower and upper bounds on l are then obtained as

$$l(Q_L^M(n)) = \frac{\sum_{i=0}^{M} q_L^n(i) \mathrm{E}[W_l | Q = id]}{\overline{\lambda} \mathrm{E}[T]} \leq l \leq \quad (23)$$

$$l(Q_H^M(n)) = \frac{\sum_{i=0}^{M} q_H^n(i) \mathrm{E}[W_l | Q = id]}{\overline{\lambda} \mathrm{E}[T]} \quad (24)$$

We have found that in this approach, the lower and upper bounds converge rapidly to the limiting occupancy[1]. This is illustrated in Figure 2, which shows the upper and lower bounds $Q_{L,H}^M(n)$ for the queue occupancy after $n = 5$, 10, and 30 iterations when the queue has been discretized using $M = 100$ bins (i.e., $d = B/100$). Note that it is possible to improve the efficiency of the discrete convolution (19) by using a fast Fourier transform (FFT) with appropriate zero-padding, which reduces the computational complexity from $O(M^2)$ to $O(M \log M)$ [27].

III. 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 buffer queues[2]. 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).

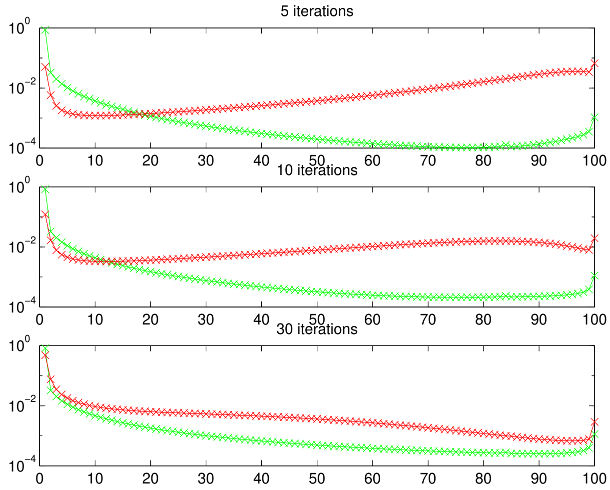

Fig. 2: An example of the discrete upper and lower bounds $Q_{L,H}^M(n)$ for $n = 5, 10, 30$ iterations and $M = 100$ (dark: upper bound; light: lower bound).

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 [23]. Each trace element is a rate averaged over a 10 ms interval.

Both traces exhibit long-range dependence. Using a Whittle or wavelet based estimator [1], we obtained $H_{MTV} \approx 0.83$ for the MTV trace and $H_{BC} \approx 0.9$ for the Bellcore trace. This latter value is consistent with the findings in [23] (refer to [23] for a detailed analysis of the Bellcore trace).

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 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

[1] In our experiments, the typical runtime was less than a second on a workstation; however, when the expected interarrival time $E[T_n]$ is very small (i.e, a small T_c and/or small θ), B is very large, and the utilization close to one, the runtime can be considerably longer.

[2] The overflow probability, i.e., the probability that the queue occupancy exceeds some amount, in an infinite buffer queue is an upper bound to the loss rate in the corresponding finite buffer queue.

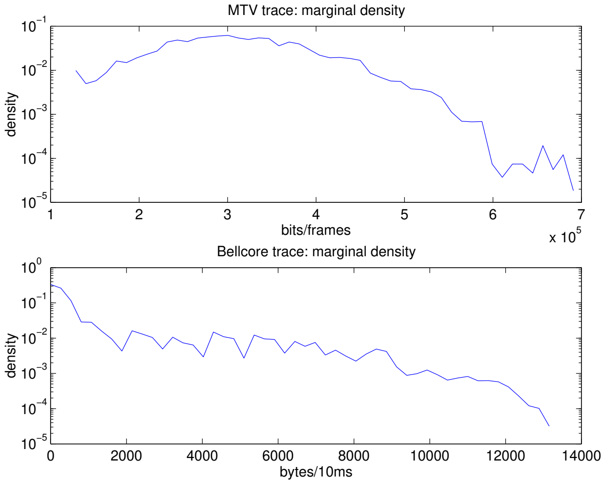

Fig. 3: The marginal distributions of the MTV and the Bellcore traces.

duration, which is given by

$$\mathrm{E}[T] = \frac{\theta}{\alpha - 1}\left[1 - \left(\frac{T_c}{\theta} + 1\right)^{1-\alpha}\right] \tag{25}$$

matches this empirical mean for $T_c = \infty$. 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 trace.

All the numerical results reported in this section have been obtained as follows. Recall that the numerical procedure described in Section II provides an upper and a lower bound on the loss rate. The numbers reported below are the average of the upper and lower bounds. We require that the distance between the upper and the lower bound be less than 20% of the average of the two bounds. If the upper bound of the loss rate falls below 10^{-10}, we consider that the loss rate is simply below a value that is of practical importance, and report zero loss. We perform iterations until either the above confidence criterion is fulfilled, in which case we stop, or until the convergence of the two bounds becomes too slow, which means that the discretization step d is too large to achieve sufficient accuracy. In this case, we double the number of "bins" M, and resume iterating.[3]

We have now completed the description of the setup for our 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 (namely 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. We have chosen different values for the utilization for the different traces in order to be in a regime of practically relevant loss rates (roughly between 10^{-1} and 10^{-10}) for both traces.

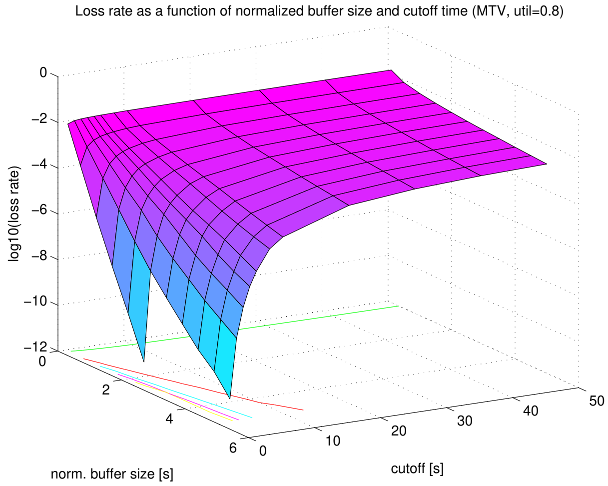

Fig. 4: The loss rate predicted by the model for the MTV trace as a function of normalized buffer size and cutoff lag, at utilization 0.8.

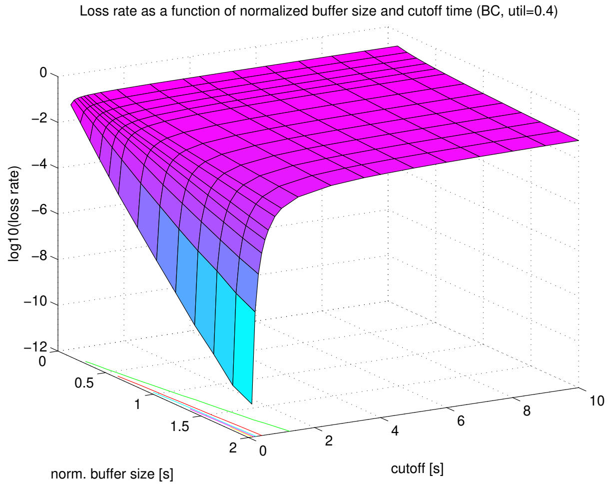

Fig. 5: The loss rate predicted by the model for the Bellcore trace as a function of normalized buffer size and cutoff lag, at utilization 0.4.

The figures bring out two important results. First, we observe that for each buffer size, the loss rate is not significantly affected if the cutoff lag increases beyond some value. We refer to this value as the *correlation horizon (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., [20]), whereby increasing buffer sizes beyond a certain value only slightly decreases loss rates. This phenomenon is not unexpected, since an input process with non negligible correlation over long lags generates occasional

[3] Note that when requantizing from M to $mM = 2M$, we do not have to reset $Q^{mM}_{L,H}(n)$ to 0 and B (as at the very beginning), but instead use the most recent values obtained with M bins, namely $Q^{M}_{L,H}(n)$. This considerably increases the efficiency of the procedure when the discretization step is small.

bursts of traffic that cannot be absorbed even by very large buffers.

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 [12]. In this procedure, a time series representing a realization of a process is divided into blocks and the blocks are shuffled (cf. Fig. 6). 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[4].

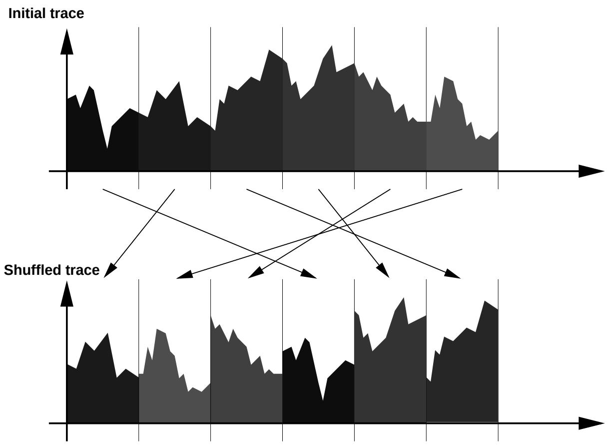

Fig. 6: Shuffling can be used to eliminate correlation beyond a given cutoff lag in a trace.

Therefore, it is natural to compare the numerical results above to a trace-driven simulation in which the input traffic is obtained by external shuffling of the MTV and the Bellcore traces. Figures 7 and 8 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). Note that the results in the figures have been obtained directly with the shuffled data used as input to a simulated queue; thus, they are completely independent of the stochastic traffic model described in Section II. 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 caused by a poor match for the residence time distribution for the fluid rates 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, can be observed 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. As an initial motivation for this set of experiments, consider the result in Figure 9, which shows the loss rate as a function of T_c for two marginal distributions matching the Bellcore and MTV distributions, all other parameters in the model being equal (normalized buffer size = 1s, utilization = 2/3, θ = 20, and H=0.9).

We observe orders of magnitude differences in the loss rate, clearly suggesting a large impact of the marginal distribution. To examine this impact further, we consider 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

[4]In fact, our choice for the interarrival time distribution function $F_T(t)$ as being a truncated Pareto function was motivated in part by the analogy with the shuffling procedure.

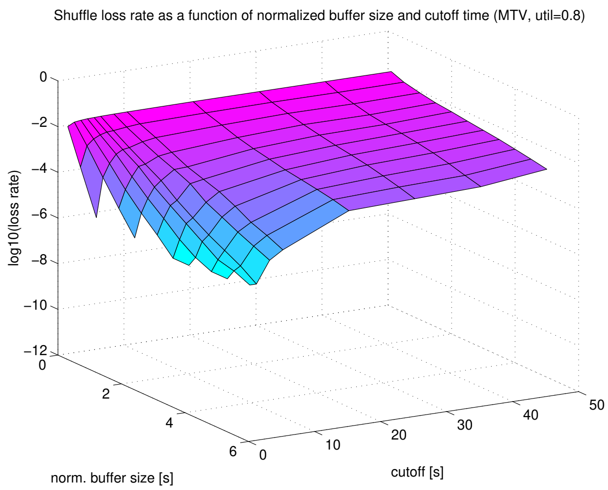

Fig. 7: The loss rate obtained with shuffling for the MTV trace as a function of normalized buffer size and cutoff lag, at utilization 0.8.

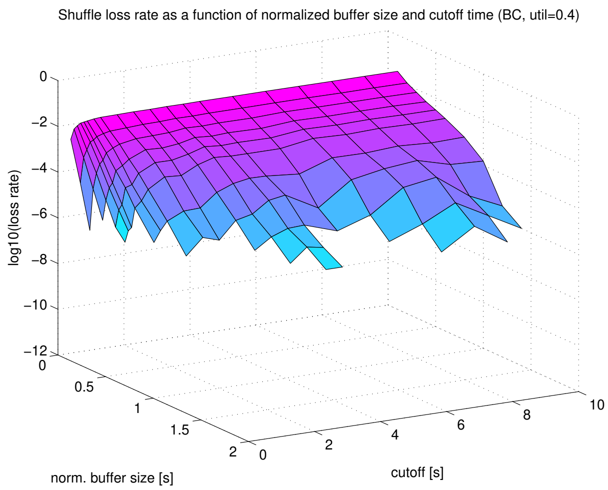

Fig. 8: The loss rate obtained with shuffling for the Bellcore trace as a function of normalized buffer size and cutoff lag, at utilization 0.4.

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 10 shows the loss rate for H in the range $(0.55, 0.95)$ and α in the range $(0.5, 1.5)$. The normalized buffer size is set to 1 s. The cutoff lag is set to infinity. Note that we use the same θ in the entire experiment, by matching the average interval length for the *nominal* Hurst parameter. We do so in order to avoid that a larger Hurst parameter decreases short-range dependence, which would perturb the assessment of the impact of the Hurst parameter.

The second transformation convolves the original distribution n times and renormalizes it to the original mean. Thus, it 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 11 shows the loss rate for H in the range $(0.55, 0.95)$ and n in the range $1, \ldots, 10$.

The figures bring out an interesting result, namely that the impact of the transformations on the loss rate for the range of parameters considered above is much greater than that of the Hurst parameter. For example, we observe in Figures 10 and 11 that changing α from 1.0 to 0.5 or superposing 5 streams decreases the loss rate by more than an order of magnitude. In contrast, changing the value of H has much less of an impact on the loss rate.

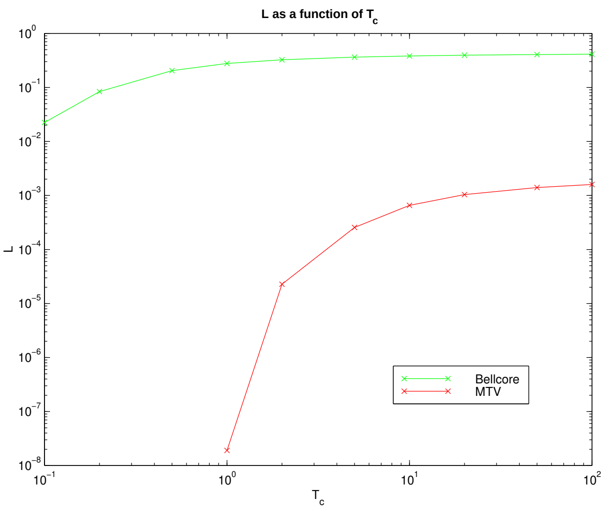

Fig. 9: The loss rate for the MTV and Bellcore marginal distributions as a function of T_c, all other parameters being equal (normalized buffer size = 1s, utilization = 2/3, θ = 20, H=0.9).

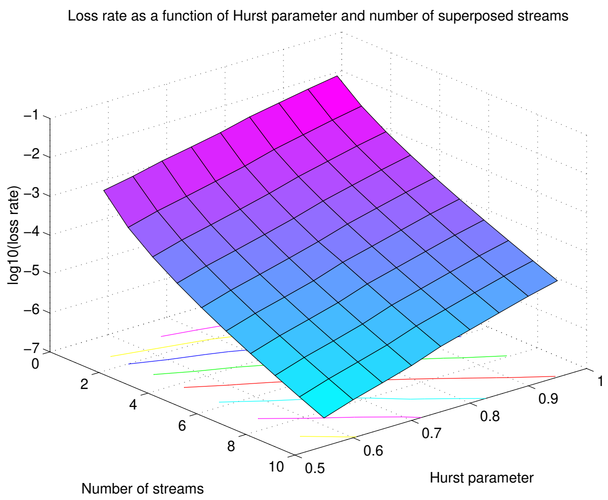

Fig. 11: The loss rate predicted by the model for the MTV trace as a function of the Hurst parameter and the marginal distribution obtained by n convolutions of the original marginal distribution, at utilization 0.8.

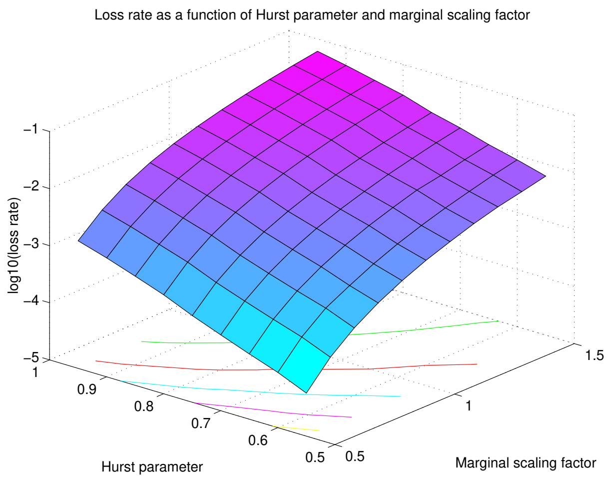

Fig. 10: The loss rate predicted by the model for the MTV trace as a function of the Hurst parameter and the marginal scaling factor, at utilization 0.8.

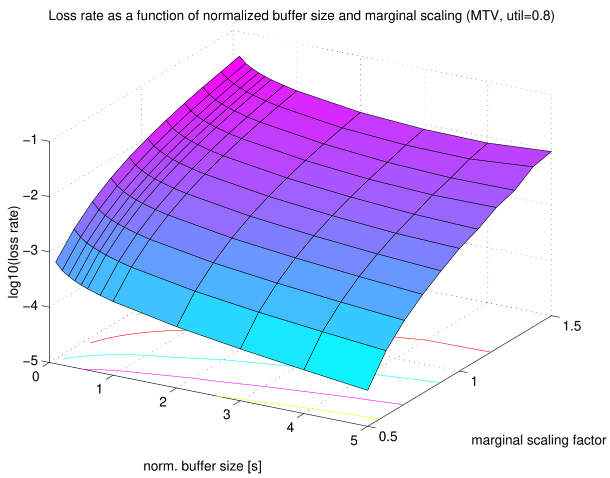

Fig. 12: The loss rate for the MTV trace as a function of normalized buffer size and marginal scaling factor for a utilization of 0.8.

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 12 and 13 show that small changes in the marginal scaling factor again yield dramatic changes in the loss rate: 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. 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 [28]). 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 [16]. The third consequence relates to traffic control. The ability to change the marginal distribution and get very different loss rates as a result suggests it would be useful to examine control mechanisms for LRD sources that modify the scaling of the marginal distribution. One example of this would be a feedback-based rate control mechanism.

IV. Discussion

We start this section by summarizing the key results from Section III. 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.
- 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.

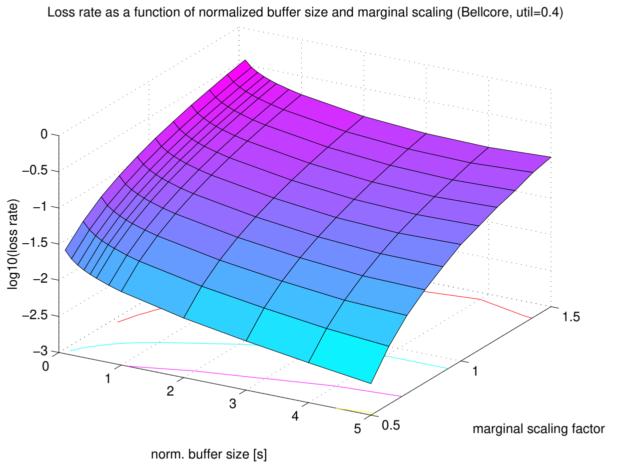

Fig. 13: The loss rate for the Bellcore trace as a function of normalized buffer size and marginal scaling factor for a utilization of 0.4.

Our observation of the existence of the correlation horizon CH helps resolve the seemingly contradictory conclusions that have been drawn in the literature from experiments with long-range dependent traffic. Indeed, mathematical analysis and simulation of queueing systems with LRD input shows that the queue occupancy exhibits an asymptotic behavior very much different from that observed with Markov sources [26], [9], [12]. However, the literature on Markov modeling reports good performance prediction for finite buffer systems even when input traffic streams are correlated over many time-scales [11], [17], [34], [18]. We have shown that there exists a correlation horizon which separates relevant and irrelevant correlation with respect to the performance measure of interest. Intuitively, the CH depends on the correlation structure of the input traffic, and on the system under study, namely the finite buffer queue. Indeed, the finite buffer queue sets a limit on the memory of the system since the the buffer "forgets" about the past as soon as it is either empty or full (this is referred to as the resetting effect in [17]). Therefore, we expect the CH to depend on the maximum queue size (we provide evidence for this below). Furthermore, while correlation on all time-scales has an impact on performance for the infinite queue, only the correlation up to the CH has an effect in the finite buffer queue. Ryu and Elwalid, in [33], independently came to similar conclusions. They use large deviations theory to derive an expression for the correlation horizon (called *Critical Time Scale*). Their traffic process is the superposition of a LRD and a SRD process. While their approach is computationally less demanding than ours, they do not study the impact on the correlation horizon of the marginal distribution and of different Hurst parameters. Also, their performance metric is the buffer overflow probability for an infinite queue, not the loss rate.

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 $\Gamma_n = \sum_{i=1}^{n} W(n) - T(n)c$ denote the sum of the excess work, assuming the server is always busy, in n consecutive intervals. Now assume that we are looking at the queue at an arrival instant τ_m and that the queue occupancy at that instant 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 $\Pr\{Q + \Gamma_n \leq 0\} + \Pr\{Q + \Gamma_n \geq B\}$. We wish to find n such that the probability $p = 1 - \Pr\{Q + \Gamma_n \leq 0\} - \Pr\{Q + \Gamma_n \geq B\}$ 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.

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 $Q + \Gamma_n$ to $B/2$. We obtain $p \leq \mathrm{erf}(\frac{B}{2\sqrt{2}n\sigma_T\sigma_\lambda})$, where erf(.) denotes the error function. Since we assume that n is large, T_{CH} can be well approximated by

$$T_{CH} = \frac{B\mu}{2\sqrt{2}\sigma_T\sigma_\lambda \mathrm{erf}^{-1}(p)} \qquad (26)$$

As suspected earlier, we find that T_{CH} scales linearly with the buffer size. Interestingly, it turns out that the trace-based simulation results we described earlier (e.g., in Figure 7) do bring out this linear scaling quite clearly. For example, Figure 14 shows the same results as Figure 7, namely the loss rate for the MTV trace as a function of the normalized buffer size and cutoff lag, but this time drawn with logarithmic scales and a different point of view. We see that the surface on the figure flattens out along a curve that parallels the straight line $B/T_c = \gamma$ on the floor of the figure (where γ is a constant), indicating a linear scaling between T_{CH} and B.

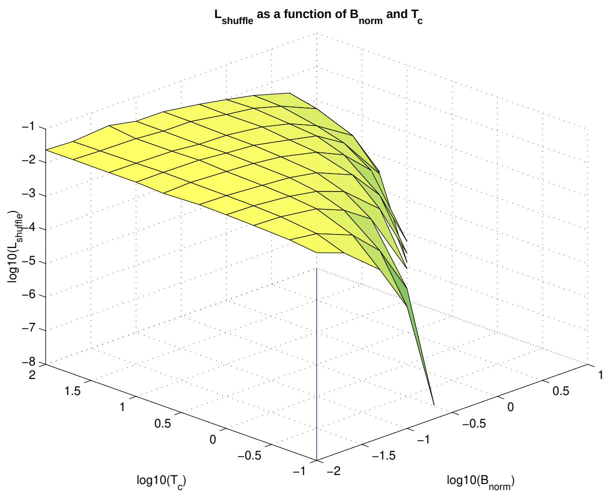

Fig. 14: The loss rate obtained with shuffling for the MTV trace as a function of the normalized buffer size and cutoff lag. The line on the (x,y)-plane (and all lines parallel to it) correspond to $B/T_c = \text{const.}$

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 are not useful after all. But instead, this means 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. In this paper, we chose the truncated Pareto model because it is self-similar when T_c is set to infinity, and because it is a parsimonious model with a simple way to control its correlation structure (via T_c). However, Markov models could have been another possible choice since they can capture correlations up to a given value CH (and since the numerical procedure developed in Section II can be used independent of the particular model).

Indeed, several studies have used an approach where Markov models approximate traffic sources with long range dependence. This approach can be used to obtain accurate performance predictions since a power law decay can be approximated arbitrarily closely by enough exponential decay functions [24]. However, the resulting Markov models typically are complex multi-state models that do not follow the principle of parsimonious modeling because every state added to such a model also adds several free parameters. This presents two problems, namely that of identifying the parameters (states and state transition rates) from experimental data, and that of obtaining closed-form analytic expressions for performance measures. The first problem is the more important one because it is often difficult and time consuming in practice to collect the data required for parameter estimation. This problem is generally used to promote instead parsimonious long-range dependent models.

However, parsimonious modeling and LRD are really orthogonal issues. Indeed, it is possible to reduce the impact of the parameter-explosion problem with multi-state Markov models, where each state models a different time scale [30].

V. 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 CH, meaning that for this type of performance problem, *we may choose any model among the panoply of available models (including Markovian and self-similar models) as long as the chosen model captures the correlation structure up to CH.*

We would like to conclude by stressing that the amount of correlation that must be taken into account in a model also depends in general on the performance metric of interest. We illustrate this with a simple example. Consider the problem of comparing 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 a loss burst and request retransmission of all packets lost in the burst 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 in this problem to accurately model the arrival and loss processes over a wide range of time-scales. Thus, a self-similar model would be appropriate here.

In general, models of self-similar processes capture the behavior of these processes over *all* time-scales using a small number of parameters. They are particularly well suited to study systems which to not have a clearly bounded time scale (such as in the evaluation of ARQ/FEC above) or to generate very long traces that match LRD behavior observed in actual networks. In this latter case, we do not want to restrict the use of these traces to a specific modeling task, and thus we cannot make any assumptions about the relevance of the model parameters. Therefore, traces should be derived from a self-similar traffic model. Of course, our results in Section III show that trace generation based on self-similar models (or any other model for that matter) must also take into account statistics such as the marginal distribution in order to achieve accurate performance prediction.

Acknowledgments

Extensive discussions with Francois Baccelli, Philippe Nain, Jim Roberts, and Darryl Veitch have been a valuable contribution to this work. We have also greatly appreciated comments from Debasis Mitra, Alain Jean-Marie, and Jean Walrand.

The MTV trace was recorded by one of the authors (MG) 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 many helpful comments and suggestions.

This work has been supported in part by a grant from France Telecom/CNET.

References

[1] P. Abry and D. Veitch. Wavelet Analysis of Long-Range Dependent Traffic. *IEEE Trans. on Information Theory*, 44(1):2–15, January 1998.

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

[3] François Baccelli and Pierre Brémaud. *Palm Probabilities and Stationary Queues*. Springer Verlag, 1996.

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

[5] 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.

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

[7] F. Brichet, J. Roberts, A. Simonian, and D. Veitch. Heavy Traffic Analysis of a Storage Model with Long-Range Dependent On/Off Sources. *Queueing Systems*, 23(2):197–215, 1996.

[8] M. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proc. ACM Sigmetrics '96*, pages 160–169, Philadelphia, PA, May 1996.

[9] 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.

[10] 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, April 1994.

[11] 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*, 13(6):1004–1016, August 1995.

[12] A. Erramilli, O. Narayan, and W. Willinger. Experimental Queueing Analysis with Long-Range Dependent Packet Traffic. *IEEE/ACM Trans. on Networking*, 4(2):209–223, April 1996.

[13] A. Erramilli, R. P. Singh, and P. Pruthi. An Application of Deterministic Chaotic Maps to Model Packet Traffic. *Queueing Systems*, 20(3):171–206, 1995.

[14] M. W. Garrett and Walter Willinger. Analysis, Modeling and Generation of Self-Similar VBR Video Traffic. In *Proc. ACM SIGCOMM '94*, pages 269–280, London, UK, August 1994.

[15] 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.

[16] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic. *IEEE/ACM Transactions on Networking*, 5(6):741–755, December 1997.

[17] D. P. Heyman and T. V. Lakshman. What are the Implications of Long-Range Dependence for VBR Video Traffic Engineering? *IEEE/ACM Trans. on Networking*, 4(3), June 1996.

[18] D. P. Heyman, T. V. Lakshman, and A. Neidhardt. A New Method for Analyzing Feedback Protocols with Applications to Engineering Web Traffic over the Internet. In *Proc. ACM SIGMETRICS '97*, Madison, WI, May 1997.

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

[20] 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.

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

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

[23] 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.

[24] 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.

[25] 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.

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

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

[28] 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.

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

[30] S. Robert and J-Y. Le Boudec. New models for self-similar traffic. *Performance Evaluation*, 30(1-2):57–68, 1997.

[31] Sheldon Ross. *Stochastic Processes (Second Edition)*. John Wiley & Sons, 1996.

[32] 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.

[33] Bong K. Ryu and Anwar Elwalid. The Importance of Long-Range Dependence of VBR Video Traffic in ATM Traffic Engineering: Myths and Realities. In *Proc. ACM SIGCOMM '96*, San Francisco, Calif., August 1996.

[34] 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.

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

[36] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson. Self-Similarity through High Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level. *IEEE/ACM Trans. Networking*, 5(1):71–96, 1997.

Matthias Grossglauser (S '92, M '99/ACM S '96, M '99) received the Diplôme d'Ingénieur en Systèmes de Communication from the Swiss Federal Institute of Technology (EPFL) in 1994, the M.Sc. degree from the Georgia Institute of Technology, in 1994, and his Ph.D. from the University of Paris 6 in 1998. He did most of his thesis work at INRIA Sophia Antipolis, France. Since November 1998, he has been a member of the Networking and Distributed Systems Laboratory of AT&T Labs - Research, Florham Park, NJ. His research interests are in network traffic analysis and modeling, resource allocation, and network management. Dr. Grossglauser received the 1998 Cor Baayen Award from the European Research Consortium for Informatics and Mathematics (ERCIM).

Jean-Chrysostome Bolot received his MS and PhD degrees from the University of Maryland at College Park in 1988 and 1991, respectively. From 1991 until 1999, he was at INRIA in Sophia-Antipolis, France. His work there focused on understanding Internet traffic, and on using this understanding to design adaptive applications, in particular interactive multicast applications such as Internet telephony and gaming applications. Since early 1999, he has been with Ensim Corp. in Mountain View, CA, working on a scalable service deployment platform for ISPs.