Experimental Queueing Analysis with Long-Range Dependent Packet Traffic

Ashok Erramilli, *Member, IEEE,* Onuttom Narayan, and Walter Willinger, *Member, IEEE*

***Abstract*— Recent traffic measurement studies from a wide range of working packet networks have convincingly established the presence of significant statistical features that are characteristic of fractal traffic processes, in the sense that these features span many time scales. Of particular interest in packet traffic modeling is a property called long-range dependence (LRD), which is marked by the presence of correlations that can extend over many time scales. In this paper, we demonstrate empirically that, beyond its statistical significance in traffic measurements, long-range dependence has considerable impact on queueing performance, and is a dominant characteristic for a number of packet traffic engineering problems. In addition, we give conditions under which the use of compact and simple traffic models that incorporate long-range dependence in a parsimonious manner (e.g., fractional Brownian motion) is justified and can lead to new insights into the traffic management of high speed networks.**

I. Introduction

In the past two to three years, large amounts of traffic measurements from working packet networks [including Ethernet local area networks (LAN's), wide area networks (WAN's), common channel signal network CCSN/SS7, integrated services digital network (ISDN), and variable bit rate (VBR) video over asynchronous transfer mode (ATM)] have been collected and analyzed. The results reported in [1], [5], [7], [13], [23], [25], [26], [33], [35], [36], and [40] have been striking for two reasons: 1) these studies demonstrate that it is possible to clearly distinguish between actual packet network traffic and traffic generated by widely employed theoretical models, and 2) in sharp contrast to the traditional packet traffic models, aggregate packet streams are statistically *self-similar* or *fractal* in nature; that is, realistic network traffic looks the same when measured over time scales ranging from milliseconds to minutes and hours. Strictly speaking, these results mostly concern the statistical nature of traffic processes, and do not address in depth issues of the relevance of these features to queueing performance or practical traffic management.

Moving beyond the statistical nature of the findings of these recent traffic studies, we demonstrate in this paper the impacts of self-similarity on queueing performance. To this end, we concentrate on the property of *long-range dependence* (LRD), one of a number of different equivalent mathematical manifestations of the property that the underlying traffic process is self-similar. A long-range dependent process is characterized by an autocorrelation function that decays as a power of the lag time, implying that the sum (over all lags) of the autocorrelations diverges. This divergence captures the intuition behind LRD, namely, that even though the high-lag autocorrelations are individually small and become negligible, their cumulative effect is of importance, giving rise to behavior that is drastically different (e.g., non-Markovian) from that of traffic processes currently considered in the teletraffic literature. The latter are almost exclusively Markovian in nature or, more generally, *short-range dependent,* i.e., the corresponding autocorrelation functions decay exponentially fast.

Performing a number of queueing simulation experiments with actual traces of Ethernet LAN traffic and certain randomly "shuffled" versions thereof, we illustrate the behavior of a simple queue as a function of the dependence structure imposed on the packet arrival stream. The main idea behind using shuffled versions of a given arrival streams is to obtain traffic traces with identical packet interarrival time distributions but with differing autocorrelation structures; in particular, we consider here shuffling experiments that make the packet interarrival times independent, retain the same short-range correlations, or exhibit the same long-range correlations as the original Ethernet traffic trace. This way, we show that LRD is not only relevant for queueing performance but that it is, in fact, a dominant characteristic for several packet traffic engineering issues, such as dimensioning of buffers and determining usable capacity. The first of this paper's principal contributions is to add to the current efforts of gaining a better understanding of queueing performance when the input to the queue is not given by a traditional traffic process but is instead fractal in nature. For recent analytic results and simulation studies in this area, see, for example, [4], [7], [8], [12], [13], [26], and [34].

Given the statistical significance of the finding of self-similarity or LRD in measured packet traffic (e.g., in case of the Ethernet data see [26]) and its demonstrated significance for queueing performance, stochastic modeling of long-range phenomena becomes of crucial importance. Traditional approaches include mimicking LRD with the help of short-range dependent models. This is equivalent to approximating a correlation function decaying as a power law by a sum of exponentials; although always possible, the number of parameters

Manuscript received September, 1994; revised September, 1995; approved by IEEE/ACM Transactions on Networking Editor H. Miyahara. This work was supported by the Institute for Theoretical Physics at Santa Barbara, CA, under National Science Foundation Grant PHY-94-07194, and by Boston University under National Science Foundation Grant NCR-9404931.

A. Erramilli and W. Willinger are with Bellcore, Morristown, NJ 07960-6438 USA.

O. Narayan is with the Physics Department, University of California at Santa Cruz, Santa Cruz, CA 95064 USA.

Publisher Item Identifier S 1063-6692(96)02714-8.

required in this approach will tend to infinity as the sample size increases. Such approaches are pursued, for example in [22], [27], and [28], and can be used successfully for solving certain queueing performance problems numerically. However, in this paper we argue strongly in favor of modeling LRD based on *the principle of parsimony,* also known as *Occam's Razor* (see, for example, [21]). The paper's second major contribution consists of giving conditions under which parsimonious traffic models that capture LRD (e.g., fractional Brownian models) are appropriate and result in accurate and practically relevant solutions to performance problems of high speed networks that carry fractal-like traffic. Self-similar traffic models address the problem of obtaining parsimonious *descriptions* of complex traffic processes, though the complete *analysis* of these models is an area for further research.

The idea behind parsimonious modeling is to explain facts in as economical a way as possible. Parsimony is essential in practical packet engineering: it is impossible in current practice to collect the large volume of operational measurements needed to fit highly parameterized models. Each additional measurement imposes a penalty on the capacity of network switching and transmission systems, as well as on the operational systems, which currently lack the capacity to collect, process, transport, and store any measurements beyond coarse time scale rate measurements. In practice, it is such considerations, and not analytical tractability, that limits the application of queueing models. Thus models that can effectively describe the behavior of the traffic in terms of a minimum number of parameters are of great practical significance. Put differently, we believe that a model that adequately represents the queueing behavior of the traffic (as we shall demonstrate in this paper) and involves only a few parameters is more likely to be a good representation of the underlying processes intrinsic to the traffic, and that its parameters should change less under a change in the operating conditions.

The rest of the paper is organized as follows. In Section II, we briefly describe the traffic measurements used in our study and summarize their statistical properties. Section III gives a description of the queueing simulation experiments that identify LRD as a crucial ingredient for queueing performance. In Section IV, we address the problem of modeling long-range dependent traffic processes in a parsimonious manner, and identify the conditions under which the model accurately describes bursty traffic. In Section V, we discuss the resulting statistical and traffic engineering implications, as well as some of the outstanding issues in the queueing analysis of long-range dependent input processes. We conclude with a summary of the paper in Section VI.

II. LRD in Actual Network Traffic

With the large scale deployment of ISDN and high speed data networks such as switched multimegabit data service (SMDS) and frame relay, and the emergence of broadband networks, there has been renewed interest in traffic measurement studies. Prominent among these recent studies are the statistical analyses of i) Bellcore's Ethernet LAN traffic measurements [25], [26], ii) WAN traffic traces collected at Berkeley [35], [36], iii) traffic measurements collected from working Common Channel Signaling (CCS) subnetworks [5], iv) packet traces collected in an ISDN office automation environment [33], and v) VBR video traces [1], [13], [20]. Note that because of the use of highest quality traffic monitoring equipment and because of current data storage capabilities that are practically unlimited, the resulting sets of traffic measurements are unique in terms of quality and size.

For the queueing experiments performed in this paper, we rely exclusively on the Bellcore Ethernet LAN traffic measurements. Similar results have been obtained using some of the other data traffic sets mentioned above. For illustration purposes, our examples below are drawn from aggregate LAN traffic recorded during August 1989. While the nature of the applications that run on the LAN have changed considerably since 1989 (e.g., web browsing and the Mbone service on the Internet now make up a significant fraction of the traffic), the essential self-similar nature of the traffic has not changed. For the August 1989 data set, the internal traffic data (consisting of all packets on the LAN, regardless of source or destination) came from a network operating at a speed of 10 Mbps and serving a laboratory of about 120 people, most of whom had a Sun-3, Sun-4, or DECstation 3100 workstation on their desk; it also served nine file servers and a small number of high-end minicomputers and was connected to the rest of Bellcore via a router. For more details about this data collection effort and the subsequent data analysis, we refer to [26] and [40].

From a statistical analysis viewpoint (for details, see [1] and [40]), the most surprising finding from the aforementioned studies concerns the ease with which it is possible to statistically distinguish between measured network traffic and traditional model-generated traffic. For example, while covariance-stationary traffic processes $X = \{X_k\}$ currently considered in the teletraffic literature are exclusively *short-range dependent,* i.e., exhibit autocorrelations $r_X(k)$ that decay exponentially fast

$$r_X(k) \sim a^{|k|}, \quad \text{as} \quad |k| \to \infty, \quad 0 < a < 1 \tag{1}$$

(here and henceforth, $\sim$ denotes that the expressions on the two sides are asymptotically proportional to each other). These studies demonstrate convincingly that the autocorrelation function $r_X(k)$, $k \geq 0$, of a realistic packet traffic process X does not decay exponentially fast but has instead the form

$$r_X(k) \sim |k|^{-\beta}, \quad \text{as} \quad |k| \to \infty \tag{2}$$

where $0 < \beta < 1$. Stochastic processes satisfying relation (2) are said to exhibit *long-range dependence* or, using Mandelbrot's terminology, the Joseph Effect.[1] Long-range dependence captures the persistence phenomenon observed in many empirical time series that manifests itself in clusters ("bursts") of consecutive large (or small) values. Long-range dependent processes are non-Markovian in nature and give rise to features that are drastically different from those of traditional short-range dependent processes. In particular, note

[1] A reference to the Biblical figure who foretold of the "seven fat years and seven lean years" that ancient Egypt was to experience.

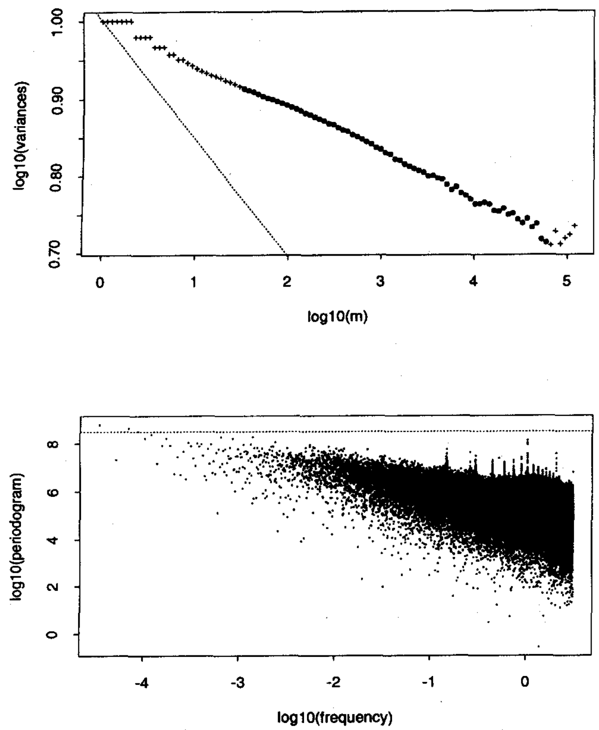

Fig. 1. Variance-time plot (top) and periodogram plot (bottom) for an hour-long Ethernet traffic trace. The dotted reference lines in the variance-time and periodogram plots have slopes -1.0 and 0.0, respectively.

that the latter give rise to a summable autocorrelation function $0 < \sum_k r_X(k) < \infty$, while LRD implies nonsummability of the correlations, i.e., $\sum_k r_X(k) = \infty$. In the frequency domain, LRD manifests itself in a spectral density $s_X(\omega) = \sum_k r_X(k)e^{ik\omega}$ that obeys a power-law near the origin ($1/f$-noise phenomenon), e.g.,

$$s_X(\omega) \sim |\omega|^{-\gamma}, \quad \text{as} \quad \omega \to 0 \tag{3}$$

where $0 < \gamma < 1$;[2] on the other hand, short-range dependent processes are characterized by a spectral density that remains finite as $\omega \to 0$. The *Hurst parameter* H is commonly used to measure the degree of LRD, and is related to the parameters β in (2) and γ in (3) by $H = 1 - \beta/2 = (1 + \gamma)/2$ (for short-range dependent processes, $H = 1/2$).

This clear statistical distinction between measured network traffic and traditional traffic processes is significant in a number of ways. On the one hand, it allows for surprisingly simple graphical methods for clearly distinguishing between measured data and traditional model-generated traffic (see, for example, Fig. 4 in [26]). On the other hand, it enables one to approach the problem of inference for data with LRD from a number of different angles, utilizing both time domain and frequency domain techniques (for details, see [1] and [26]). For example, Fig. 1 depicts the *variance-time plot* and *periodogram plot* obtained from the time series describing the number of Ethernet packets per 10 ms during a normal traffic hour from the August 1989 traffic measurements. The variance-time plot is obtained by plotting $\log\{\text{var}[X^{(m)}]\}$ against $\log(m)$, where for each $m = 1, 2, \cdots$, the *aggregated process* $X^{(m)} = \{X_k^{(m)}\}$ is obtained by averaging the original traffic process X over nonoverlapping intervals of size $10m$ milliseconds. Then for short-range dependent traffic processes

$$\text{var}[X^{(m)}] \sim m^{-1}, \quad \text{as} \quad m \to \infty \tag{4}$$

while processes with long-range dependence can be characterized by

$$\text{var}[X^{(m)}] \sim m^{-\beta}, \quad \text{as} \quad m \to \infty, \quad 0 < \beta < 1. \tag{5}$$

The variance-time plot technique exploits this difference in the rate of the decay of $\text{var}[X^{(m)}]$; while values of the estimate $-\hat{\beta}$ of the resulting asymptotic slope between -1 and 0 indicate LRD, short-range dependent processes are characterized by an asymptotic slope of -1 (dotted reference line). An estimate for the Hurst parameter H is then given by $\hat{H} = 1 - \hat{\beta}/2$. For the Ethernet traffic data in Fig. 1, the asymptotic slope parameter is readily estimated to be about -0.40, resulting in a Hurst parameter estimate of $\hat{H} \approx 0.80$. In case of the periodogram plot, we rely on the fact that for a given set of observations $(X_1, X_2, \cdots, X_n)$ from a long-range dependent process X, the corresponding periodogram $I_X(\omega) = (2\pi n)^{-1}|\sum_j X_j e^{ij\omega}|^2, 0 \le \omega \le \pi$, has an expectation value of $s_X(\omega)$. Using (3), we see that this decreases linearly (at least for the low frequencies) in $\log-\log$ plots against ω with a negative slope. In contrast, short-range dependent processes result in periodogram plots that are flat (i.e., zero slope) around the origin (dotted reference line). Letting $-\hat{\gamma}$ denote the estimate of the resulting asymptotic slope in the periodogram plot near the origin ($1 > \hat{\gamma} > 0$), an estimate of the Hurst parameter of X is given by $\hat{H} = (1 + \hat{\gamma})/2$. The periodogram plot in Fig. 1 resulting from the Ethernet data yields $\hat{\gamma} \approx 0.60$, and hence the same Hurst parameter estimate of about 0.80. Additional estimation techniques, including R/S-analysis and Whittle's method (e.g., see [26]), can be used to confirm and refine aspects related to statistical inference for data with LRD.

Mathematically (for $1/2 < H < 1$), autocorrelations that decay hyperbolically [i.e., satisfy relation (2)], variances of the aggregated processes that decrease more slowly than the reciprocal of the sample size [see relation (5)], and spectral densities that exhibit the $1/f$-noise phenomenon [i.e., see (3)] are different manifestations of the property that the underlying traffic process is statistically self-similar. Another characteristic that can be said to span many time scales is the *heavy-tailed* nature of the densities describing traffic processes such as packet interarrival times and burst lengths. We will refer collectively to these properties as the *fractal properties* of measured packet traffic, using the popular notion of fractals to describe phenomena that span many length or time scales. Mathematically, a covariance stationary packet traffic process X satisfying (5) is called *asymptotically second-order self-similar* with self-similarity parameter $H = 1 - \beta/2$ if for all sufficiently large m, $r_X^{(m)}(k) \sim r_X(k)$, as $k \to \infty$, where $r_X^{(m)}$ denotes the autocorrelation function of the aggregated

[2] From the correlation function in the frequency domain, it is clear why $\beta < 1$ for long-range dependent processes, since $\beta = 1 - \gamma$ and $\gamma > 0$.

process $X^{(m)}$. *Exactly second-order self-similar processes* have the asymptotic proportionality strengthened to an exact equality for all k and all m. An example of an exactly self-similar process with self-similarity parameter H is *fractional Gaussian noise* with parameter $1/2 < H < 1$ [see (see [31] for more details); *fractional ARIMA* (p, d, q) *models* (see [15], [19]) are examples of asymptotically self-similar processes with self-similarity parameter $H = d + 1/2$, $0 < d < 1/2$. For other approaches to characterizing the fractal properties of measured packet traffic that make explicit use of certain fractal dimension descriptors, see [7], [9], and [10].

III. Experimenting with Measured Ethernet Traffic Traces

While the studies mentioned above convincingly establish the presence of LRD over a wide range of time scales in packet traffic processes (see also Fig. 1), its significance to queueing performance and traffic engineering may not be clear at this point. For example, it has been argued by some that since queueing performance is determined by features in arrival processes of the time scales of a queueing system's busy period, long-range dependence has no practical impact and need not be incorporated into performance models. In this section, we describe a series of simulation experiments that demonstrate—contrary to this argument—the practical significance of LRD in queueing performance.

A. Three Simple Queueing Experiments

We consider a queueing system with the following characteristics: infinite waiting room, deterministic service times, a single server, and arrivals taken from *actual* Ethernet traffic traces. The input traces consist of the measured interarrival times of actual Ethernet traffic; while the results shown below are obtained with a single 30-min long Ethernet traffic trace, similar results (not shown here) also hold for other Ethernet traffic data sets, as well as for some of the other traffic traces (e.g., ISDN) mentioned in Section II. A 30-min interval is representative of the time scales over which traffic in packet networks is currently measured (15–60 min); for example, traffic levels are typically reported at 30-min intervals, and rate and utilization measurements over these intervals are used as baselines in "load-service" curves. The underlying assumption in engineering practice is that the traffic environment is stationary over such time scales. While this assumption is not always satisfied in practice, it does appear to be a reasonable hypothesis for the 30-min trace used in our studies. In particular, the variability of a number of relevant traffic statistics estimated across subsets of the 30-min trace is, within confidence limits, consistent with the stationarity hypothesis. In our experiments, we also perform various transformations on these Ethernet traffic traces. For this purpose, we choose to work with interarrival time traces, primarily to preserve the marginal interarrival time distribution throughout the different queueing experiments. By keeping the packet interarrival time distributions of the traces used in our experiments identical, we are able to isolate the possible effects of the underlying dependence structure on queueing and separate them from those due to the distributional aspects of the interarrival times. A similar experiment investigating the queueing impact of the latter by holding the former fixed would be of interest but is not part of the present study. To achieve different utilizations of the queue for a given input trace, we adjust the service time of the deterministic server, and consider the average waiting time, as well as the asymptotic form of the queue length distribution.

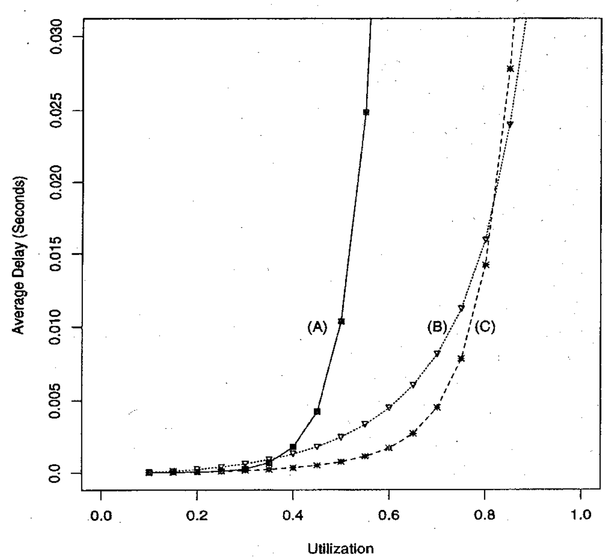

Fig. 2. Average delay (in seconds) versus utilization plot for original trace (A), QNA-based approximation (B), and fully shuffled trace (C).

Curve (A) in Fig. 2 is the average delay versus utilization plot obtained with the original trace. From a traffic engineering perspective, the "knee of the curve" is of particular interest. As can be seen, there is a sharp rise in the average delay around 50% utilization. By way of a comparison, Fig. 2 also shows [curve (B)] the delay curve predicted by the Queueing Network Analyzer (QNA). QNA uses a set of GI/G/1 approximations (based on two moment characterizations of the input traffic) that are widely applied in practice. In contrast to the delay curve obtained with the actual trace, curve (B) predicts useful capacities in excess of 80% utilization. To identify the features of the input trace that contribute to the sharp rise in delays at relatively low utilizations, we repeat the simulation experiments with transformations of the original traffic trace that preserve some aspects of statistical characteristics normally associated with burstiness, while eliminating others.

In the second experiment, we repeat the simulation with an input trace that is obtained by shuffling the time series of interarrival times. Note that by randomizing the set of interarrival times, we preserve the *marginal distribution* of the interarrival times, while destroying all *correlations* between them. The delay curve obtained with this input is the curve marked (C) in Fig. 2. As can be seen, the delay performance obtained with this renewal process is somewhat different from that predicted by QNA. More significantly, the disparity between curves (A) and (C) indicates that the best renewal model will grossly underestimate queueing delays at moderate and high utilizations. In particular, this experiment suggests that the burstiness of traffic *cannot* simply be explained by a renewal arrival process with a singular distribution. Such models have been considered in [38], following earlier work by Mandelbrot and co-workers [29] in a different context. In these models, the probability $P(T > t)$ that the interarrival time T is greater than t decays slowly with t when t is large

$$P(T > t) \sim t^{-\alpha}, \quad \text{as} \quad t \to \infty \tag{6}$$

with $1 > \alpha > 0$. Since $\alpha < 1$, the mean interarrival time, $\int_0^\infty dt\, P(T > t)$, *diverges,* so that the mean traffic density approaches zero when measured over a sufficiently long interval. On the other hand, since $-dP/dt$ diverges at $t = 0$, most of the interarrival times will be quite small.[3] There are thus bursts of activity separated by long periods of inactivity; as one observes the system for longer and longer times, one runs into larger and larger inactive stretches, leading to the vanishing of the mean traffic density. Despite the fact that the interarrival time process is renewal—one interarrival time is not correlated with any other—the "heavy" tail of the distribution of T leads to unusual features; in particular, delays increase sharply at surprisingly low utilization factors. This effect certainly *is* present in the Ethernet data: $P(T > t)$ for the measured Ethernet data does decay slowly at large t, as in singular renewal models, resulting in weak long-range dependence (in the time series of counts) in even the shuffled trace used to generate curve (C). We believe that this is responsible for the eventual intersection of the curves (B) and (C). However, Fig. 2 also establishes—curve (C) is much closer to (B) than to (A)—that any aspect of the interarrival time distribution (whether the residual long-range dependence created by it or something else) is not the *main* source of the unexpectedly high waiting times for the traffic. This is not to say, however, that *aggregates* of singular renewal models will not be useful in modeling packet traffic, as discussed further in Section V-A.

Given that the properties of the packet interarrival time marginals do not determine the sharp increase in delays seen with the original trace, it is apparent that correlations in the arrival process are responsible for this behavior. We now wish to distinguish between one-step correlations, short-term correlations, and long-range dependence. The traffic traces do show complex short-term correlations, including striking one-step correlations. This is apparent by studying phase plots of the time series of interarrival times, i.e., T_{i+k} versus T_i. Fig. 3 is a phase plot of T_{i+1} versus T_i. If the interarrival times were independent, Fig. 3 would at most have only vertical and horizontal lines, and more complicated structures (like the diagonal lines seen) would be absent. Explaining how such structures arise in a packet stream is in itself an intriguing exercise; these features are probably not Ethernet protocol specific, because similar phase plots are obtained with ISDN packet traffic traces as well. Given that there is some degree of quantization in the interarrival times, the horizontal lines correspond to the first packet in a burst, while the vertical lines correspond to the last packet in a burst. The diagonal lines are probably caused by a superposition of two or more

[3] Of course, in reality (6) must be rounded off for very small t, but this will not affect our results.

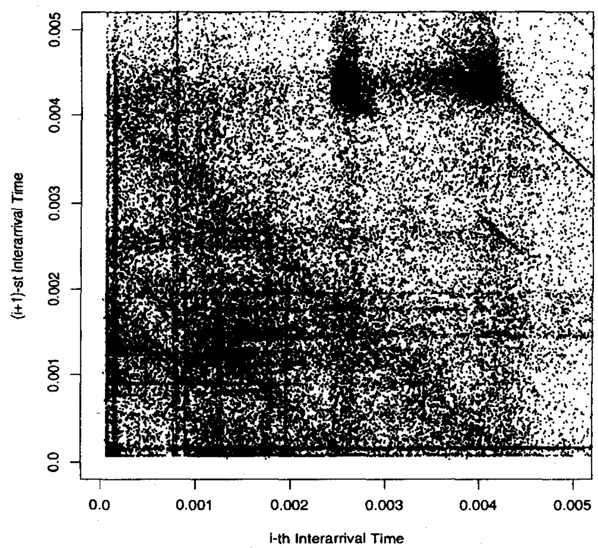

Fig. 3. Phase plot of interarrival times T_{i+1} versus T_i (in seconds) for 200 000 packets.

sources which generate deterministically spaced interarrival times when they are on.

To consider the performance implications of the one-step correlations in Fig. 3, we repeat the simulation with an input trace that is derived from the phase plot. For a given T_i, interarrival time T_{i+1} is randomly picked from the possible set of interarrival times indicated by the phase plot. The input trace thus preserves properties of the marginal distribution, as well as the one-step correlations. Curve (D) in Fig. 4 shows the resulting delay performance, and the disparity between this and the original trace [curve (A)] is once more considerable. Thus, the one-step correlations are not responsible for the sharp rise in the delay curve observed with the original trace, which must arise from the short-term and/or long-term correlations.

B. Two Experiments with Shuffled Data

A modified version of the above test can be used to ascertain the importance of long-range and short-range correlations in the data. Considering the sequence of interarrival times corresponding to the same set of Ethernet data as above, we divide the sequence into blocks of size m; with N interarrival times, there are N/m such blocks. We then perform an "external shuffle," i.e., the order of the blocks is shuffled, while preserving the sequence inside each block. For choices of m in the range 10–100, this has the effect of essentially preserving the short-range correlations while eliminating the long-range correlations. Note that for a given m, the number of packets in a block is fixed, while the time duration of a block can vary significantly. Experiments with fixed block durations (and variable numbers of packets) are briefly described later.

Curve (E) in Fig. 4 is the delay curve obtained with an input trace resulting from an external shuffle ($m = 25$). For the 30 min trace used in the experiments described here, $m = 25$ corresponds to an average block duration of about 76 ms, with block durations ranging from 14–629 ms. Even though this trace captures the properties of the marginal distribution, as

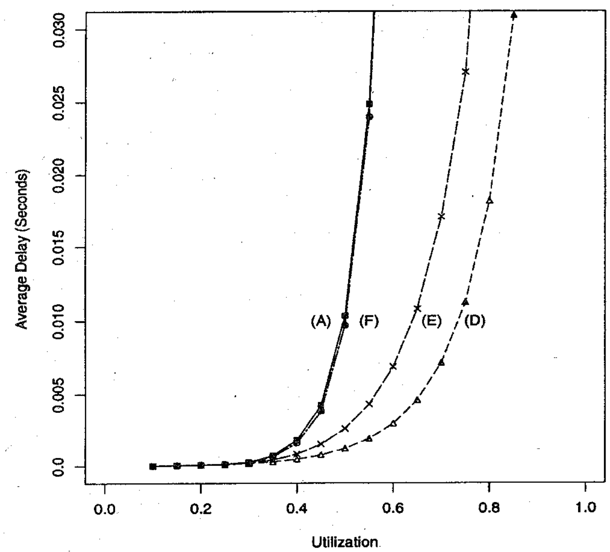

Fig. 4. Average delay (in seconds) versus utilization plot for original trace (A), trace with identical one-step correlations (D), externally shuffled trace (E), and internally shuffled trace (F).

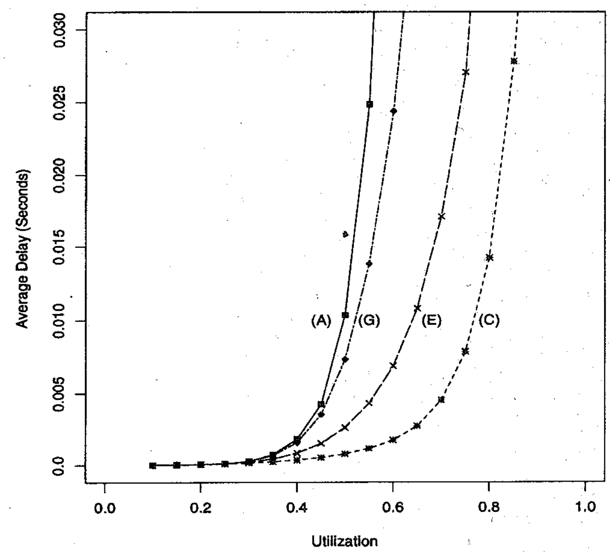

Fig. 5. Average delay (in seconds) versus utilization plot for original trace (A), fully shuffled (i.e., external shuffle with block size $m = 1$) trace (C), externally shuffled trace with block size $m = 25$ (E), and externally shuffled trace with block size $m = 500$ (G).

well as short term correlations up to lags of 25, the discrepancy with the delay curve (A) of the original trace is considerable. The input trace is representative of short-range dependent models that do not model correlations beyond a small number of lags, and the results indicate that such models may not capture significant aspects of the queueing performance. We next consider an "internal shuffle" (again with $m = 25$) of the sequence, in which the sequence *within* each block is randomized, while the *order* of the blocks is unchanged. This has the effect of destroying the *short-range* correlations in the data, while preserving the long-range correlations. This trace is then input to the simulation, and curve (F) in Fig. 4 is the resulting delay curve. Strikingly, this is almost exactly coincident with the delay characteristic of the original curve, even though in statistical terms, the two traces are very different. This demonstrates that LRD is not merely relevant for queueing performance; it is a dominant characteristic for determining several issues of traffic engineering concern, such as dimensioning of buffers, and determining usable capacity. Note that while other traffic characteristics (e.g., marginal distributions of time series packet counts) can be of similar significance for traffic engineering practice, they are not the focus of this study.

Obviously, an external shuffle with sufficiently large m does not affect the data significantly, because the correlations across such large blocks will be weak. On the other hand, one would expect an internal shuffle with sufficiently small m to also leave the data unaffected, because the traffic density is more correlated over short time intervals, and is thus not changed by local rearrangements of the interarrival times. Note that the fully random shuffle we considered earlier is the extreme case of an external shuffle with $m = 1$ or, equivalently, an internal shuffle with $m = N$. Fig. 5 depicts the average delay for the original data [curve (A)] and for three "external shuffles" using block sizes of $m = 1$ [curve (C)], $m = 25$ [curve (E)], and $m = 500$ [curve (G)], respectively. For the 30-min data set used here, a block size of $m = 500$ corresponds to block durations that are on an average about 1.5 s long, and ranging from a minimum of 0.5 s to a maximum of 4.2 s. It can be seen that even for $m = 500$, the delay characteristic is perceptibly different from that obtained for the original data, emphasizing the fact that correlations over extremely long time scales in the data have *measurable and practical* consequences. A second conclusion from these simulation experiments is that, for the problem at hand, the precise structure of the short-term correlations is relatively unimportant; thus a description in terms of arrival counts over a small time interval is adequate, even though it will not include characteristics of the traffic below this time scale. Equally, descriptions of the traffic process that are continuous in time (and assume a continuum of values) are valid.

C. Discussion

Similar results are obtained when other performance metrics are considered, such as delay percentiles or queue length distributions. Fig. 6 shows the queue length distributions obtained at a utilization of 50% in the simulation experiments described above. The behavior of the system at this utilization is of particular interest in engineering, because this is in the region of the "knee" of curve (A) in Figs. 2, 4, and 5. The logarithm of the complementary distribution $\log[P(V > x)]$ is plotted against x. For any finite Markov state model, this plot should be asymptotically linear (see for example [14]), indicating that the complementary distribution is of the form $\exp(-\eta x)$. The asymptotic form of the queue length distribution is important in itself in traffic engineering; for example, it forms the basis of equivalent bandwidth schemes as considered in [16] and [6], and is used in call admission controls. However, curve (A) obtained using the same Ethernet

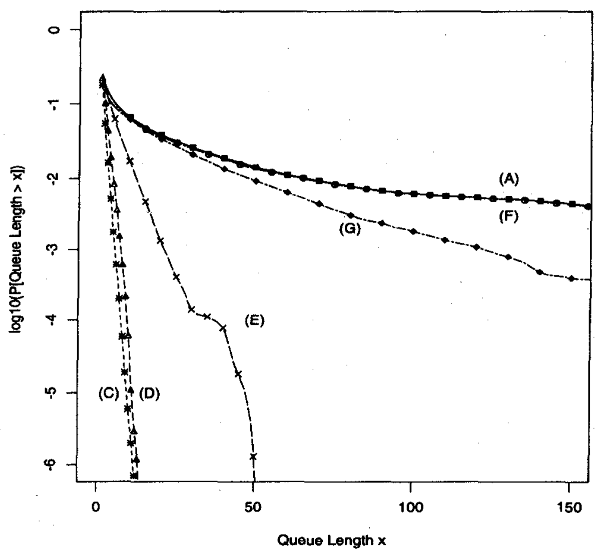

Fig. 6. Complementary queue length distributions for original trace (A), fully shuffled trace (C), trace with identical one-step correlations (D), externally shuffled trace with block size $m = 25$ (E), externally shuffled trace with block size $m = 500$ (G), and internally shuffled trace with block size $m = 25$ (F).

traffic trace as above is emphatically *not* linear. We will discuss the significance of this departure from linearity in more detail in the next section. Curves (C), (D), and (E) correspond to the cases of shuffled data (complete shuffle; one-step correlations preserved; external shuffle with $m = 25$, i.e., all long-term correlations removed). It can be seen that the queue length distributions decay much more rapidly than in the original data set. In fact, the curves obtained in these cases appear linear, indicating that the tails of the queue length distributions decay exponentially. In contrast, curve (F) which is obtained by an internal shuffle with $m = 25$, is almost coincident with the distribution obtained with the original trace. Curve (G), the queue length distribution obtained with an external shuffle using $m = 500$ also decays slowly, but can be distinguished from the plot obtained with the original trace. Thus, one can state that tails of queue length distributions obtained with actual data traces appear to be much heavier than is indicated by exponential decay, and that this is attributable to LRD. Beyond the range displayed in the plot, the queue length distributions (A), (F), and (G) are marked by the same sharp drop observed at the end of curve (E). This is an artifact and occurs in the region where simulation-based results become unreliable due to insufficient observations; simulations with longer traces move this region further out.

Experiments with counts are more in the spirit of the empirically observed LRD property of the counts processes considered in the traffic studies mentioned in Section II. Nevertheless, the variance-time plots in Fig. 7 demonstrate convincingly that the "internal shuffle" and "external shuffle" experiments on the interarrival times have indeed the desired effect of preserving or destroying the LRD property of the corresponding time series of counts.

It is therefore to be expected that similar results are also obtained if the data sets consist of time series of counts, and

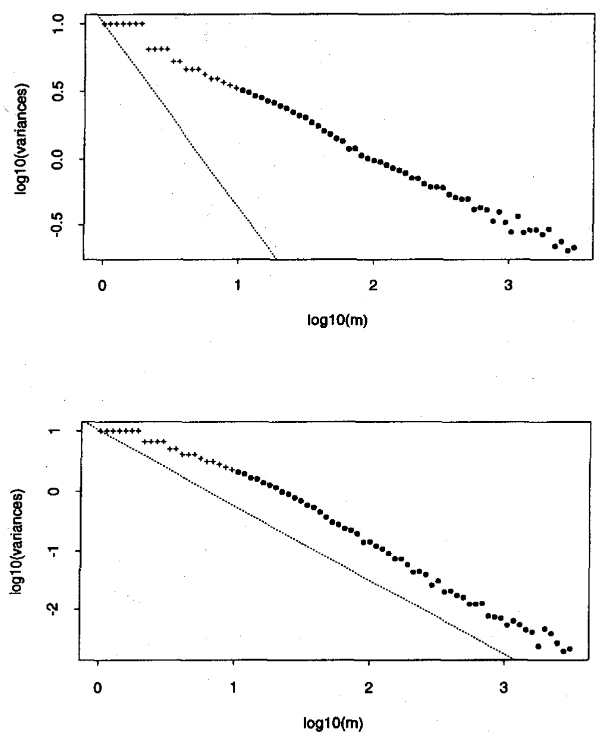

Fig. 7. Variance-time plots for "internally shuffled" (top) and "externally shuffled" (bottom) Ethernet data. The dotted reference lines correspond to lines with slopes -1.0.

the blocks are divided on the basis of equal time intervals, rather than constant numbers of packets. Fig. 8 shows results obtained when the data is represented by a time series of counts over 30 millisecond intervals with: the original trace [curve (A)]; complete shuffle [curve (C)]; external shuffle with $m = 10$, i.e., block duration is 300 milliseconds [curve (E)]; and internal shuffle with $m = 10$ [curve (F)]. The counts are randomly distributed over the 30 ms interval for the queueing simulations. Thus, performing the experiments with either blocks of interarrival times, or blocks of counts, yields similar results. The two approaches are equivalent when the high frequency structure (which is lost when the packets are aggregated over a time interval) is unimportant.

Our experimental results are also qualitatively consistent with the results obtained from the frequency-domain based approach considered in [27] and [28], where it is noted that low frequencies in the power spectra dominate queueing behavior. Recall that the LRD manifests itself as a sharp divergence in the low frequency region of the power spectrum.

IV. Parsimonious Traffic Modeling

One can conclude from the previous section that conventional short range models which do not incorporate LRD can be significantly in error when used in traffic engineering. This finding does not necessarily imply, however, that complicated and highly parameterized traffic models have to be employed to guarantee accurate and relevant solutions to practical engineering problems. In fact, it is well-known that parsimony in models can be achieved by abstracting out features that

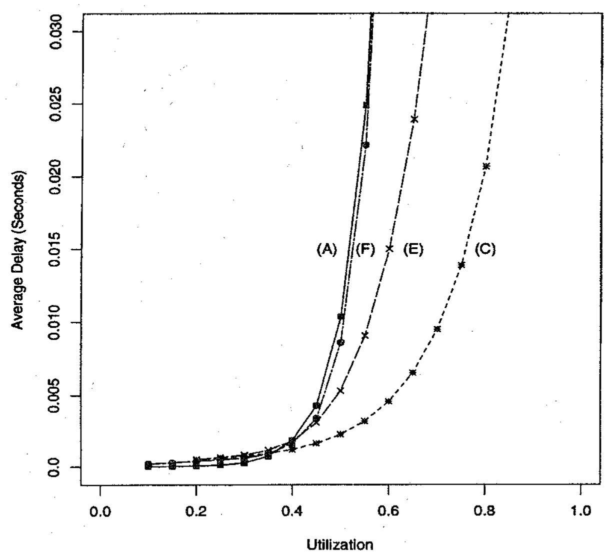

Fig. 8. Average delay (in seconds) versus utilization plot (using time series of counts as opposed to time series of interarrival times) for original trace (A), fully shuffled trace (C), externally shuffled trace with block size $m = 10$ (E), and internally shuffled trace with block size $m = 10$ (F).

do not contribute significantly to queueing performance. For example, the one-step correlations in interarrival times are striking, but for infinite buffer systems at least, they need not be modeled. In this context, it is essential to know what statistical aspects of network traffic can be ignored and when. We next consider this key problem in realizing parsimonious models, and i) give conditions under which second order statistical descriptions of traffic processes are sufficient for traffic engineering purposes, ii) discuss Fractional Brownian Motion models, which parsimoniously capture long-range dependence, and iii) show that the modeled features are robust with respect to traffic shaping.

A. Relevance of Second-Order Properties

It is reasonable to assume in a network environment that traffic flows are generated by multiple users who act independently; any two successive packets typically come from different and uncorrelated users. Thus, a characterization in terms of interarrival times is less useful for describing traffic aggregated from multiple users. (The importance of descriptions based on interarrival time distributions arises from their widespread use in classical GI/G/. theory, which is still the foundation of current traffic engineering practice. The practical utility of modeling the features of time series of interarrival times is, however, limited by the need to model operations such as splitting and merging of traffic streams.) On the other hand, as we shall see, it is precisely in the case where the number of users is large that a description of the traffic in terms of a continuous density variable, $\lambda(t)$, expressed as a function of time, works well. The results of Section III demonstrate that the shortest time scale fluctuations in the traffic density can be ignored. This validates the use of a theory based on the notion of a continuous traffic density, which we now construct.

For an aggregation of n independent users, the traffic density can be decomposed as the sum of n independent terms

$$\lambda(t) = \sum_{i=1}^{n} \lambda_i(t) \tag{7}$$

where λ_i and λ_j are assumed to be uncorrelated for different users i and j. If we also assume that n is large, a form of the central limit theorem is obtained; since with long-range correlations in the data this may not seem obvious, we outline the derivation here. The expectation value of any moment of the deviation from the mean of the traffic density

$$\delta\lambda(t) = \lambda(t) - \overline{\lambda} \tag{8}$$

is given by

$$\langle \delta\lambda(t_1)\delta\lambda(t_2) \cdots \delta\lambda(t_l)\rangle = \sum_{i_1=1}^{n} \cdots \sum_{i_l=1}^{n} \langle \delta\lambda_{i_1}(t_1) \cdots \delta\lambda_{i_l}(t_l)\rangle. \tag{9}$$

We *assume* that, in the *time* domain, all correlation functions of the form $\langle \delta\lambda_i(t_1) \cdots \delta\lambda_i(t_r)\rangle$ are finite. This is explicitly verified for $r = 2$ from the data (and is true for all r for a simple ON–OFF model that we discuss later).[4] If the number of distinct elements in the set $\{i\}$ is L, in the *large n limit* there are $O(n^L)$ such terms in the sum on the right hand side of (9), which is therefore dominated by the case when $L = l/2$ (the largest value possible). Thus *all* moments of the distribution for λ are given in terms of the second moment, which is the signature of a Gaussian distribution. If

$$\langle \delta\lambda_i(t_1)\delta\lambda_i(t_2)\rangle = \frac{1}{n} r(t_1 - t_2) \tag{10}$$

where r is the autocorrelation function defined earlier, then the probability distribution for the sum of a *large* number of such users is

$$P(\lambda) = \int [d\lambda] \exp\left\{-\tfrac{1}{2}\int dt_1\, dt_2\, r^{-1}(t_1 - t_2)\delta\lambda(t_1)\delta\lambda(t_2)\right\}. \tag{11}$$

It is simpler to express this equation in the Fourier domain, where we obtain

$$P(\lambda) = \int [d\lambda] \exp\left\{-\tfrac{1}{2}\int d\omega\, \delta\lambda(\omega) s^{-1}(\omega)\delta\lambda(-\omega)\right\} \tag{12}$$

where $s(\omega)$ is the spectral density. As has been seen earlier, for long-range dependent traffic processes, the low frequency behavior of $s(\omega)$ is

$$s(\omega) \sim |\omega|^{1-2H}, \quad \text{as} \quad \omega \to 0. \tag{13}$$

Notice that although $s(\omega)$ diverges for small ω (for $H > 1/2$), $r(t)$ is finite for all t: the maximum value is at $t = 0$, where $r(t = 0) = \int d\omega\, s(\omega)$, which is finite for $H < 1$.

[4] Strictly speaking, the correlation function at lag zero is unbounded if the traffic is viewed as a sequence of delta functions. The "data smoothing" we carry out in constructing the traffic densities removes this singularity.

Thus the assumption of finiteness of moments, necessary for second order properties to be adequate in the large n limit, is not contradicted. Also note that in this analysis the singular form of the spectrum at small frequencies is a property of each source, and is not a result of the multiplexing. Recent studies in [41] of Ethernet traffic traces suggest that single sources can be represented by ON–OFF models in which the sojourn time distributions of the ON and OFF states decay as power laws. Similar conclusions were made earlier ([33]) regarding the ON–OFF behavior of single ISDN D-channel sources. Thus the low frequency form of the spectrum in (13) arises due to the occasional sustained ON and/or OFF activity indicated by the power law sojourn time distributions.

Since in reality the number n of users on the system is finite, it is necessary to verify how far this $n \to \infty$ limit is applicable. In order to do this, we note that (12) is left unchanged under the transformation

$$\lambda(\omega) \to \lambda(\omega) \exp[i\theta(\omega)]$$

and

$$\lambda(-\omega) \to \lambda(-\omega) \exp[-i\theta(\omega)] \qquad (14)$$

with a random phase $\theta(\omega)$ that is separately chosen for each ω. (The joint condition on $\lambda(\pm\omega)$ preserves $\mathrm{Im}\,\lambda(t) = 0$.) We proceed in the following manner: from a five min subset of our 30 min experimentally measured data set, a smoothed density function is constructed, by dividing the total time into suitably chosen intervals, and measuring the (integer) number of arrivals in each such interval. The criterion for choosing the size of such an interval will be discussed shortly. From the sequence of integers X_k thus obtained, a discrete Fourier transform is constructed. This is now modified according to the prescription of (14), with $\theta(\omega)$ chosen independently for each $\pm\omega$, from a distribution uniform over $[0, 2\pi)$. The resulting altered Fourier transform is inverted, and rounded to the nearest integer to generate a sequence of integers, X'_k. To the interval k we assign X'_k arrivals, spaced randomly over the interval (other approaches to spacing the traffic over the interval, such as deterministic, uniform spacing, do not significantly change the outcome for such small intervals). For $X'_k < 0$, the interval is left empty. The criterion for choosing the size of the discretization interval is that it should be as large as possible, in order to minimize negative values of X'_k, without destroying the long-range variations in the density that are important. From the earlier tests on shuffling the data, we expect that an interval size of the order of ten times the mean interarrival time should be reasonable.

Figure 9 shows the effect of such a transformation on the queueing performance. As a function of the processor utilization, we plot the average delays for the original data [curve (A)] as well as the transformed data [curve (H)]. We see that there is a discernible difference between the two cases; this is to be expected, in view of the fact that the transformed data set is equivalent to the original only in the $n \to \infty$ limit. Such discrepancies (within 5% on the utilization axis in this case) may be acceptable in practical traffic engineering, though additional safety margins have to be incorporated into capacity estimates, given that this error may be on the optimistic side. These discrepancies arise because the transformed data differs from the original data set in a number of ways, including in distribution, higher order statistics, as well as the random interpolation scheme used to assign arrivals within an interval. The effects of the distribution and higher order statistics are expected to decrease as the number of independent users becomes larger. Indeed, Fig. 10 shows the average delay versus utilization curve obtained by aggregating 10 traces, derived from a five min Ethernet trace using different initial offsets chosen randomly. For a carefully chosen set of initial offsets, this has the effect of increasing n, and as can be seen, the discrepancy is reduced. This provides some empirical justification for representing the arrival process by first and second order statistics.

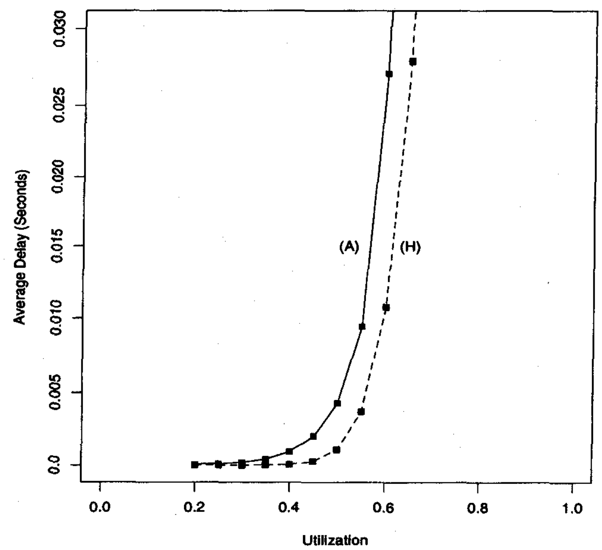

Fig. 9. Average delay (in seconds) versus utilization plotfor original trace (A) and transformed trace (H).

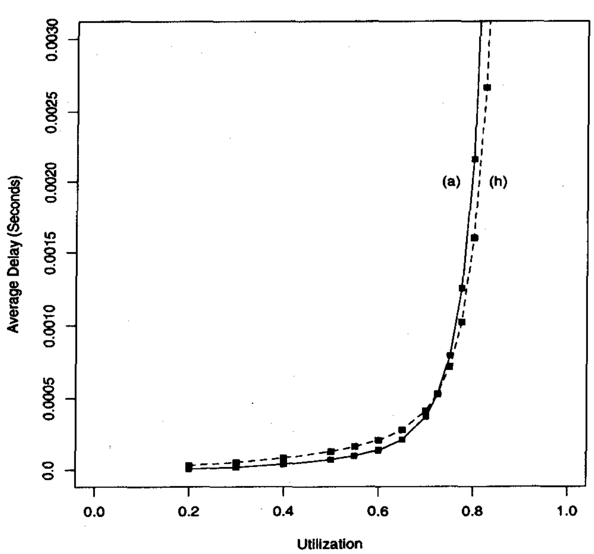

Fig. 10. Average delay (in seconds) versus utilization plot for 10 multiplexed original traces (a) and the corresponding transformed trace (h).

B. Fractional Brownian Motion Models

In principle, one can also use a high-order Markov model to approximate the power spectrum by using a rational function to match it, as suggested in [28]. As we have emphasized before, it is essential in practice for the description to be parsimonious. Parsimony is achieved by using the power law representation in (13) of the spectrum, which is the form indicated by all the data sets we have analyzed. Thus traffic can be characterized by the Hurst parameter H, a magnitude term representing the strength of the fluctuations $\delta\lambda$, and the average rate $\overline{\lambda}$.

This is the approach effectively adopted by Norros [34], using a Fractional Brownian Motion (FBM) model of the arrival process. FBM can be viewed as an extension of the standard Brownian Motion models that have been used with some success in heavy traffic analysis. In standard Brownian Motion models, the cumulative arrival process $A(t)$ is modeled by random fluctuations about a mean rate

$$A(t) = nt + \sqrt{an}Z(t) \tag{15}$$

where the process $Z(t)$ has independent Gaussian increments, and a is a *peakedness* term that describes the magnitude of fluctuations. (We are using units of time in which $\overline{\lambda}(t) = n$, so that each source has, on the average, $\overline{\lambda}_i = 1$.) In FBM, the increments of $Z(t)$ are taken to be long-range dependent. $Z(t)$ satisfies the conditions $Z(t=0) = 0$ and $\langle Z(t)\rangle = 0$, and

$$\langle Z(t_1)Z(t_2)\rangle = \tfrac{1}{2}(t_1^{2H} + t_2^{2H} - |t_1 - t_2|^{2H}). \tag{16}$$

The distribution of $Z(\alpha t)$ is identical to that of $\alpha^H Z(t)$. From (16), one can obtain the correlation structure of the increments of $Z(t)$: for $t_1 < t_2 < t_3 < t_4$, with $t_1 \approx t_2$ and $t_3 \approx t_4$, we have

$$\begin{aligned}&\langle [Z(t_4) - Z(t_3)][Z(t_2) - Z(t_1)]\rangle = \\ &H(2H-1)(t_4 - t_2)^{2H-2}(t_2 - t_1)(t_4 - t_3)\end{aligned} \tag{17}$$

which decays as a power-law in terms of $t_4 - t_2$. By setting the Hurst parameter H to 0.5, (16) reduces to the form for standard Brownian motion, and the right hand side of (17) is zero. Note that FBM is an exactly self-similar model, with the above scaling relation applicable for all choices of $\{t_1, t_2\}$. In practice, data traffic shows this scaling behavior over a wide range of time scales, though there are lower cutoffs (for example, about 10 milliseconds for Ethernet traffic) below which short range correlations dominate. The FBM model is expected to be valid under the following conditions: i) the time scales of interest in the queueing processes coincide with the scaling region, ii) the traffic is aggregated from a large number of independent users, and iii) the effect of flow controls on any one user is negligible. In such a setting, our analysis has shown that the complex short-range correlations can be ignored from a traffic engineering perspective. The long-range correlations, or equivalently, the low-frequency structure of the power spectrum, which is relevant for many aspects of practical traffic engineering, can be parsimoniously modeled using the three parameter FBM model.

Norros has derived several results for the queueing behavior obtained by driving a deterministic service time queue with an FBM process, including the following asymptotic lower bound for the probability P(V > x) that the queue-length V exceeds x

$$P(V > x) \sim \exp[-cx^{2-2H}] \tag{18}$$

with

$$c = \frac{n^{2H-1}(\frac{1-\rho}{\rho})^{2H}}{2a} \cdot \left[\left(\frac{1-H}{H}\right)^{H} + \left(\frac{H}{1-H}\right)^{1-H}\right]^2. \tag{19}$$

Here ρ is the utilization, or ratio of $\overline{\lambda}$ to the processor speed. It is expected that the full expression for $P(V > x)$ will differ from (18) by at most a power-law prefactor. In fact, Duffield and O'Connell [4] have shown that the Norros bound is asymptotically tight.

Figure 11 shows a plot of log $P(V > x)$ versus x, using trace-driven simulations [i.e., curve (A) from Fig. 6]. Here we have considered a utilization of 0.5, corresponding to the knee of curve (A) in Fig. 2. The dashed curve (I) is the asymptotic form from (18) and (19). The values of H and a used to obtain the fit are within the confidence limits of the parameters estimated from Fig. 1. The two curves should be asymptotically the same; although for small x we expect in general corrections to the asymptotic form, it seems to work surprisingly well even in this regime. At very large x, beyond $x = 100$, the measured queue length distribution falls off *faster* than the dashed curve, due to limitations of the length of the simulation. In any finite set of data it is inevitable that the queue length distribution will eventually fall-off faster than the predicted form for large x. Note that it is always possible to find a range of queue lengths where an exponential form (i.e., linear on the semi-log plot), with an appropriately chosen small decay constant, provides a good fit to the queue length distribution. While the effects of the long-range dependence, whose importance we have demonstrated in this paper (as in Fig. 6), can thus be modeled over a limited range of x values by exponential decays with small decay constants, with longer data traces, discrepancies become increasingly apparent (in addition, the range and decay constant depend on the length of the data trace). Due to the substantial weight at large x in the distribution, these discrepancies will have practical implications.

Heavy-tailed queueing behavior is also reported in [9], based on numerical and analytical studies of models that formulate ON–OFF source and queueing behavior in terms of chaotic maps. Thus, there is considerable support to the simulation results indicating that the tails of the queue length distributions decay much more slowly than the exponential rate predicted by Markov models.

This property can be expected to have considerable impact in engineering. For example, buffers sized on the basis of conventional model may result in underprovisioning. Early experiences in data services over ATM indicate heavier than expected losses [3]. A combination of vastly increased buffer sizes and reduced utilizations may be necessary to achieve acceptable loss rates. The heavy-tailed nature of the queue

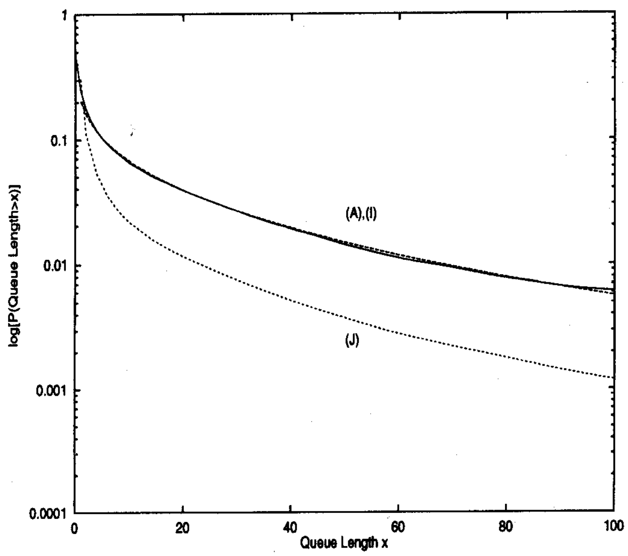

Fig. 11. Complementary queue length distributions for original trace (A), asymptotic form given in (18) (I), and finite buffer simulation with original trace (J).

length distributions may introduce additional errors in admission strategies that are based on equivalent bandwidth schemes.

Hwang and Li [17] report a phenomenon called "buffer ineffectiveness," in which increasing buffer sizes beyond a certain value results in only a slight decrease in loss rates. This is, in fact, a consequence of the "stretched exponential" nature of (18) [34].

The FBM model also provides insights into the sharp rise in average waiting times observed with actual traffic traces [curve (A) in Figs. 2, 4, and 5]. Using (18) to estimate the average queue length $\overline{x}$, it is readily shown that $\overline{x}$ diverges as $(1-\rho)^{-H/(1-H)}$ as ρ approaches one. When $H = 0.5$, $\overline{x}$ behaves as $(1-\rho)^{-1}$, which is familiar from standard queueing models with constant or exponential service times. For $H = 0.8$, $\overline{x}$ increases much more rapidly [as $(1-\rho)^{-4}$] which explains the sharp rise in delays observed in the simulations. Once again, this behavior is unprecedented in short-range dependent models with constant or exponential service times, and shows that the use of standard queueing approximations to set network operating points may lead to gross performance problems.

C. Effect of Traffic Shaping

Network elements can interact with traffic flows in a variety of passive and active ways. The characteristics of traffic flows can be shaped by the buffering action of a queue (which is an example of passive interactions) or in a more active fashion by rate controls that limit each source to a prescribed set of peak and mean rates. In both cases, excess traffic is buffered and/or dropped. It has been suggested that "traffic shaping" by passive or active means can reduce or eliminate LRD in network traffic. We show in the following that shaping traffic sources to the extent of eliminating LRD will require substantial buffering at the edge of the network, with attendant prohibitive delay penalties; also see [28] for an analysis on the limitations of rate controls in shaping low frequencies.

Let $V(t)$ be the length of the queue or backlog at a passive network element or a shaper that enforces a prescribed set of peak and mean rates, at a given time instant t. Further, let $\lambda_I(t)$ and $\lambda_O(t)$ be the input and output traffic rates, respectively. Then it follows from flow balance that

$$\lambda_O(t) = \lambda_I(t) - \frac{dV(t)}{dt}. \quad (20)$$

Transforming into the Fourier domain and taking the square of the magnitude of the result, one obtains

$$\begin{aligned}\langle \lambda_O(\omega)\lambda_O(-\omega)\rangle &= \langle \lambda_I(\omega)\lambda_I(-\omega)\rangle + \omega^2\langle V(\omega)V(-\omega)\rangle \\ &\quad - 2\,\mathrm{Im}\left[\omega\langle V(\omega)\lambda_I(-\omega)\rangle\right] \quad (21)\end{aligned}$$

where we have used the fact that $V(t)$ and $\lambda_I(t)$ are real functions and $\mathrm{Im}[z]$ denotes the imaginary part of z.

Examining the third term we have

$$\begin{aligned}-\mathrm{Abs}\left[\omega\langle V(\omega)\lambda_I(-\omega)\rangle\right] &< \mathrm{Im}\left[\omega\langle V(\omega)\lambda_I(-\omega)\rangle\right] \\ &< \mathrm{Abs}\left[\omega\langle V(\omega)\lambda_I(-\omega)\rangle\right] \quad (22)\end{aligned}$$

where Abs denotes absolute value. By Schwarz's inequality, the RHS of this inequality is bounded above by

$$\begin{aligned}&|\omega|\langle \mathrm{Abs}[V(\omega)]\mathrm{Abs}[\lambda_I(-\omega)]\rangle \\ &< |\omega|\sqrt{\langle \mathrm{Abs}[V(\omega)]^2\rangle\langle \mathrm{Abs}[\lambda_I(-\omega)]^2\rangle} \\ &= |\omega|\sqrt{\langle V(\omega)V(-\omega)\rangle}\sqrt{\langle \lambda_I(\omega)\lambda_I(-\omega)\rangle}. \quad (23)\end{aligned}$$

Given that the probability distribution of V is dominated by a stretched exponential (18), $\langle V(t)V(t)\rangle$ is finite. In particular, if $\langle V(\omega)V(-\omega)\rangle$ diverges at low frequency, it must diverge slower than $1/|\omega|$. Thus, the expression in (21) is dominated at low frequencies by the divergence of the first term, and the output of the queue or traffic shaper has essentially the same low frequency structure as the input traffic. Conversely, if we require the output from the shaper to have less singular long-range dependence than the input, we see that this can only be achieved if mean square queue length, $\langle V(t)V(t)\rangle$, diverges.

Note that this result does not depend on the *mechanism* by which the shaping is achieved, which will only affect the functional dependence of $\lambda_O(t)$ on $\lambda_I(t)$ and $V(t)$ (and, perhaps, other parameters). All it requires is that flow balance should be satisfied. For a finite buffer, finiteness of $\langle V(t)V(t)\rangle$ is achieved at the expense of dropping some of the packets; by arguments similar to those above it can be shown that the dropped data must have the same long-range dependence as the input stream, which we expect to lead to unacceptably high losses.

Two conclusions can be drawn from this analysis. First, the FBM description is robust and is not qualitatively altered as the traffic flows through the network under normal operation. Secondly, traffic shaping will not eliminate LRD. At best, controlling LRD by traffic shaping will have the effect of transferring buffer requirements from within the network to the edge of the network, without an improvement in performance. In fact, given that there is potential for considerable statistical multiplexing within the network, the net effect may be to vastly increase buffer requirements at the edge of the network. We stress that this is not meant to detract from the importance of network controls such as traffic shaping, which are essential to monitor *atypically heavy* sources, and to prevent them from monopolizing network resources. Note that LRD is a typical characteristic of *generic* traffic sources.

V. Additional Issues

In the following, we will discuss miscellaneous problems, from engineering implications of LRD to the impact of various factors on the conclusions we have drawn so far. A detailed discussion of these issues is beyond the scope of this paper, and we mention them in part to stimulate a more in-depth investigation of these issues.

A. Effect of Finite Buffers

While all the experiments described so far in this paper assume infinite buffer systems, in Fig. 11 we also present the results of a finite buffer simulation (same Ethernet trace as above, 50% utilization). Clearly, the functional form (i.e., Weibull) of the queue length distributions corresponding to the finite [curve (J)] and infinite [curve (A)] buffer systems is the same, with an initial offset that results in the well-known dominance of the infinite buffer queue length distribution over its finite buffer counterpart. Intuitively, the effect of finite buffers is in some sense to limit the time scales of queueing interest in the correlation structure. Thus, it is to be expected that queues with finite buffers will not show the sharp degradation in delay performance observed with infinite buffer systems, but analytical results in support of this argument are presently nonexistent. This comes at the expense, however, of increased packet losses. In this context, one can keep packet losses acceptable by operating network resources at such low utilizations that the queueing system never enters the regime in which correlations impact performance, but this may not be economical. Because of the way transport protocols such as TCP/IP respond to packet losses, many data applications are tolerant of delays, but highly sensitive to losses (see, for example, [32]). Current ATM switches designed to support data traffic are incorporating much larger buffers than switches used in early trials. Thus, it is expected that the infinite buffer regime is of interest in practice.

B. Physical Basis for LRD

While there are clearly long-range correlations in the traffic, with demonstrated impacts on performance and engineering, it is also important to understand *how* they arise in data traffic. Recent analysis in [41] of single sources from the Ethernet data sets provide insights into this issue. In terms of the familiar ON–OFF source abstraction, real sources differ from existing models in that the sojourn times in these states are characterized by power-law or Pareto distributions. Similar findings of highly variable periods of activity and idleness at the application level form the basis for the self-similar features observed in the WAN traffic traces collected at Berkeley [35], [36]. It is the (occasional) sustained inactivity, or activity implied by power law distributions, that manifests itself as the $1/f$ noise spectra or LRD. Extending the results in [30] and [37], it can be shown (see [41]) that aggregating a large number of such independent ON–OFF sources [with ON–OFF periods characterized by $P(T > t) \sim t^{-\alpha}$] will in fact generate FBM of index $H = (3 - \alpha)/2$. (Thus for $1 > H > 1/2$, we have $2 > \alpha > 1$, so that the distribution has a *finite mean,* but *infinite variance.*) These models are different from the singular renewal models discussed in Section III-A, where the aggregated traffic is viewed as an ON–OFF source (with each ON period consisting of only one data packet). By shuffling the interarrival times, we saw that the queueing implications of such models for the aggregate traffic would be relatively minor. In addition, since the aggregated traffic is modeled rather than individual sources, it would be very difficult to obtain the dominance of second-order properties that exists in real traffic. Note that, with a traffic density λ_1 in the ON state and λ_0 in the OFF state ($\lambda_1 > \lambda_0$), all cumulants of the traffic density, with different time spacings, are finite, as mentioned in Section IV-A. When the ON–OFF distributions of different sources are distinct, the sources with heavier tail behavior may be expected to dominate. Though there are no formal results in this case, an analysis similar to that in the next subsection should be possible.

C. Statistical Multiplexing Gains

The Hurst parameter H is preserved under multiplexing of identical sources. When heterogeneous sources, with identical mean traffic rate $\overline{\lambda}_i$ and "burstiness" a_i, but with a distribution of possible values of H, are multiplexed, the largest H will dominate. The interpretation of the Hurst parameter as a measure of burstiness, along with this nondecreasing property, have led some to conclude that multiplexing does not i) reduce the burstiness of traffic, and ii) as such multiplexing gains are not feasible with long-range dependent traffic. On the contrary, the FBM model does predict significant multiplexing gains when a large number of independent sources are multiplexed. This is because burstiness is characterized not just by the correlations in the fluctuations parameterized by H, but also by their relative magnitude characterized by $\sqrt{a/n}$. When n independent sources are multiplexed, the relative magnitude is reduced by $\sqrt{n}$, and therefore environments in which the capacity of a server is significantly greater than that of the magnitude of fluctuations of a single source should realize substantial multiplexing gains. In particular, the lower bound on the queue length distribution, (18), explicitly decreases with increasing n (provided that the processor speed is increased proportionately, so that ρ is fixed).

A conservative approach to the problem of which H to choose when the sources are heterogeneous is to use the largest H value of all the sources. Given that the form for the queue-length distribution obtained in (18) depends sensitively on the Hurst parameter, H, this may prove to be too conservative. While taking the largest H value is indeed correct for the strictly asymptotic behavior, in practice a small number of users with unusually high values of H will only affect the distribution for extremely large x. With n_1 users with $H = H_1$ and $n_2 \ll n_1$ users with $H = H_2 > H_1$, an estimate of the crossover from one form to the other can be obtained

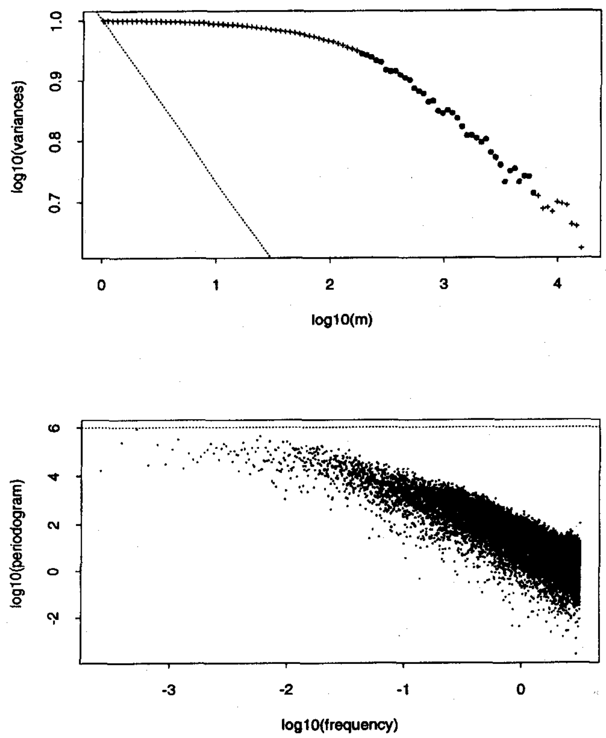

Fig. 12. Variance-time plot (top) and periodogram plot (bottom)for a 1/2-hour VBR video traffic trace. The dotted reference lines in the variance-time and periodogram plots have slopes −1.0 and 0.0, respectively.

by considering either $n_1 + n_2$ users with $H = H_1$, or only n_2 users with $H = H_2$, and obtaining from (18) the value of x for which the two cases yield the same $P(V > x)$. This is clearly an issue of practical importance, and there is considerable scope for further work.

D. VBR Video Traffic

As demonstrated in [1], VBR video traces also demonstrate LRD, though there are some fundamental differences in the correlation structures of data and video traffic. A full discussion of these differences and their impact is beyond the scope of this paper, and we mention them in passing. Fig. 12 shows the variance-time plot and periodogram plot of a VBR video trace, representing a video-conferencing scene. It can be seen that for a range of intermediate time scales, the variance-time plot shows very little decay before entering the asymptotic regime that shows the slowly decaying variances indicative of LRD. This feature of a variance-time plot suggests the presence of strong short-term correlations in the data, which makes VBR traffic only "asymptotically" self-similar (see [2]). In contrast, data traffic shows essentially the same structure of the variance-time plot over all but the shortest time scales, and can be modeled over time scales of engineering interest by exactly self-similar processes.

The strong short-term correlations in VBR video data are modeled by means of both Markovian and non-Markovian models (see for example [18], which reports that this satisfactorily reproduces queueing behavior of the original traffic traces). The impacts of LRD in this context are still under investigation. Parsimony is a primary consideration here as well, motivating the study of models that capture both the short-range as well as long-range correlation structures. Possible alternatives are Fractional ARIMA models, and extensions of the FBM model that incorporate lag dependent correlations. In any event, the effect of the additional strong short term correlations can be expected to *aggravate* the heavy queueing behavior predicted by LRD models. This can be seen specifically using extended FBM models for a large number of multiplexed video sources: if we approximate the variance-time plot in Fig. 12 as

$$\operatorname{var}[X^{(m)}] = \begin{cases} a_0, & m < m_0 \\ a_0 \left(\dfrac{m_0}{m}\right)^{2-2H}, & m > m_0 \end{cases}$$

we obtain the same asymptotic form for the queue length distribution as with FBM models with no cutoff, with the parameter c in (19) inversely proportional to $a_0 m_0^{2-2H}$. Thus, a large flat section in Fig. 12 implies (for the same a_0) a small prefactor in the exponent of (18), leading to a heavier weight in the tail of the distribution. In fact, it is in principle possible to obtain the entire queue length distribution for this multiple source case, by performing numerical simulations with the *full* form of the variance-time plot, using an (appropriately modified) FBM model.

E. Estimating Model Parameters

As discussed earlier, one of the basic problems limiting the application of many theoretical traffic models in practice is the difficulty of assigning model parameters, especially when parsimony has not been a major concern at the modeling stage. Beyond statistical considerations, parsimonious models are clearly preferred over highly parameterized models when faced with the task of assigning model parameters in practice. For example, the FBM description of packet traffic considered in Section IV requires only three parameters: the mean rate m, the peakedness parameter a, and the Hurst parameter H, all of which can be estimated from high time-resolution traffic measurements. Regarding the availability of estimation procedures (especially for a and H), however, some words of caution are in place. To illustrate, there exist numerous methods for estimating H from a time series of packet counts (see for example, [40], and the references therein); while the statistical properties of some estimates (e.g., Whittle's estimate) are well understood within a Gaussian framework, much less is known when the Gaussian assumption is violated. Other estimates for H (e.g., based on variance-time analysis) are mostly heuristic in nature, and their distributional properties and statistical features are typically unknown. Similar comments hold with respect to the estimation of the peakedness parameter a: while it can be directly obtained from a variance-time plot its statistical properties are largely unknown, which makes this approach again mostly heuristic. Clearly, there is considerable scope for improvements and a pressing need for innovative statistical approaches.

While high time-resolution traffic traces were used in this paper, it is prohibitive in practice to collect such measurements on an operational basis. In principle, the empirically observed self-similarity of our measured traffic can be exploited to reduce such measurement overhead by noting that the value of H for a time series of counts over coarse time scales is the same as that obtained from high resolution traces (the value of a can also be inferred). Currently, some switching systems have the capability to report traffic counts over one second intervals during an engineering period; in principle, H can be estimated from such counts. For a more detailed description of these traffic measurement problems, see [11].

VI. Conclusions

Recent studies involving large sets of actual traffic measurements from working packet networks (e.g., LAN's, ISDN's, and CCSN's) have illustrated that these high time-resolution data sets are consistent with the assumption of long-range dependence in packet traffic. In this paper, we use these traffic measurements to demonstrate (via trace-driven simulation experiments) that, beyond its omnipresence and statistical significance in our measured data, long-range dependence is a traffic characteristic that i) has measurable and practical impact on queueing behavior, ii) is of crucial importance for a number of packet traffic engineering problems (e.g., buffer sizing, admission control, and rate control), and iii) if ignored, typically results in overly optimistic performance predictions and inadequate network resource allocations.

From a modeling viewpoint, our work shows that finding long-range dependence in traffic measurements from modern high speed networks does not necessitate complicated and highly parameterized traffic models. In fact, we provide conditions under which compact and parsimonious Fractional Brownian motion-based models can be expected to describe packet traffic in tomorrow's networks realistically and predict their performance accurately. In the presence of traffic measurements from modern high speed networks, we strongly argue in favor of and illustrate parsimonious modeling approaches that focus on essential traffic characteristics, i.e., properties that i) have a dominant and practical impact on network design and performance, ii) have a meaningful physical basis in the network context, and iii) have been uncovered by detailed statistical analyses of measured network traffic and can be efficiently quantified (that is, estimated from the data). A prime example of such a traffic characteristic is long-range dependence which we have demonstrated in this paper to be of practical relevance. An area of considerable future research is assessing its full impact on network design, management, and operations.

Acknowledgment

The authors thank I. Norros, of VTT Finland and N. Duffield for making available the preprints of their work. They also wish to thank their colleagues M. Aizenman, K. Krishnan, S. Majumdar, and A. Neidhardt for many useful discussions. In particular, they acknowledge the many helpful comments by two anonymous referees and by Vern Paxson on an earlier draft of this paper. O. Narayan acknowledges the support of the Society of Fellows at Harvard University and also the hospitality of AT&T Bell Laboratories.

References

[1] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger, "Long-range dependence in variable-bit-rate video traffic," *IEEE Trans. Commun.*, vol. 43, pp. 1566–1579, 1995.

[2] D. R. Cox, "Long-range dependence: A review," in *Statistics: An Appraisal,* H. A. David and H. T. David, Eds. Ames, Iowa: Iowa State Univ. Press, 1984, pp. 55–74.

[3] *Communications Week,* May 16, 1994.

[4] N. G. Duffield and N. O'Connell, "Large deviation and overflow probabilities for the general single-server queue, with applications," preprint, Dublin Institute for Advanced Studies, 1993.

[5] D. E. Duffy, A. A. McIntosh, M. Rosenstein, and W. Willinger, "Statistical analysis of CCSN/SS7 traffic data from working subnetworks," *IEEE J. Select. Areas Commun.,* vol. 12, no. 3, pp. 544–551, 1994.

[6] A. I. Elwalid and D. Mitra, "Effective bandwidth of general markovian traffic sources and admission control of high speed networks," *IEEE/ACM Trans. Networking,* vol. 1, no. 3, pp. 329–343, 1993.

[7] A. Erramilli and W. Willinger, "Fractal properties in packet traffic measurements," in *Proc. St. Petersburg Regional ITC Seminar,* St. Petersburg, Russia, 1993, pp. 144–158.

[8] A. Erramilli, J. Gordon, and W. Willinger, "Applications of fractals in engineering for realistic traffic processes," in *The Fundamental Role of Teletraffic in the Evolution of Telecommunications Networks, (Proc. ITC-14, Antibes Juan-les-Pins, France, June 1994),* J. Labetoulle and J. W. Roberts, Eds. Amsterdam: Elsevier, 1994, pp. 35–44.

[9] A. Erramilli, R. P. Singh, and P. Pruthi, "An application of deterministic chaotic maps to model packet traffic," *Queueing Syst.,*, vol. 20, pp. 171–206, 1995.

[10] ——, "Chaotic maps as models of packet traffic," in *The Fundamental Role of Teletraffic in the Evolution of Telecommunications Networks (Proc. ITC-14, Antibes Juan-les-Pins, France, June 1994),* J. Labetoulle and J. W. Roberts, Eds. Amsterdam: Elsevier, 1994, pp. 329–338.

[11] A. Erramilli and J. Wang, "Monitoring packet traffic levels," in *Proc. IEEE Globecom'94,* 1994.

[12] H. J. Fowler and W. E. Leland, "Local area network traffic characteristic, with implications for broadband network congestion management," *IEEE J. Select. Areas Commun.,* vol. 9, no. 7, pp. 1139–1149, 1991.

[13] M. W. Garrett and W. Willinger, "Analysis, modeling and generation of self-similar vbr video traffic," in *Proc. ACM Sigcomm'94,* London, UK, 1994, pp. 269–280.

[14] P. W. Glynn and W. Whitt, "Logarithmic asymptotics for steady-state tail probabilities in a single-server queue," *J. Appl. Prob.,* vol. 31, pp. 131–156, 1994.

[15] C. W. J. Granger and R. Joyeux, "An introduction to long-memory time series models and fractional differencing," *J. Time Series Anal.,* vol. 1, pp. 15–29, 1980.

[16] R. Guerin, H. Ahmadi, and M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high speed networks," *IEEE J. Select. Areas Commun.,* vol. 9, pp. 968–981, 1991.

[17] C. L. Hwang and S. Q. Li, "On input state space reduction and buffer noneffective region," in *Proc. IEEE Infocom'94,* 1994, pp. 1018–1028.

[18] D. P. Heyman, A. Tabatabai, and T. V. Lakshman, "Statistical analysis and simulation study of video teleconference traffic in ATM networks," *IEEE Trans. Circuits Syst.*, vol. 2, pp. 49–59, 1992.

[19] J. R. M. Hosking, "Fractional differencing," *Biometrika,* vol. 68, pp. 165–176, 1981.

[20] C. Huang, M. Devetsikiotis, I. Lambadaris, and A. R. Kaye, "Self-similar modeling of variable bit rate compressed video: A unified approach," in *Proc. ACM Sigcomm'95,* Cambridge, MA, 1995, pp. 114–125.

[21] W. H. Jefferys and J. O. Berger, "Ockham's razor and Bayesian analysis," *Amer. Scientist,* vol. 80, pp. 64–72, 1992.

[22] D. L. Jagerman and B. Melamed, "The transition and autocorrelation structure of tes processes part I: General theory," *Stochastic Models,* vol. 8, pp. 193–219, 1992.

[23] S. Klivansky, A. Mukherjee, and C. Song, "Factors contributing to self-similarity over NSFNet," preprint, Georgia Institute of Technology, 1994.

[24] W. E. Leland and D. V. Wilson, "High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection," in *Proc. Infocom'91,* Bal Harbour, FL, 1991, pp. 1360–1366.

[25] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic," in *Proc. ACM Sigcomm'93,* San Francisco, CA, 1993, pp. 183–193.

[26] ——, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Trans. Networking,* vol. 2, no. 1, pp. 1–15, 1994.

[27] S. Q. Li and C. L. Hwang, "Queue response to input correlation functions: discrete spectral analysis," *IEEE/ACM Trans. Networking,* vol. 1, pp. 522–533, 1993.
[28] ——, "Queue response to input correlation functions: Continuous spectral analysis," *IEEE/ACM Trans. Networking,* vol. 1, pp. 678–692, 1993.
[29] B. B. Mandelbrot, "Self-similar error clusters in communication systems and the concept of conditional stationarity," *IEEE Trans. Communications Technol. COM-13,* pp. 71–90, 1965.
[30] ——, "Long-run linearity, locally gaussian processes, h-spectra and infinite variances," *Int. Econom. Review,* vol. 10, pp. 82–113, 1969.
[31] B. B. Mandelbrot and J. W. Van Ness, "Fractional brownian motions, fractional noises and applications," *SIAM Review,* vol. 10, pp. 422–437, 1968.
[32] G. Meempat, "Interactions of reliable stream data transport protocols with rate control algorithms," in *Proc. 3rd Int. Conf. Comput. Commun. Networks,* 1994, pp. 159–164.
[33] K. Meier-Hellstern, P. E. Wirth, Y.-L. Yan, and D. A. Hoeflin, "Traffic models for isdn data users: Office automation application," in *Teletraffic and Data Traffic in a Period of Change (Proc. 13th ITC, Copenhagen, 1991),* A. Jensen and V. B. Iversen, Eds., North Holland, 1991, pp. 167–172.
[34] I. Norros, "A storage model with self-similar input," *Queueing Syst.,* vol. 16, pp. 387–396, 1994.
[35] V. Paxson and S. Floyd, "Wide-area traffic: The failure of poisson modeling," in *Proc. of the ACM Sigcomm'94,* London, UK, 1994, pp. 257–268.
[36] ——, "Wide area traffic: The failure of poisson modeling," *IEEE/ACM Trans. Networking,* vol. 3, pp. 226–244, 1995.
[37] M. S. Taqqu and J. B. Levy, "Using renewal processes to generate long-range dependence and high variability," in *Dependence in Probability and Statistics,* E. Eberlein and M. S. Taqqu, Eds. Boston, MA: Birkhauser, 1986.
[38] D. Veitch, "Novel methods of broadband traffic," in *7th Australian Teletraffic Research Seminar,* Murray River, Australia, 1992.
[39] W. Willinger, "Traffic modeling for high speed networks: Theory versus practice," in *Stochastic Networks,* F. P. Kelly and R. J. Williams, Eds. New York: Springer-Verlag, The IMA Volumes in Mathematics and its Applications, 1995, vol. 71, pp. 395–409.
[40] W. Willinger, M. S. Taqqu, W. E. Leland, and D. V. Wilson, "Self-similarity in high speed packet traffic: Analysis and modeling of ethernet traffic measurements," *Stat. Sci.,*, vol. 10, pp. 67–85, 1995.
[41] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: Statistical analysis of ethernet LAN traffic at the source level," in *Proc. ACM Sigcomm'95,* Cambridge, MA, 1995, pp. 100–113.

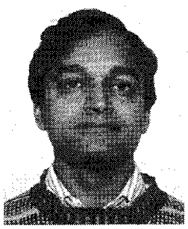

Ashok Erramilli (M'83) obtained the B.Sc. degree in physics from the University of Poona, India, the B.E. degree in electrical communications engineering from the Indian Institute of Science, Bangalore, and the M.S. and Ph.D. degrees in electrical engineering from Yale University.

He is Director of the Network Design and Traffic Analysis Research Group in Bellcore. His current research interests are in the analysis and control of fractal network traffic flows, broadband traffic management methods, and traffic engineering economics.

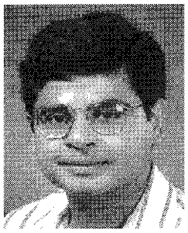

Onuttom Narayan received the Ph.D. degree in physics from Princeton University, Princeton, NJ.

He is on the Physics faculty of the University of California at Santa Cruz. He is currently working on the behavior of spatially extended disordered systems that are out of equilibrium.

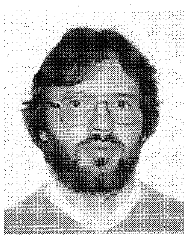

Walter Willinger (M'95) received the Diplom (Dipl. Math.) from the ETH Zürich, Switzerland, and the M.S. and Ph.D. degrees from the School of ORIE, Cornell University, Ithaca, NY.

He is a Member of Technical Staff at Bellcore, where he works in the Network Design and Traffic Analysis Research Group. He has been a leader of the work on self-similar traffic modeling at Bellcore.

Dr. Willinger is co-recipient of the 1996 IEEE W.R.G. Baker Prize Award and the 1994 W.R. Bennett Prize Paper Award from the IEEE Communications Society for the paper titled "On the Self-Similar Nature of Ethernet Traffic."