IEEE/ACM Transactions on Networking, Vol. 5, No. 1, 1997, 71-86.

Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level

Walter Willinger, Murad S. Taqqu, Robert Sherman and Daniel V. Wilson*

April 15, 1997

Abstract

A number of recent empirical studies of traffic measurements from a variety of working packet networks have convincingly demonstrated that actual network traffic is *self-similar* or *long-range dependent* in nature (i.e., bursty over a wide range of time scales) – in sharp contrast to commonly made traffic modeling assumptions. In this paper, we provide a plausible physical explanation for the occurrence of self-similarity in LAN traffic. Our explanation is based on new convergence results for processes that exhibit *high variability* (i.e., infinite variance) and is supported by detailed statistical analyses of real-time traffic measurements from Ethernet LAN's at the level of individual sources. This paper is an extended version of [53] and differs from it in significant ways. In particular, we develop here the mathematical results concerning the superposition of strictly alternating *ON/OFF* sources.

Our key mathematical result states that the superposition of many *ON/OFF* sources (also known as *packet trains*) with strictly alternating *ON*- and *OFF*-periods and whose *ON*-periods or *OFF*-periods exhibit the *Noah Effect* (i.e., have high variability or infinite variance) produces aggregate network traffic that exhibits the *Joseph Effect* (i.e., is self-similar or long-range dependent). There is, moreover, a simple relation between the parameters describing the intensities of the Noah Effect (high variability) and the Joseph Effect (self-similarity). An extensive statistical analysis of high time-resolution Ethernet LAN traffic traces (involving a few hundred active source-destination pairs) confirms that the data at the level of individual sources or source-destination pairs are consistent with the Noah Effect. We also discuss implications of this simple physical explanation for the presence of self-similar traffic patterns in modern high-speed network traffic for (i) parsimonious traffic modeling, (ii) efficient synthetic generation of realistic traffic patterns, and (iii) relevant network performance and protocol analysis.

1 Introduction

Starting with the extensive analyses of traffic measurements from Ethernet LAN's over a 4-year period described in [27], there have been a number of recent empirical studies that provide evidence of the prevalence of *self-similar* or *fractal* traffic patterns in measured traffic from today's high-speed networks. Prominent among these studies are the in-depth statistical analysis of large amounts of wide-area traffic measurements reported in [37, 38] and the detailed investigation of traffic data collected at the packet level from multiple NSFNET core switches presented in [22]. One of the most surprising findings from these and other studies concerns the ease with which it is possible to statistically distinguish between measured network traffic and traditional traffic models: actual traffic exhibits correlations over a wide range of time scales (i.e., has *long-range dependence*), while traditional traffic models typically focus on a very limited range of time scales and

*W. Willinger, R. Sherman and D.V. Wilson are with Bellcore, Morristown, NJ 07960-6438. M.S. Taqqu is with the Department of Mathematics, Boston University, Boston, MA 02215.

are thus *short-range dependent* in nature. Although such findings can in general be expected to favor the use of self-similar models (with underlying long-range dependent structure) over traditional models, there has been considerable resistance toward self-similar traffic modeling on practical grounds. One of the major reasons for this resistance has been the absence of satisfactory answers to the following 2 questions. (1) What is the physical "explanation" for the observed self-similar nature of measured traffic from today's packet networks? and (2) What is the impact of self-similarity on network and protocol design and performance analysis?

In this paper, we present an answer to question (1) by providing the appropriate mathematical results and by validating our findings with detailed statistical analyses of high time-resolution Ethernet traffic measurements. In particular, we provide a plausible and simple explanation for the observed self-similarity of measured Ethernet LAN traffic in terms of the nature of the traffic generated by the individual sources or source-destination pairs that make up the aggregate packet stream. Developing an approach originally suggested by Mandelbrot [31], we show that the superposition of many (strictly alternating) independent and identically distributed (i.i.d.) *ON/OFF* sources, each of which exhibits a phenomenon called the "Noah Effect", results in self-similar aggregate traffic. Here, by a strictly alternating *ON/OFF* source, we mean a model where the *ON*- and *OFF*-periods strictly alternate, where the *ON*-periods are independent and identically distributed (i.i.d.), the *OFF*-periods are i.i.d., and where the ON- and OFF-period sequences are independent from one another. The ON- and OFF-periods do not need to have the same distribution. By presenting the results in the well-known framework of these *ON/OFF* source models (also known as "packet train models"), we identify the Noah Effect as the essential point of departure from traditional to self-similar traffic modeling. Intuitively, the Noah Effect for an individual *ON/OFF* source model results in *ON*- and *OFF*-periods, i.e., "train lengths" and "intertrain distances" that can be very large with non-negligible probability. In other words, the Noah Effect guarantees that each *ON/OFF* source individually exhibits characteristics that cover a wide range of time scales. The Noah Effect is synonymous with the *infinite variance syndrome* – the empirical observation that many naturally occurring phenomena can be well described using distributions with infinite variance (for references, see [44, 42, 54]). Mathematically, we use *heavy-tailed* distributions with infinite variance (e.g., Pareto or truncated stable distributions) to account for the Noah Effect, and the parameter α describing the "heaviness" of the tail of such a distribution gives a measure of the intensity of the Noah Effect. We also provide a simple relation between α and the Hurst parameter H, where the latter has been suggested in [27] as a measure of the degree of self-similarity (or equivalently, of the "Joseph Effect") of the aggregate traffic stream.

In sharp contrast to our findings, traditional traffic modeling, when cast in the framework of *ON/OFF* source models, without exception assumes finite variance distributions for the *ON*- and *OFF*-periods (e.g., exponential distribution, geometric distribution). These assumptions drastically limit the *ON/OFF* activities of an individual source, and as a result, the superposition of many such sources behaves like white noise in the sense that the aggregate traffic stream is void of any significant correlations, except possibly some in the short range. This behavior is in clear contrast with measured network traffic (for details, see for example [28]). Note that the results of the present study suggest yet another, equally simple, statistical method for distinguishing between traditional and self-similar traffic: an analysis of network traffic that checks for the presence or absence of the Noah Effect in the traffic generated by the individual sources or source-destination pairs. To demonstrate the effectiveness of such an analysis, we used two Ethernet traffic traces, generated by about 100 and 3,200 individual sources (resulting in about 700 and 10,000 active source-destination pairs), respectively. The data were collected at the Bellcore Morristown Research and Engineering Center (MRE). One of the data sets is representative of Ethernet LAN traffic (consisting of all internal Ethernet packets), was collected in August of 1989, has been studied extensively in the past (at the aggregate packet level), and was, in fact, part of the analysis presented in [27]. The second data set represents a recent (December 1994) collection of high time-resolution WAN traffic measurements (consisting of all "remote" Ethernet packets, i.e., all packets destined for points outside of Bellcore ot for Bellcore from the outside), and includes applications such as WWW and Mbone. The motivation for including this second data set was to check whether WAN traffic self-similarity can be explained the same way as LAN traffic self-similarity or requires a different approach.

Although the main objective of this paper is to provide an answer to question (1) (physical "explanation"), our results concerning individual source behavior are clearly significant for answering question (2) (possible impact of self-similarity on network and protocol design and performance analysis). Starting with the work by Norros [35], there has been mounting evidence that clearly shows that the performance of queueing models with self-similar inputs can be radically different from the performance predicted by traditional traffic models, especially by Markovian models (e.g., see [9, 8, 12]). Here we complement this evidence by illustrating the practical relevance of our findings for (i) parsimonious traffic modeling for high-speed networks, (ii) efficient simulation of actual network traffic, and (iii) analyzing queueing models and protocols under realistic traffic scenarios.

Two previous studies of LAN traffic measurements are of particular relevance in the present setting. Jain and Routhier [19] used packet data collected at a ring network at MIT and proposed a "packet train" (or *ON/OFF*) source model in order to capture the observed burstiness in actual packet streams. In this context, our results show that packet train models are consistent with measured Ethernet LAN traffic collected at the level of individual source-destination pairs – once the Noah Effect for the packet-train lengths and the inter-train distances has been accounted for. By doing so, some of the shortcomings of the original packet train modeling approach (e.g., lack of any physical interpretation, arbitrary choice of crucial parameter values) are remedied; in particular, the use of infinite variance distributions for packet-train lengths and inter-train distances implies that packet trains (and inter-train distances) can be sensibly defined on all (or a wide range of) time scales. As a result, the need to pre-select a time scale that lacks physical interpretation and is often arbitrary can be avoided. Of particular importance to our work are Gusella's extensive studies [14, 15, 16] of traffic measurements from a 10 Mb/s Ethernet LAN. In view of the results discussed in the present paper, Gusella's work falls strictly within the traditional approach to traffic modeling: phenomena like the Joseph and Noah Effects are attributed to non-stationarity in the data and are ignored in subsequent data modeling. Naturally, the resulting models, based on burstiness characterizations using indices of dispersion, are adequate only over a limited range of time scales. Note that certain types of non-stationarity can indeed potentially imitate long-range dependence. However, applying new methods for testing the long-range dependence hypothesis against certain non-stationarity alternatives (e.g., see [50, 52]) shows that the long-range dependence hypothesis cannot be rejected for the majority of Ethernet traffic measurements considered in [27]. Thus, our approach suggests a viable alternative: by expanding the range of traditional traffic models to account for the Joseph and Noah Effects, it is possible to describe these phenomena in a stationary setting. The benefits for doing so include new insights into the time dynamics of high-speed network traffic, and the applicability of simple models for the very complex traffic patterns observed in today's networks.

The rest of the paper is organized as follows. In Section 2, we consider the superposition of many *ON/OFF* source models and present the convergence theorems that form the basis of our approach. In contrast to [53] where we assumed a special kind of *ON/OFF* sources, we consider here the more commonly used *ON/OFF* source model with strictly alternating *ON*- and *OFF*-periods. Due to space limitation, a proof of the main result and generalizations thereof will appear elsewhere (see the companion paper [48]). In Section 3, we discuss the available traffic measurements and present our statistical analysis of these data, concentrating on detecting the Noah Effect in LAN traffic generated by individual source-destination pairs. Finally, in Section 4 we illustrate the significance of the presence of the Noah Effect at the source level and its implications for aggregate traffic streams with a number of examples that are of practical importance for the design and performance analysis of modern communication networks and protocols. We conclude in Section 5 by outlining our current work focusing on an application-level based physical explanation of traffic self-similarity in a WAN environment.

Terminology: The main themes of this paper are "long-range dependence" and "self-similarity." In general, these two notions are not equivalent; the former involves the tail behavior of the autocorrelation function of a stationary time series, while the latter typically refers to the scaling behavior of the finite dimensional distributions of a continuous time or discrete time process. However, Cox [5] introduced the term "exactly second-order self-similar" for stationary sequences whose aggregated processes possess the same non-degenerate autocorrelation functions as the original process ("asymptotically second-order self-similar" has been defined in a corresponding way). In view of this definition, we use the terms "long-range dependence" and "(exactly or asymptotically second-order) self-similarity" in an interchangeable fashion, because both refer to the tail behavior of the autocorrelations and are essentially equivalent. In particular, when dealing with Gaussian processes – as we do in this paper – we call fractional Brownian motion as well as its increment process (i.e., fractional Gaussian noise) self-similar; while in the former case, self-similarity refers to the scaling behavior of the finite dimensional distributions of a continuous time process, in the latter case it is understood to mean exact second-order self-similarity and is synonymous with long-range dependence. (For a related discussion, see also Taqqu, Teverovsky and Willinger [47].) The context in which these notions appear will typically resolve any potential confusion. Note also that long-range dependence (Joseph Effect), self-similarity (regardless of the specific definition), as well as infinite variance (Noah Effect) are convenient mathematical idealizations (just as Markov processes and Brownian motions are) and can never be fully validated from finite data sets. As we will show, however, these idealizations offer simplification and clarity and capture in a parsimonious manner important characteristics of the data at hand. The terms "Joseph Effect" and "Noah Effect" were coined by Mandelbrot [33].

2 Self-Similarity Through High-Variability

In [53], we presented an idealized *ON/OFF* source model which allows for long packet trains ("*ON*" periods, i.e., periods during which packets arrive at regular intervals) and long inter-train distances ("*OFF*" periods, i.e., periods with no packet arrivals). In that model, however, the *ON*- and *OFF*-periods did not strictly alternate: they were i.i.d. and hence an *ON*-period could be followed by other *ON*-periods, and an *OFF*-period by other *OFF*-periods. The model was a relatively straightforward extension of the one first introduced by Mandelbrot [31] and Taqqu and Levy [46]. The processes we describe here have strictly alternating *ON*- and *OFF*-periods and agree therefore with the *ON/OFF* source models commonly considered in the communications literature. The *ON*- and *OFF*-periods, moreover, may have different distributions, either with infinite or finite variance (a partial treatment of the finite variance case can be found in [24]). Although our main result is essentially the same as in [53], namely, that the superposition of many such packet trains exhibits, on large time scales, the self-similar behavior that has been observed in the Ethernet LAN traffic data and WAN traces (see [27, 37]), the case of strictly alternating *ON/OFF* sources is much more delicate, and we provide a rigorous proof in the Appendix. Motivated by our earlier work in [53], Heath et al. [17] have independently obtained a proof for the asymptotics of the tail decay of the autocorrelation function of a heavy-tailed *ON/OFF* source that is essentially identical to the one presented in the Appendix; they also give the precise rate of decay. For related work, see also [29].

2.1 Homogeneous Sources

Suppose first that there is only one source and focus on the stationary binary time series $\{W(t),\ t \geq 0\}$ it generates. $W(t) = 1$ means that there is a packet at time t and $W(t) = 0$ means that there is no packet. Viewing $W(t)$ as the reward at time t, we have a reward of 1 throughout an *ON*-period, then a reward of 0 throughout the following *OFF*-period, then 1 again, and so on. The length of the *ON*-periods are i.i.d., those of the *OFF*-periods are i.i.d., and the lengths of *ON*- and *OFF*-periods are independent. The *ON*- and *OFF*-period lengths may have different distributions. An *OFF*-period always follows an *ON*-period, and it is the pair of *ON*- and *OFF*-periods that defines an interrenewal period.

Up to this point, we have considered only one source. Suppose now that there are M i.i.d. sources. Since each source sends its own sequence of packet trains it has its own reward sequence $\{W^{(m)}(t),\ t \geq 0\}$. The superposition or cumulative packet count at time t is $\sum_{m=1}^{M} W^{(m)}(t)$. Rescaling time by a factor T, consider

$$W_M^*(Tt) = \int_0^{Tt} \left(\sum_{m=1}^{M} W^{(m)}(u) \right) du,$$

the aggregated cumulative packet counts in the interval $[0, Tt]$. We are interested in the statistical behavior of the stochastic process $\{W_M^*(Tt),\ t \geq 0\}$ for large M and T. This behavior depends on the distributions on the *ON*- and *OFF*-periods, the only elements we have not yet specified. Motivated by the empirically derived fractional Brownian motion model for aggregate cumulative packet traffic in [53], or equivalently, by its increment process, the so-called fractional Gaussian noise model for aggregate traffic (i.e., number of packets per time unit), we want to choose these distributions in such a way that, as $M \to \infty$ and $T \to \infty$, $\{W_M^*(Tt),\ t \geq 0\}$ adequately normalized is $\{\sigma_{\lim} B_H(t),\ t \geq 0\}$, where $\sigma_{\lim}$ is a finite positive constant and B_H is *fractional Brownian motion*, the only Gaussian process with stationary increments that is self-similar. By self-similar, we mean that the finite-dimensional distributions of $\{T^{-H} B_H(Tt),\ t \geq 0\}$ do not depend on the chosen time scale T. The parameter $1/2 \leq H < 1$ is called the *Hurst parameter* or the *index of self-similarity*. Fractional Brownian motion is a Gaussian process with mean zero, stationary increments and covariance function $EB_H(s)B_H(t) = (1/2)\{s^{2H} + t^{2H} - |s-t|^{2H}\}$. Its increments $G_j = B_H(j) - B_H(j-1),\ j = 1, 2, \ldots$ are called *fractional Gaussian noise*. They are strongly correlated:

$$EG_H(j)G_H(j+k) \sim H(2H-1)k^{2H-2} \quad \text{as} \quad k \to \infty,$$

where $a_k \sim b_k$ means $a_k/b_k \to 1$ as $k \to \infty$. The power law decay of the covariance characterizes long-range dependence. The higher the H the slower the decay. For more information about fractional Brownian motion and fractional Gaussian noise, refer for example to Chapter 7 of Samorodnitsky and Taqqu [42].

To specify the distributions of the *ON*- and *OFF*-periods, let

$$f_1(x), \quad F_1(x) = \int_0^x f_1(u)du, \quad F_{1c}(x) = 1 - F_1(x), \quad \mu_1 = \int_0^\infty x f_1(x)dx, \quad \sigma_1^2 = \int_0^\infty (x-\mu_1)^2 f_1(x)dx$$

denote the probability density function, cumulative distribution function, complementary (or tail) distribution, mean length and variance of an *ON*-period, and let $f_2, F_2, F_{2c}, \mu_2, \sigma_2^2$ correspond to an *OFF*-period. Assume as $x \to \infty$,

$$\text{either} \quad F_{1c}(x) \sim \ell_1 x^{-\alpha_1} L_1(x) \text{ with } 1 < \alpha_1 < 2 \quad \text{or} \quad \sigma_1^2 < \infty,$$

and

$$\text{either} \quad F_{2c}(x) \sim \ell_2 x^{-\alpha_2} L_2(x) \text{ with } 1 < \alpha_2 < 2 \quad \text{or} \quad \sigma_2^2 < \infty,$$

where $\ell_j > 0$ is a constant and $L_j > 0$ is a slowly varying function at infinity, that is $\lim_{x\to\infty} L_j(tx)/L_j(x) = 1$ for any $t > 0$. For example, $L_j(x)$ could be asymptotic to a constant, to $\log x$, to $(\log x)^{-1}$, etc. Since the function L_j will be used as normalization in (1) below, it is preferable not to absorb the constant ℓ_j into it. (We also assume that either probability densities exist or that $F_j(0) = 0$ and F_j is non-arithmetic, where F_j is called arithmetic if it is concentrated on a set of points of the form $0, \pm\lambda, \pm 2\lambda, \cdots$.) Note that the mean μ_j is always finite but the variance σ_j^2 is infinite when $\alpha_j < 2$. For example, F_j could be Pareto, i.e. $F_{jc}(x) = K^{\alpha_j} x^{-\alpha_j}$ for $x \geq K > 0$, $1 < \alpha_j < 2$ and equal 0 for $x < K$, or it could be exponential. Observe that the distributions F_1 and F_2 of the *ON*- and *OFF*-periods are allowed to be different. One distribution, for example, can have a finite variance, the other an infinite variance.

In order to state the main result, we need to introduce some notation. When $1 < \alpha_j < 2$, set $a_j = \ell_j(\Gamma(2-\alpha_j))/(\alpha_j - 1)$. When $\sigma_j^2 < \infty$, set $\alpha_j = 2, L_j \equiv 1$ and $a_j = \sigma_j^2$. The normalization factors and the limiting constants in the theorem below depend on whether

$$\Lambda = \lim_{t \to \infty} t^{\alpha_2 - \alpha_1} \frac{L_1(t)}{L_2(t)}$$

is finite, 0, or infinite. If $0 < \Lambda < \infty$, set $\alpha_{\min} = \alpha_1 = \alpha_2$,

$$\sigma_{\lim}^2 = \frac{2(\mu_2^2 a_1 \Lambda + \mu_1^2 a_2)}{(\mu_1 + \mu_2)^3 \Gamma(4 - \alpha_{\min})}, \quad \text{and} \quad L = L_2;$$

if, on the other hand, $\Lambda = 0$ or $\Lambda = \infty$, set

$$\sigma_{\lim}^2 = \frac{2\mu_{\max}^2 a_{\min}}{(\mu_1 + \mu_2)^3 \Gamma(4 - \alpha_{\min})}, \quad \text{and} \quad L = L_{\min},$$

where min is the index 1 if $\Lambda = \infty$ (e.g. if $\alpha_1 < \alpha_2$) and is the index 2 if $\Lambda = 0$, max denoting the other index. We claim that under the conditions stated above the following holds:

Theorem 1. *For large M and T, the aggregate cumulative packet process $\{W_M^*(Tt),\ t \geq 0\}$ behaves statistically like*

$$TM \frac{\mu_1}{\mu_1 + \mu_2} t + T^H \sqrt{L(T)M} \sigma_{\lim} B_H(t)$$

where $H = (3 - \alpha_{\min})/2$ and $\sigma_{\lim}$ is as above. More precisely,

$$\mathcal{L} \lim_{T \to \infty} \mathcal{L} \lim_{M \to \infty} T^{-H} L^{-1/2}(T) M^{-1/2} \left(W_M^*(Tt) - TM \frac{\mu_1}{\mu_1 + \mu_2} t \right) = \sigma_{\lim} B_H(t), \tag{1}$$

where $\mathcal{L}\lim$ means convergence in the sense of the finite-dimensional distributions.

Heuristically, Theorem 1 states that the mean level $TM(\mu_1/(\mu_1 + \mu_2))t$ provides the main contribution for large M and T. Fluctuations from that level are given by the fractional Brownian motion $\sigma_{\lim} B_H(t)$ scaled by a lower order factor $T^H L(T)^{1/2} M^{1/2}$. As in [46], it is essential that the limits be performed in the order indicated. Also note that $1 < \alpha_{\min} < 2$ implies $1/2 < H < 1$, i.e., long-range dependence. Thus, the main ingredient that is needed to obtain an $H > 1/2$ is the heavy-tailed property

$$F_{jc}(x) \sim \ell_j x^{-\alpha_j} L_j(x), \quad \text{as } x \to \infty,\ 1 < \alpha_j < 2 \tag{2}$$

for the *ON*- or *OFF*-period; that is, a hyperbolic tail (or power law decay) for the distributions of the *ON*- or *OFF*-periods with an α between 1 and 2. A similar result obtains if $W_M^*(Tt)$ is replaced by the cumulative number of bytes in $[0, Tt]$. For a proof of Theorem 1, a generalization of the results to allow for the case of heterogeneous sources, for a detailed description of the limiting behavior in (1) when F_1 and F_2 satisfy special properties, and for a weak convergence analogue of Theorem 1, see [48].

3 Ethernet Traffic Measurements at the Source Level

In this section, we first describe two sets of Ethernet traffic measurements that will be analyzed in detail. The first data set is used to support our claim that the self-similar nature of Ethernet LAN traffic is caused by the presence of the Noah Effect in the traffic generated by the individual source-destination pairs that make up the aggregate packet stream. The data consist of one hour worth of Ethernet LAN traffic and is representative of the Ethernet LAN traces considered in [27] which have been shown to be consistent with (second-order) self-similarity. The other Ethernet LAN traces analyzed in [27] reveal a similar behavior (not shown here) in terms of the Noah Effect of the underlying source-destination pairs and could have been used for our purpose as well.

The motivation to include a second data set is as follows. Recall that some of the traces analyzed in [27] are comprised entirely of "remote" Ethernet packets and hence represent what is typically referred to as WAN traffic. Since these WAN traces have also been shown in [27] to be consistent with (asymptotic second-order) self-similarity, it is natural to pose the question whether WAN traffic self-similarity can be explained in the same way as LAN traffic self-similarity, namely in terms of the Noah Effect exhibited by the underlying source-destination pairs that make up the aggregate traffic, or whether WAN traffic points toward a different physical explanation, possibly at the application level. To this end, we could have considered one of the WAN traces previously analyzed in [27], but we decided instead to choose as our second data set a more recent, hour-long WAN traffic trace that contains WWW as well as Mbone traffic. The findings reported below, however, apply to the earlier WAN traces as well.

The two data sets result in about 500 and 10,000 active source-destination pairs, respectively. These numbers present a considerable challenge when trying to investigate in a statistically rigorous manner the presence of the Noah Effect in the traffic streams generated by all or a large part of these individual active source-destination pairs. Clearly, a compromise is needed between making the analysis "fully compelling" and keeping at the same time the amount of work at a reasonable level. Being "persuasive" is perhaps all we can achieve under these circumstances, in addition to providing mathematical and statistical rigor, setting a reasonable standard for future empirical and statistical analysis of such data, and motivating other researchers to reproduce our results in different network environments. Thus, one of the main objectives of this section is to illustrate the use of exploratory data analysis tools that can assist in extracting essential information out of an abundance of traffic measurements without an extraordinary effort. While some of the tools applied below are well-known, others are less familiar and will be explained in more detail as they are used.

A delicate compromise is also needed between "theory" and "practice" of parameter estimation for heavy-tailed data. Parameter estimation for heavy-tailed distributions is an active area of research, the few known estimation procedures (see below), which are theoretically reasonably well understood can perform quite erratically in practice, and some theoretical results (e.g., confidence intervals) for these procedures are known to hold only under conditions that often cannot be validated in practice. The resulting balancing act favors data-intensive heuristics over unfounded statistical rigor, and typically results in strong empirical evidence at the macroscopic level (i.e., for or against the presence of the Noah Effect in a given data set), but only approximate results at the microscopic level (i.e., estimating the intensity of the Noah Effect, that is, point estimates for α). In this paper, we have made a conscious decision in favor of data-intensive heuristics. As a result, precise point estimates for the index α appearing in equation (2) and measuring the exact intensity of the Noah Effect for a given source or source-destination pair are of secondary interest; we are primarily concerned with determining ranges of α-values that are consistent with the data representing individual source-destination pairs.

3.1 Traffic Measurements

The first set of traffic measurements is the busy hour of the August 1989 Ethernet LAN measurements presented and analyzed (and denoted by AUG89.HB and AUG89.HP in Table 1) in [27] (for further details about this data set, see also [28]). In addition to the information about time stamp and size (in bytes) of every (internal) Ethernet packet seen during this hour, the data set also contains the source and destination address of each recorded packet. During this busy hour, 105 hosts sent or received packets over the network (out of 121 hosts that were active during the whole 27 hour long monitoring session). The source-destination matrix (based on Ethernet addresses) corresponding to this set of traffic data is depicted in Figure 1, where we also indicated the activity level (i.e., total number of bytes sent by an individual source or source-destination pair, normalized by the total numbers of bytes sent by all hosts during this period) for each source (see marginals along x-axis), for each destination (marginals along y-axis), and for each individual source-destination pair. Notice that out of 11,025 possible source-destination pairs, only 748 or about 6.8% were actually sending or receiving packets (this effect has also been observed in previous traffic studies, e.g., [7, 36]). The most active hosts were sources 1, 7, 11, 27, 32, 58 (6 Sun-3 fileservers), sources 2 and 47 (2 DEC 3100 fileservers), source 34 (a Sun-4 server), sources 6, 15, 20, 25, 30, 63 (6 diskless Sun-3 clients), source 8 (a DEC 3100 client), and sources 10 and 17 which served as routers. Only about 5% of the traffic on this network was *external*, i.e., destined for machines on other networks or outside the company.

While the first data set consists of LAN traffic, the second data set is made up entirely of *remote* traffic, i.e., of packets destined for points on the Internet outside of Bellcore or for Bellcore from the outside, and represents what is usually referred to as WAN traffic. More precisely, the second data set consists of a "typical" hour of

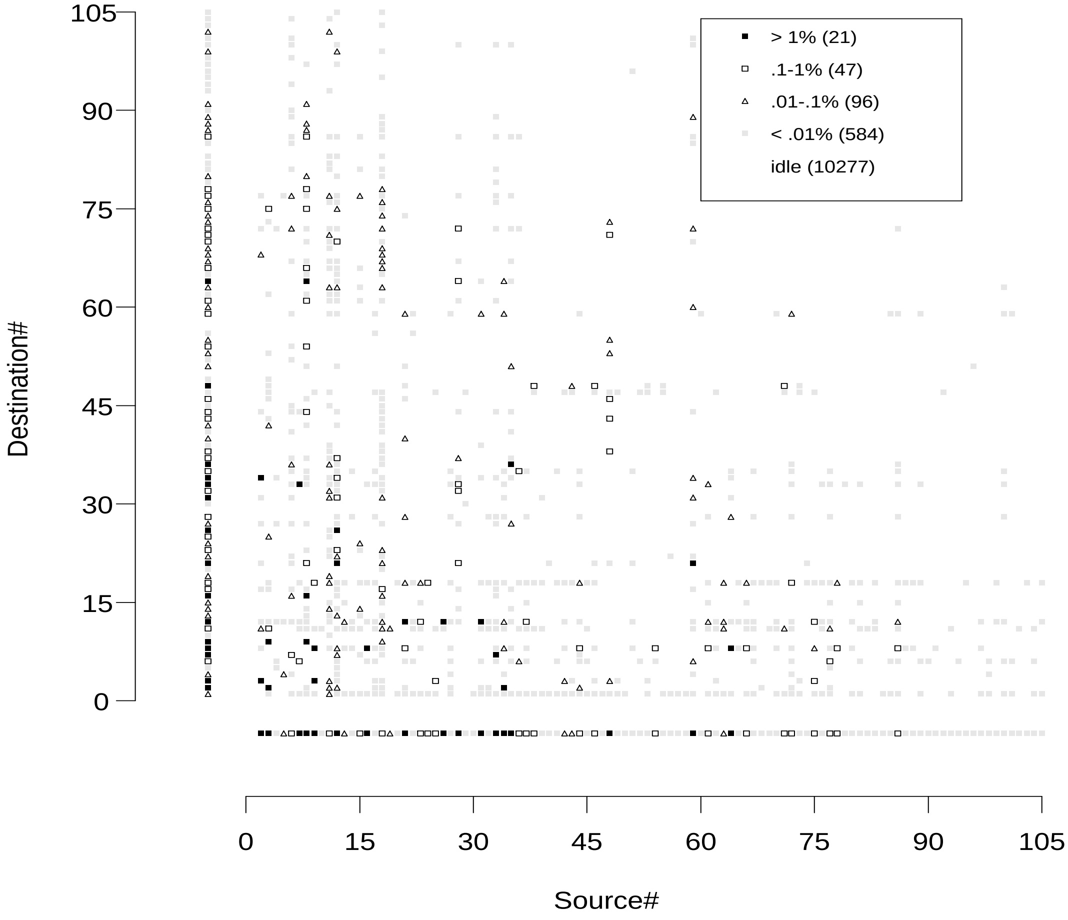

Figure 1: Source-Destination matrix for the August 1989 traffic data set.

traffic gathered from the stub Ethernet between the router provided by Bellcore's Internet service provider and a second Bellcore-controlled router that enforces security. For this data set, the number of active hosts (based on IP addresses) turns out to be about 3,500, while the percentage of active to possible source-destination pairs is about 0.25%. The most active host in this data set was the machine outside of Bellcore that sent *Mbone* packets (see below and Section 3.6 for more details regarding *Mbone*). Also included in the most active machines were four machines outside of Bellcore supplying data in response to file transfer (FTP) sessions, along with one Bellcore host supplying file transfer, E-mail, and Domain name service to the outside world. Other active hosts included one Bellcore host supplying Network News to the outside and three machines supplying news to Bellcore. Of the two Bellcore machines mentioned above, one is a Sun Sparcserver 690MP and the other (the network news supplier) is a Sparc-1. Note that all remote traffic is bandwidth limited by Bellcore's 1.5 Mb/s link to the outside world. A brief investigation of what applications generated this second hour-long data set revealed that a new Internet service called *Mbone* (see for example, [11]) was responsible for over 50% of the recorded traffic (in bytes). Another service, the World Wide Web (WWW) information retrieval service (e.g., see [1]), made up 9.4% of the total traffic. Neither Mbone nor WWW traffic was present in the first data set, nor in any of the earlier WAN traces studied in [27]. Services such as file transfer (14.5%), telnet/rlogin (2.8%), electronic mail (SMTP) (3.2%) and Network News transfer (NNTP) (12.2%) still present significant components of the total traffic but no longer dominate it.

3.2 Textured Plots and the Packet Train Assumption

We consider here the first data set that has been shown in [27] to be consistent with second-order self-similarity, with a Hurst parameter of $H \approx 0.90$ for the time series representing the packet counts per 10 milliseconds. This conclusion was reached by treating the Ethernet packets as black boxes, i.e., without using any information contained in the packet header fields. In contrast, for the present study, we extracted from the header field of each packet monitored during this hour the corresponding pair of source-destination addresses. This process resulted in 105 individual time series representing the packet arrivals on the Ethernet from the 105 hosts that were active (i.e., sent or received packets) during this hour. Furthermore, separating the packets generated by a given source depending on the packet's destination address yields a total of 748 time series corresponding to the number of active source-destination pairs. In view of the results presented in Section 2, we are thus faced with the challenging task of analyzing 748 time series with sufficient statistical rigor and accuracy to conclude whether or not these data support our physical explanation for self-similarity, i.e., whether or not the data are consistent (i) with the *ON/OFF* traffic model assumption for individual sources or source-destination pairs and (ii) with the crucially important assumption of the Noah Effect for the corresponding *ON*- and *OFF*-periods. To this end, our goal is not to provide a single point estimate for the intensity α of the Noah Effect, but to examine if there is evidence for the Noah Effect in the data and if so, to determine the "typical" range of α-values. Note that because of the basic relation $H = (3 - \alpha)/2$ (see Theorem 1), the earlier findings in

[27] of $H \approx 0.90$ for the time series of (aggregate) packet counts suggests the presence of the Noah Effect with a low α-value of about 1.20.

For the purpose of checking the appropriateness of the *ON/OFF* traffic modeling assumption for individual sources or source-destination pairs, we first make use of a simple exploratory data analysis tool called *textured dot strip plot* or simply *textured plot*, originally proposed in [51] and subsequently incorporated into *XGobi*, a tool for interactive dynamic graphics and analysis of multi-dimensional data (e.g., see [43]). Intuitively, the idea of textured plots is to display one-dimensional data points in a strip in an attempt to show all data points individually. Thus, if necessary, the points are displaced vertically by small amounts that are partly random, partly constrained. The resulting textured dot strip facilitates a visual assessment of changing patterns of data intensities in a way other better-known techniques such as histogram plots, one-dimensional scatterplots, or box-plots are unable to provide, especially in the presence of extreme values. To illustrate the effectiveness of textured plots for assessing the bursty or *ON/OFF* nature of traffic generated by an individual source or source-destination pair, we display in Figure 2 six textured plots associated with source 10 (other sources result in similar plots). Each point in the plots represents the time of a packet arrival. Serving as a router, this source contributed 1.85% to the overall number of packets and sent data to 25 different destinations. The top plot in Figure 2 represents the textured dot strip corresponding to the arrival times of all packets originating from source 10 (there are 26,330 packets), and the subsequent 5 panels result from applying the textured plot technique to the arrival times of all packets originating from source 10 and destined for sources 1, 18, 70, 13 and 17, respectively. These 5 source-destination pairs were responsible for 5,901, 4,050, 3,407, 2,135, and 1,918 packets, respectively, and make up about 66% of all the packets generated by source 10.

Figure 2 supports two important observations regarding the bursty behavior of traffic generated by (i) a reasonably active individual source (e.g., source 10) and (ii) a "typical" individual source-destination pair (e.g., source-destination pair 10-18). First, a close look at the top plot of Figure 2 (source 10) clearly reveals the burstiness expected from actual packet traffic, but offers little hope for supporting the *ON/OFF* nature of the underlying traffic that gave rise to this strip plot. The plot is even more discouraging from the point of view of hoping for some "objective" criterion for identifying *ON/OFF* periods. However, a glance at the 5 source-destination plots in Figure 2 makes the *ON/OFF* behavior of the traffic generated by the individual source-destination pairs immediately apparent. There is little question of what is meant by an *ON*- or *OFF*-period, and subsequent extraction of the lengths of the *ON*- and *OFF*-periods from a given source-destination traffic trace is greatly facilitated by use of of the corresponding textured plot of packet arrival times.

3.3 Checking for the Noah Effect

Next we illustrate the techniques used for determining the presence or absence of the Noah Effect for the *ON*- and *OFF*-periods derived from the traffic data generated by individual sources or source-destination pairs. In the case where the data are found to be consistent with the Noah Effect, these techniques also allow for fast

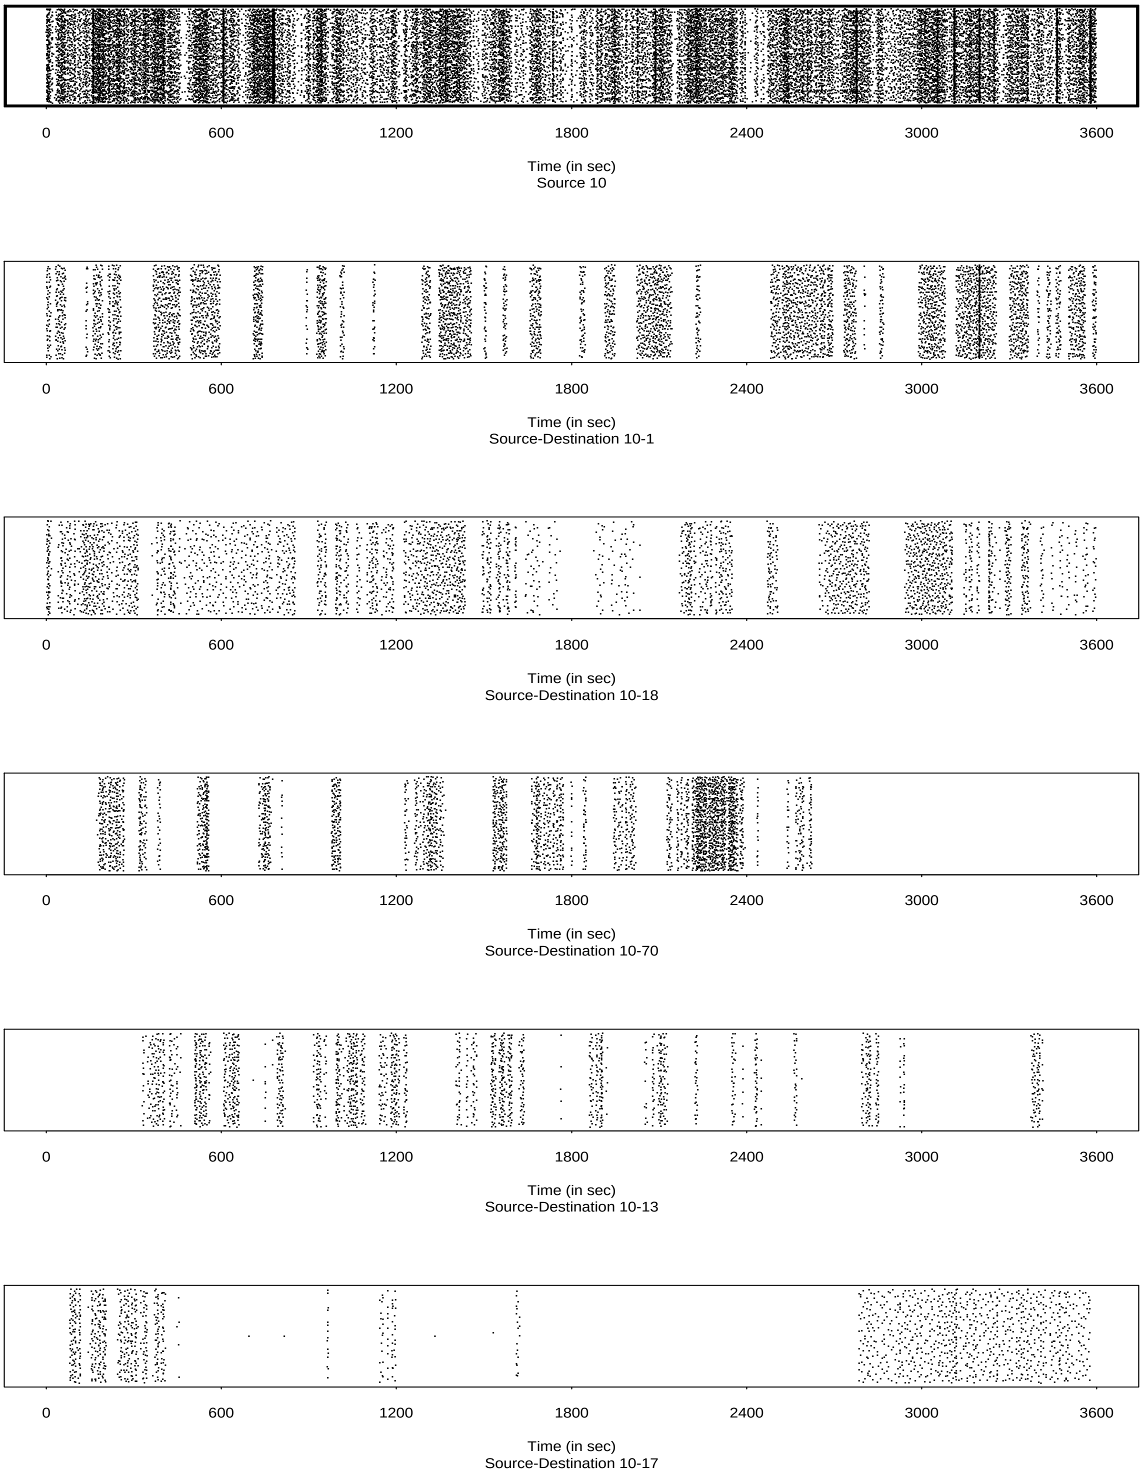

Figure 2: Textured plots of packet arrival times for (top to bottom) source 10 and source-destination pairs 10-1, 10-18, 10-70, 10-13 and 10-17.

procedures (partly heuristic, partly rigorous) for estimating the intensity α of the Noah Effect. As we will demonstrate, these techniques work best when applied in combination with each other and with attention to the physical structure of the data. Specifically, we make extensive use of *complementary distribution plots* (related to the *qq-plot* method [23]) and *Hill's method* [18, 41] for estimating α.

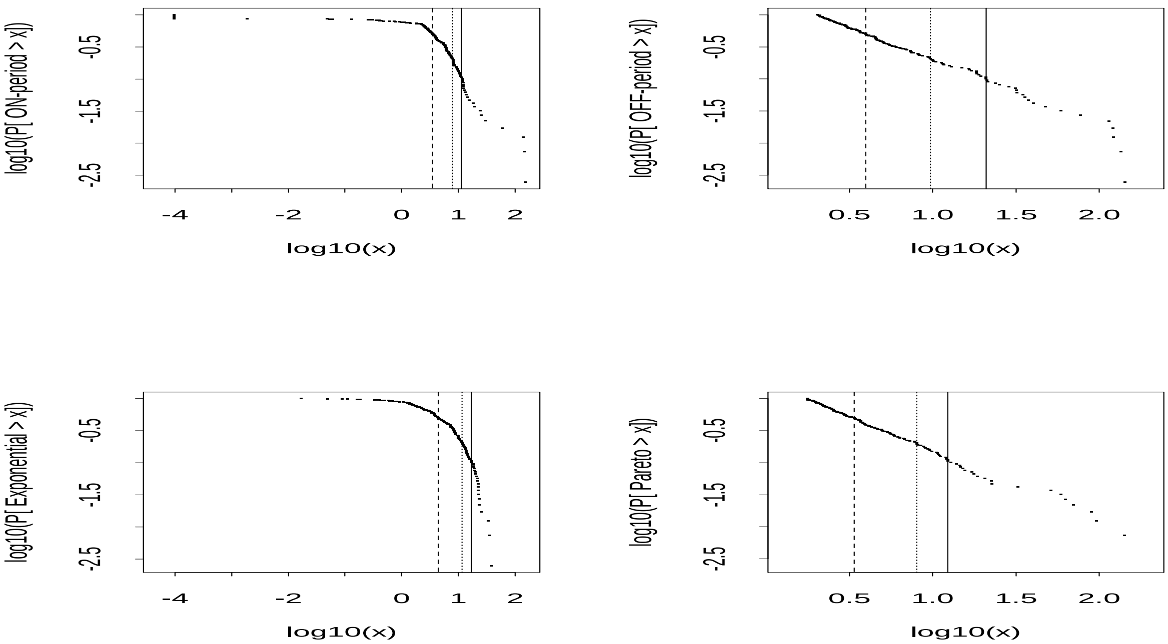

Figure 3: Complementary distribution plots for *ON*-periods (top left) and *OFF*-periods (top right) for the source-destination pair 10-18, using a threshold value of $t = 2s$; for a sample from an *exponential distribution* that matches the mean of the *ON*-periods (lower left), and for a sample from a *Pareto distribution* that matches the mean of the *OFF*-periods (lower right). (The vertical solid, dotted and dashed lines indicate that 10%, 20% and 50% of all data points are to the right of the respective lines.)

In order to determine the presence or absence of the Noah Effect in a given data set, we take logarithms of both sides of relation (2), obtaining

$$\log F_{jc}(x) \sim log(l_j) \ - \alpha_j \ log(x), \text{ as } \ x \to \infty, \ 1 < \alpha_j < 2. \tag{3}$$

Using complementary distribution plots, i.e., plotting (on a $\log - \log$ scale) the complementary empirical distribution function of a sample that was presumably drawn from a distribution that exhibits hyperbolic tails (i.e., satisfies (2)), results in an approximately straight line for large x-values, with a slope of $-\alpha, 1 < \alpha < 2$. To illustrate the effectiveness of this technique for the data at hand, we concentrate on source-destination pair 10-18 (Figure 2, panel 3); other source-destination pairs yield similar results. Based on its textured plot, we define an *OFF*-period to be *any interval of length* $t \geq 2$ *seconds that does not contain any packet*; this, in turn, defines the *ON*-periods unambiguously and results in a total of 202 *ON*-periods and the same number of *OFF*-

periods for this source-destination pair. (We will return to the issues of the particular choice of the threshold t and of the robustness of our results under different threshold values in Section 3.4 below.) Figure 3 depicts the complementary distribution plots (on $\log-\log$ scale) of the *ON*-periods (top left) and *OFF*-periods (top right) and indicates a straight line behavior for large x-values, i.e., a hyperbolic tail distribution satisfying (2) for the *ON*- and *OFF*-periods. In fact, a heuristic estimate (obtained by "eyeballing" a straight line through the points to the right of the dashed vertical line) yields $\alpha \approx 1.7$ for the *ON*-periods and $\alpha \approx 1.2$ for the *OFF*-periods. To compare, Figure 3 also includes the complementary distribution plot of (i) an exponential distribution with the same mean of 7.2 s as the *ON*-periods (bottom left), and (ii) a Pareto distribution with $\alpha = 1.2$ and the same mean of 10.5 s as the *OFF*-periods (bottom right). Clearly, when compared to the distribution of the *ON*-periods, the exponential distribution as a whole as well as its right tail are concentrated on a more narrower range of x-values; in addition, the tail of the exponential distribution falls off much faster than the tail of the *ON*-periods. On the other hand, the Pareto distribution covers practically the same range of time scales as the empirically observed *OFF*-periods and matches the straight line behavior of the data over practically the whole x-axis.

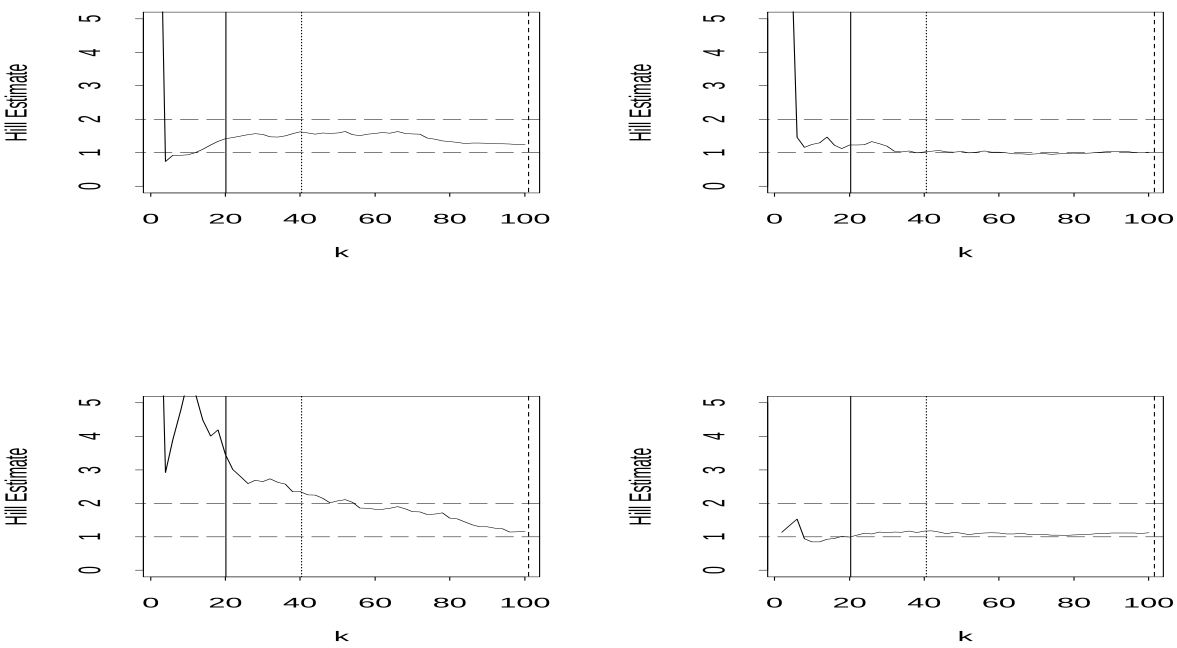

Figure 4: Hill estimate plots for *ON*-periods (top left) and *OFF*-periods (top right) for the source-destination pair 10-18, using a threshold value of $t = 2s$; for a sample from an *exponential distribution* that matches the mean of the *ON*-periods (lower left), and for a sample from a *Pareto distribution* that matches the mean of the *OFF*-periods (lower right). (The vertical solid, dotted and dashed lines indicate that 10%, 20% and 50% of the largest order statistics have been included in the Hill estimation calculation.)

While complementary distribution plots often provide solid evidence for or against the Noah Effect in a given data set, the eyeballing method described above for producing a rough estimate for α is cumbersome and unsatisfactory. A statistically more rigorous method for estimating the intensity of the Noah Effect is known as *Hill's estimator* and is described in [18] (see also [41]). Briefly, let $U_1, U_2, \ldots, U_n$ denote, for example, the observed *ON*-periods, and write $U_{(1)} \leq U_{(2)} \leq \ldots \leq U_{(n)}$ for the corresponding order statistics. The Hill estimator of α is

$$\hat{\alpha_n} = \left(1/k \sum_{i=0}^{i=k-1} (\log U_{(n-i)} - \log U_{(n-k)}) \right)^{-1}, \tag{4}$$

where the choice of $1 < k \leq n$ indicates how many of the largest observations enter into the calculation of formula (4). In practice, one plots the Hill estimator $\hat{\alpha_n}$ as a function of k, for a range of k-values. In the presence of a tail behavior in the data that is consistent with (2), a typical *Hill plot* varies considerably for small values of k (i.e., only a small fraction of the largest observations are considered), but becomes more stable as more and more data points in the tail of the distribution are included (often up to a cut-off value, to the left of which (2) no longer holds). An apparent straight line behavior for large x-values in the complementary distribution plot corresponds to a region of k-values where the Hill's estimate remains stable. In the absence of such a straight line behavior, the Hill's estimate will continue to decrease as k increases, a strong indication that the data are not consistent with the hyperbolic tail assumption (2). Although a number of theoretical properties of the Hill estimator are known (e.g., see [40]), they typically require extra assumptions on the underlying distribution which are essentially unverifiable in practice. This is especially true for the asymptotic normality property of the Hill estimator, which is used to compute confidence intervals.

This is why, in this paper, we prefer data-intensive heuristics for specifying ranges for the Hill estimator over the use of confidence intervals that are theoretically exact but rely on conditions that cannot be verified for a given data set. Figure 4 depicts the Hill estimate plots corresponding to the data used in Figure 3. Recalling that each data set contains 202 observations, the top left plot depicts the Hill estimate for the *ON*-periods and should be viewed together with the top left plot in Figure 3; note the region of stability in the Hill plot (k-values between 20 to about 70), i.e., the tail of the distribution that is consistent with the hyperbolic decay as given in (2) contains about 40% of all the observations. Moreover, the Hill estimate can be readily read off from the y-axis and yields $\hat{\alpha} \approx 1.7$. In the case of the Hill plot for the *OFF*-periods (top right) and the fitted Pareto model (bottom right), the situation is obvious and in agreement with the information contained in the corresponding complementary distribution plots in Figure 3. Finally, the bottom left plot in Figure 4 illustrates the typical behavior of the Hill plot when the data are inconsistent with assumption (2); the plot does not settle down but continues to decrease as more and more of the smaller order statistics are included in the calculation of the estimator. Intuitively, this behavior is caused by the concave shape (throughout the whole x-axis) of the complementary distribution plot (see bottom left plot in Figure 3).

3.4 A Robustness Property of the Noah Effect

Before checking for the presence of the Noah Effect in the traffic traces generated by the remaining source-destination pairs, we first point out a robustness feature of the Noah Effect that should greatly diminish any reluctance toward using *ON/OFF* or packet train models at the source or source-destination level. In the past, such reluctance has typically been based on a lack of physical interpretation or intuition for defining objectively the notion of an *OFF*-period (or, using the notation of Section 3.3, for selecting the "right" threshold value t). In the packet train terminology, the problem is to decide in a coherent manner on the "appropriate" intertrain distance, i.e., on deciding when the "departure" of the previous train took place and when the "arrival" of the next train occurs. Here we show, that *as far as the Noah Effect is concerned, it does not matter how the OFF-periods or intertrain distances (and subsequently, the ON-periods or packet train lengths) have been defined*. In other words, the Noah Effect is robust under a wide range of choices for the threshold value t that we used in Section 3.3 to explicitly define *OFF*-periods as any interval of length t seconds or larger that sees no packet arrival.

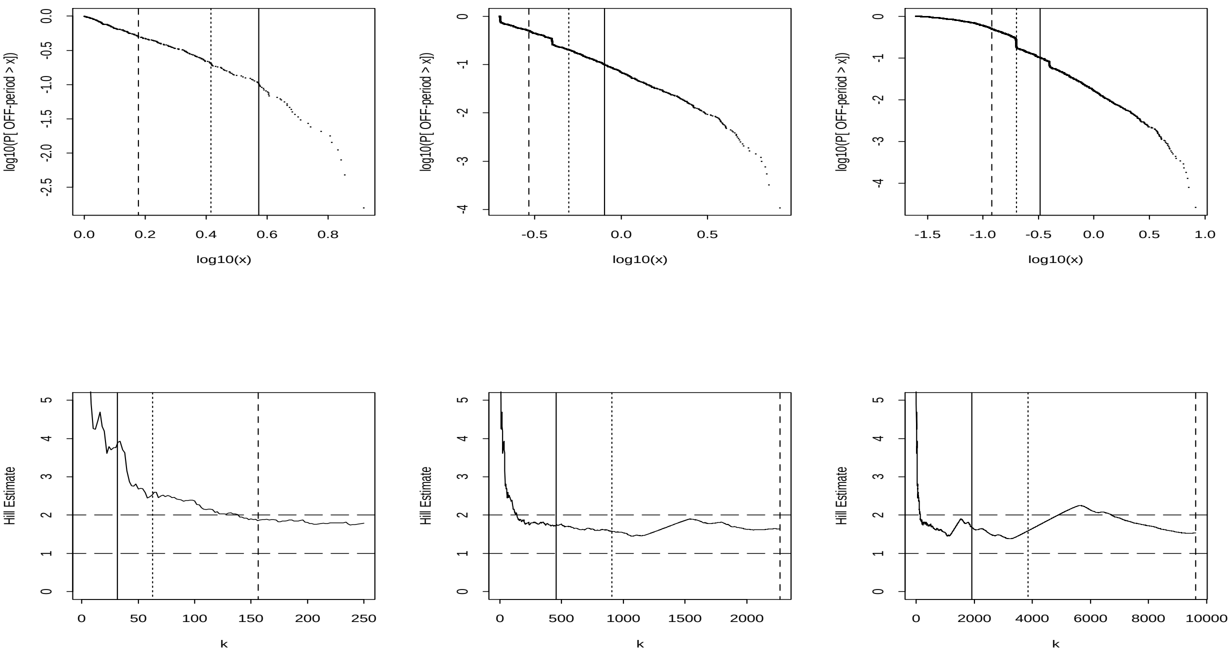

Figure 5: An illustration of the robustness property of the Noah Effect for the *OFF*-periods (using source 10). For threshold values $t = 1.0s, 0.20s$ and $0.025s$, the top row gives the complementary distribution plots and the bottom row the corresponding Hill plots.

The reason behind this insensitivity of the Noah Effect for (non-degenerate) *OFF*-periods to different choices of the threshold t is the well-known scaling property of distributions that satisfy the hyperbolic tail condition (2). Here, by scaling property we mean that if the distribution of the random variable U satisfies

(2) and t denotes a threshold value, then for sufficiently large u, t with $u > t$,

$$P(U > u | U > t) \sim (\frac{u}{t})^{-\alpha}, 1 < \alpha < 2, \tag{5}$$

(see [32], and also [37, 38] where this property is discussed and used for an intuitive explanation of the Joseph Effect in measured TELNET traffic traces). Thus, the tail behavior of the (conditional) distributions of U given $U > t$, for different choices of the threshold t, differs only by a scaling factor and hence gives rise to complementary distribution plots (on $\log - \log$ scale) with identical asymptotic slopes but different intercepts. This appealing robustness property of the Noah Effect for the *OFF*-periods with respect to the choice of t is illustrated in Figure 5 where we show the complementary distribution plots (top row) and corresponding Hill estimate plots (bottom row) for three different ways of defining the *OFF*-periods for the traffic associated with source 10 (see top plot of Figure 4). More specifically, we chose t-values that span 3 orders of magnitude, namely $t = 1s$ (left column, 313 observations), $t = 0.2s$ (middle column, 4,537 observations), and $t = 0.025s$ (right column, 19250 observations). Figure 5 confirms the robustness property of the Noah Effect under the different choices of t, with an estimated intensity between $1.6 - 1.9$. Recall, that our objective is not to come up with a precise point estimate for α, but to identify a range of α-values that is consistent with the given data.

A similar convincing robustness property of the Noah effect can be shown to hold for *ON*-periods (see [53], Fig. 6), although slightly different arguments from the ones we used in the case of the *OFF*-periods are needed to explain the insensitivity of the Noah Effect of the *ON*-periods to the different choices of the threshold t (see [53], pp. 107-108). Thus, while there is no "natural" division into *ON/OFF*-periods at the source level, such a division becomes apparent at all (a wide range of) time scales; moreover, these divisions appear in a consistent manner.

3.5 Self-Similarity and the Noah Effect in LAN Traffic

To facilitate the full-fledged statistical analysis of the first data set (i.e., the busy hour of the August 1989 traffic measurements) at the source-destination level, we consider in detail only the 181 most active (out of a total of 748 active) source-destination pairs. Together, these 181 source-destination pairs generated more than 93% of all the packets seen on the Ethernet during this hour and represent more than 98% of the overall traffic (in bytes). We chose to neglect all source-destination pairs that generated fewer than about 300 packets during the whole hour. Our statistical analysis of this abundance of traffic data benefited tremendously from the observed robustness property of the Noah Effects for the *ON*- and *OFF*-periods and from the availability of graphical tools that allow for effective visualization of complex data structures.

The summary plots in Figure 6 were obtained by checking, for each of the 181 source-destination pairs, for the presence or absence of the Noah Effect in their corresponding sequences of *ON*- and *OFF*-periods. For each source-destination pair, we typically generated 5 sequences of *ON*- and *OFF*-periods using 5 different threshold values, ranging between 5 s and 0.01 s. As a result of this data-intensive thresholding procedure, we obtain

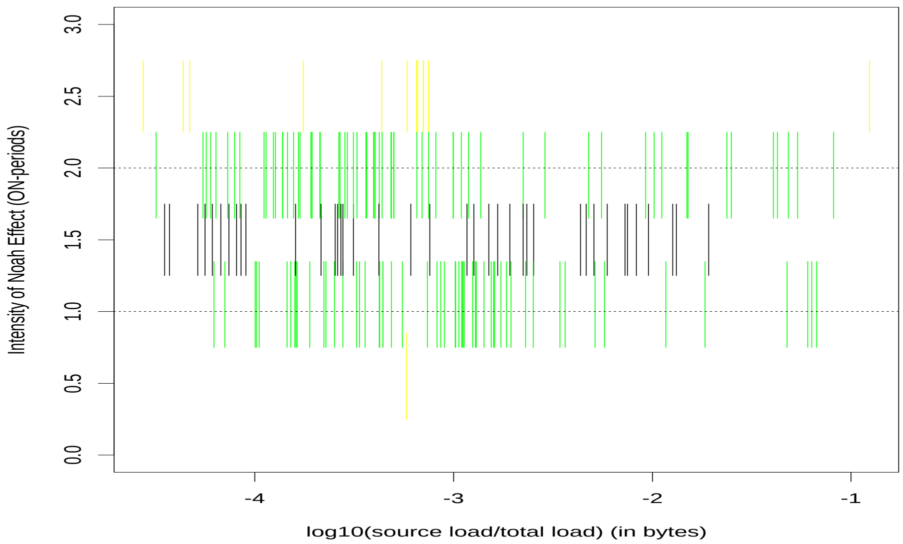

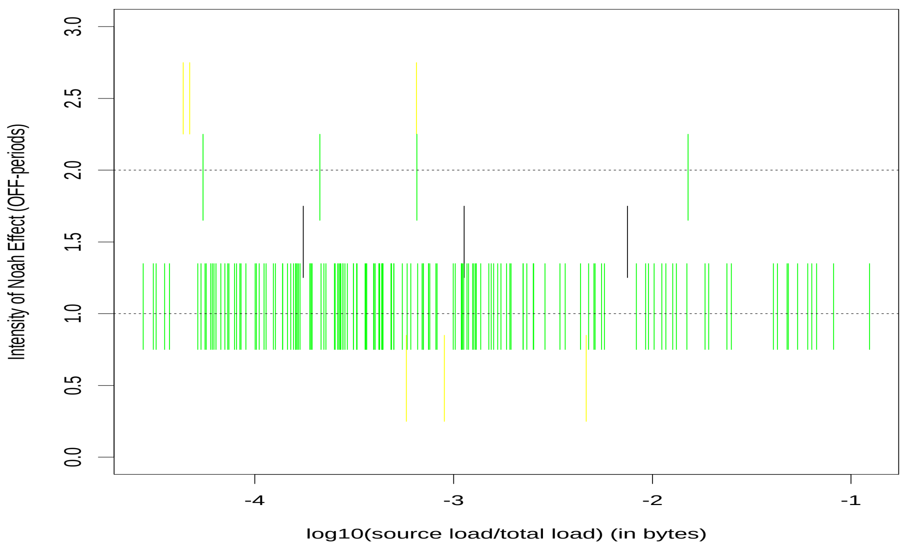

Figure 6: Summary plot of ranges for the α-estimates for the *ON*-periods (top) and *OFF*-periods (bottom) of the 181 most active source-destination pairs, as a function of their loads (in bytes, on log scale).

for each source-destination pair two ranges of α-values (one for the corresponding *ON*-periods, another for the *OFF*-periods) that are consistent with the data and insensitive to the particular definition of *ON/OFF*-periods. More precisely, we categorize the *ON/OFF*-nature of each source-destination pair, depending on whether the α-estimates obtained from the corresponding five different ways of defining the *ON/OFF*-periods are more or less the same and are consistent with α-estimates in the intervals $(0, .85), (.75, 1.35), (1.25, 1.75), (1.65, 2.25)$ or $(2.25, 2.75)$, representing the intuitively easy to define cases "definitely below 1.0", "around 1.0", "somewhere in the middle of the interval (1,2)", "around 2.0", and "definitely above 2.0 or inconclusive", respectively. For each source-destination pair considered, this categorization is based on a combination of (i) textured plots for visual assessment of *ON/OFF* nature of the traffic, (ii) complementary distribution plots as a quick heuristic method for checking the tail behavior of a distribution, and (iii) a careful interpretation of Hill plots (typically in connection with information obtained via (ii)). For the vast majority of source-destination pairs, the categorization process worked well, slightly better for the *OFF*-periods than for the *ON*-periods. Moreover, not all *ON/OFF*-periods fitted this framework, but the number of inconclusive cases was insignificant. The results are shown in Figure 6, where we plot for each of the 181 source-destination pairs its load (in bytes, on log scale) against the range of α-values that is consistent with its traffic trace. As can be seen, in the case of the *ON*-periods (top plot), the α-estimates consistent with the data cover pretty much the whole interval $(1, 2)$. In comparison, the bottom plot in Figure 6 shows that in the case of the *OFF*-periods, α-estimates in the lower part of the interval $(1, 2)$ clearly dominate the picture. Note that the mathematical results in Section 2 readily apply to this case involving different α-values for the *ON*- and *OFF*-periods.

When combined, the two plots in Figure 6 provide strong statistical evidence in favor of our proposed physical explanation for the empirically observed self-similarity property of aggregate Ethernet LAN traffic, in terms of the nature of traffic generated by each individual source-destination pair that sent packets over the Ethernet. In particular, our analysis shows that the data at the source-destination level are consistent with an *ON/OFF* modeling assumption for individual sources or source-destination pairs, and are in strong agreement with the assumption of the Noah Effect for the distributions of the corresponding *ON/OFF*-periods. (We have also done extensive testing of the independence assumption for the *ON*-periods and *OFF*-periods and have found the data to be in full agreement with it; for example, *ON/OFF* event counts over disjoint time intervals show, in general, no significant correlations.) In fact, one of the most astonishing findings from our analysis has been the extremely widespread and often very obvious presence of the Noah Effect, expressed via relation (2), in measured source-level LAN traffic data, regardless of whether the source represents a fileserver or a client machine. Possible explanations for this phenomenon typically refer to application-level characteristics and include (i) an empirically observed hyperbolic tail behavior for the file sizes residing in file systems such as file servers (see the discussion and references in [37]), (ii) a Pareto-like tail behavior for measured CPU time used by a typical Unix process (see [26]), (iii) measurements studies of an ISDN office automation application reported in [34] that suggests that human-computer interactions occur over a wide range of time scales and thus, may require models based on infinite variance distributions, and – in the case of more recent LAN measurements – (iv) empirically observed infinite variance properties for the sizes of documents that reside on many of today's WWW-servers (see [6]). A similar, though less thorough analysis at the level of individual source-destination pairs (not shown here) of some of the other Ethernet LAN traces considered in [27] shows full agreement with the above-mentioned findings and demonstrates the robustness of traffic characteristics such as the Joseph Effect (for aggregate LAN traffic) and the Noah Effect (for individual source-destination traffic) under a variety of changes (e.g., with regard to network configuration, host population, hardware and software upgrades, user applications) that working LANs experience over time.

3.6 Self-Similarity and the Noah Effect in WAN Traffic

In addition to discovering the self-similar nature of Ethernet LAN traffic, Leland et al. [27] also encountered self-similarity in traces that were comprised entirely of "remote" Ethernet packets (i.e., packets destined for points outside of Bellcore, or for Bellcore from the outside world) and represent what is commonly called WAN traffic. Naturally, the question arises whether WAN traffic self-similarity can be explained in the same way as LAN traffic self-similarity, namely in terms of the highly variable *ON/OFF*-nature of the traffic generated by the individual source-destination pairs, or whether WAN traffic points toward the need for a physical explanation at a different level. In addition, the question of robustness of traffic characteristics such as the Joseph and Noah Effect also arise in the WAN context, because WANs are known to experience drastic changes at the user and application level within relatively short periods of time (e.g., see [36].

To answer these questions, we considered some of the WAN traces that were collected between 1989 and 1992, analyzed in [27], and found to be consistent with (asymptotic second-order) self-similarity. In addition, we considered another data set that is described in Section 3.1 and represents an hour worth of WAN traffic collected in December 1994. For each of these data sets, we performed an analysis at the level of individual source-destination pairs that was similar – though less detailed – to the one described in Section 3.5 for LAN traces. Although we use mainly the December 1994 trace to illustrate our findings, the results described below apply to the earlier WAN traces as well and are based on an analysis that combines techniques introduced in [27] for analyzing aggregate traffic streams and methods illustrated earlier in Section 3.5 for dealing with individual source-destination traffic traces.

Given the information provided in Section 3.1 about the second data set, we first split the hour-long trace into two subsets; the first subset represents the traffic sent from the machine that furnishes *Mbone* traffic to Bellcore. It makes up about 52% of the total traffic (in bytes). The second subset consists of the remaining, i.e., all *non-Mbone* packets recorded during the given hour. Concentrating first on the *non-Mbone* traffic, an analysis along the lines of [27] of the aggregate traffic (number of packets per 10 milliseconds) reveals that this data set is consistent with (second-order asymptotic) self-similarity, and as an estimate of the degree of the Joseph Effect (i.e., the Hurst parameter), we obtain an H-value between 0.85 and 0.90. After separating the aggregate traffic into traffic traces generated by individual source-destination pairs, we get that the 300 most active (out of a total of about 10,000) source-destination pairs are responsible for 83% of the *non-Mbone* traffic. Analyzing these 300 traces in the same way as above shows again overall consistency of the data with the *ON/OFF* source model assumption. However, in contrast to the LAN traffic traces considered in Section 3.5, our findings for the WAN traces strongly suggest α-values for the Noah Effects for the *OFF*-periods that are typically around 1.0 and often even below 1.0 (implying infinite mean). On the other hand, the α-values corresponding to the *ON*-periods are typically around 2.0 (i.e., on the borderline between finite and infinite variance).

Upon closer examination of the behavior of the *OFF*-periods, we observe that a typical source-destination pair extracted from a WAN traffic trace starts transmission at some time into the data set, transmits packets (in some fashion) for a random duration, and then ceases transmitting packets for the remainder of the data set. Clearly, the long periods of inactivity at the beginning/end of a traffic trace generated by such a source-destination pair give rise to *OFF*-periods whose distributions exhibit extreme heavy right tails. This property of the *OFF*-periods is typical for all the WAN traces at hand and indicates that WAN traffic self-similarity cannot be explained in the same simple manner as LAN traffic self-similarity, namely in terms of the highly variable *ON/OFF*-nature of the individual source-destination pairs that make up the aggregate traffic. At the same time, the observed behavior of a typical source-destination pair in a WAN environment is consistent with the ground-breaking work by Floyd and Paxson [37] on WAN traffic characterization and suggests a different – equally simple – plausible physical explanation for the empirically observed self-similar nature of measured WAN traffic, but at the level of individual applications instead of individual source-destination pairs. We will expand on this explanation of WAN traffic self-similarity in Section 6.

In the case of the *Mbone* traffic data set, only an analysis of the aggregate packet stream was performed. The results indicate that *Mbone* traffic is asymptotically self-similar, with an H-value in the high 0.9 range. Its distinctive feature, however, is that only after aggregation levels beyond 100 ms does the strong intensity of the Joseph Effect become obvious, i.e., does the correlation structure of *Mbone* traffic remain unchanged as aggregation levels increase further. Based on our current understanding (see Section 2) and on extensive simulation studies (see Section 4 and, especially [39]), this property of *Mbone* traffic suggests the absence of the Noah Effect for the *ON*-periods (i.e., a corresponding intensity level for the *ON*-period α that exceeds 2.0) and at the same time, the presence of a strong (i.e., α-values closer to 1.0 than to 2.0) intensity of the Noah Effect for the *OFF*-periods of the individual user applications that typically run over *Mbone*.

4 Implications of the Noah Effect in Practice

Recall that the empirically observed self-similarity property in measured LAN traffic allows for a clear distinction – on statistical grounds – between traditional traffic models and actual traffic collected from working networks. The proposed physical explanation based on the Noah Effect enables us to identify the essential difference between self-similar and traditional traffic modeling in the setting of the well-known *ON/OFF* source models: traditional traffic modeling assumes finite variance distributions for the *ON*- and *OFF*-periods (in fact, exponential or geometric distributions are used almost exclusively), while self-similar modeling is based on the assumption of the Noah Effect, i.e., requires infinite variance distributions. Moreover, traditional traffic modeling becomes a special case of the self-similar approach by chosing α-values bigger than 2.0. From a more applied viewpoint, questions related to the impact of self-similarity in practice (e.g., generating realistic network traffic, performance of networks, protocols, and controls) can be reduced to the more basic question of the practical implications of the Noah Effect. In this section, we illustrate its impact with examples concerning traffic modeling, synthetic traffic generation, and network performance analysis.

4.1 Traffic Modeling and Generation

There is no question that today's network traffic is complex. Often, this is interpreted as saying that a mathematical model of this traffic must be complicated in nature, i.e., must be highly parameterized in order to realistically account for the observed complexity. One of the main results of this paper is that although network traffic is intrinsically complex, *parsimonious modeling* is still possible; we demonstrate, moreover, that it gives rise to a physical explanation for the self-similarity phenomenon that is simple and fully consistent with actual traffic measurements from a LAN environment. Thus, for aggregate traffic measurements, insistence on parsimonious modeling has lead to the use of self-similar (or long-range dependent) processes for traffic modeling at the aggregate level. In this paper, the desire for a "phenomenological" explanation of self-similarity in LAN traffic has resulted in new insights into the nature of traffic generated by the individual sources that contributed to the aggregate stream. We identified the Noah Effect as an essential ingredient, thus describing an important characteristic of the traffic in today's networks by essentially a single parameter, namely the intensity α of the Noah Effect in the *ON*- and *OFF*-periods of a "typical" network host. Whether we consider an idealized setting involving i.i.d. *ON*- and *OFF*-periods (see [53, Theorem 1]) or strictly alternating *ON/OFF* sources (see Section 2, Theorem 1) is not important for this finding. Generalizations accommodating more realistic conditions are possible (see [48]), maintain the simplicity of the basic result, and may require the addition of only a small number of physically meaningful parameters.

Explaining, and hence modeling self-similar phenomena in the traffic context in terms of the superposition of many *ON/OFF* sources with infinite variance distribution for the lengths of their *ON/OFF*-periods, leads to a straightforward method for generating long traces of self-similar traffic within reasonable (i.e., linear) time – assuming a parallel computing environment. Indeed, the results (e.g., Theorem 1) are tailor-made for parallel computing: letting every processor of a parallel machine generate traffic according to an alternating *ON/OFF* model (same α), simply adding (i.e., aggregating) the outputs over all processors produces self-similar traffic. For example, producing a synthetic trace of length 100,000 on a MasPar MP-1216, a massively parallel computer with 16,384 processors, takes on the order of a few minutes. In fact, Figure 8 shows the result of a

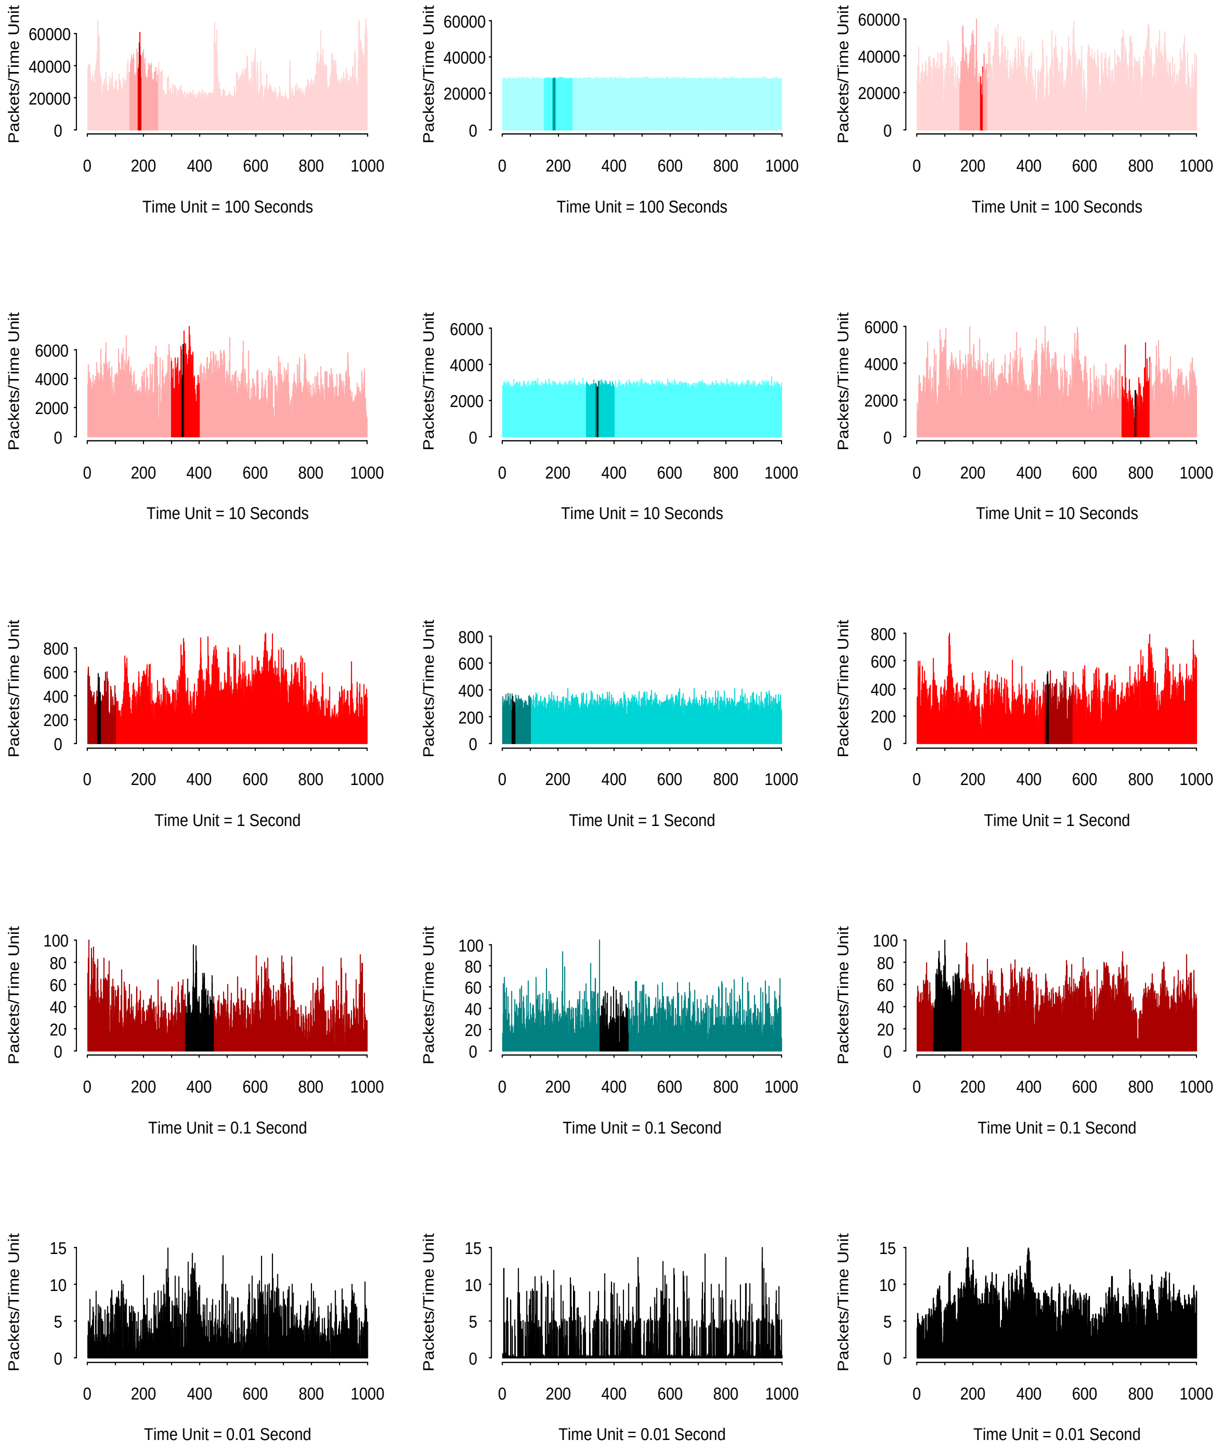

Figure 7: Actual Ethernet traffic (left column), synthetic trace generated from an appropriately chosen traditional traffic model (middle column), and synthetic trace generated from an appropriately chosen self-similar traffic model with a single parameter (right column) – on five different time scales. Different gray levels indicate the same segments of traffic on the different time scales.

simulation where we used this method to generate 27 hours worth of Ethernet-like traffic at the 10 millisecond time scale (i.e., a time series of approximately 10,000,000 observations). More precisely, our objective here was to experimentally "verify" Theorem 1 in the context of the August 1989 traffic measurements; i.e., we chose $\alpha = 1.2$ (corresponding to the estimated Hurst parameter of $H = 0.9$ that is consistent with the August 1989 data set), $M = 500$ (number of processors used to generate traffic, corresponding roughly to the number of active source-destination pairs during the observed period), and strictly alternating *ON/OFF* sources with the same α-value for the distributions of the *ON*- and *OFF*-periods. To check whether or not the resulting synthetic traffic trace "looks like" actual Ethernet LAN traffic as measured in August 1989, we plot in Figure 8 (right most column) the synthetic trace on 5 different time scales, the same way it was done in [28], the original traffic measurements (left most column), and a synthetic trace (middle column) generated from an appropriately matched batch Poisson process (the latter was taken as representative of traditional traffic modeling). As can be seen, our synthetic traffic passes the "visual" test easily, with the possible exception of the plot in the top row (the effect of the diurnal cycle in the 27 hour trace of Ethernet traffic on the 100s time scale becomes noticeable, especially because it is – by definition – not part of the stationary model that gave rise to the top right plot). On a more rigorous level, the trace also fits the data well in a statistical sense, i.e., the estimated Hurst parameter matches the one from the data. Similarly striking agreement between synthetically generated traffic and actual Ethernet LAN traces was obtained in a number of different scenarios, e.g., chosing $M = 16,000$ (close to the total number of processors on the MasPar machine), allowing for different source types (see [48]), selecting different α-values for the *ON*- and *OFF*-period distributions (including different combinations of finite/infinite variance scenarios), and generating under the i.i.d. and alternating renewal assumptions, respectively (see Section 2). Recall that the Ethernet-like behavior of the synthetically generated trace in Figure 7 has essentially been accomplished *with only one parameter*, namely the intensity α of the Noah Effect for the *ON/OFF*-periods of the traffic generated by a "typical" user. Thus Figure 7 is testimony to parsimonious modeling at its best, and proof that today's complex network traffic dynamics can be modeled in a simple manner without requiring highly parameterized mathematical models.

4.2 Performance and Protocol analysis

The practical benefits of parsimonious modeling of measured network traffic become especially apparent when focusing on the potential impacts of traffic characteristics such as the Joseph and Noah Effects on queueing and network performance, protocol analysis, and network congestion controls. Clearly, the appeal lies in the small number of physically meaningful parameters whose practical impacts need to be investigated. Starting with the empirical finding of self-similarity in Ethernet LAN traffic data reported in [27], there has been mounting evidence for the practical importance of the Hurst parameter H for traffic engineering purposes. In particular, work in [35, 9] (see also [12, 3]) demonstrates a significant difference in queueing performance (expressed in terms of the queue length distribution) between traditional (Markovian) traffic models and those exhibiting the Joseph Effect. More specifically, while the queue length distribution of the former decreases exponentially fast, that of the latter decreases much more slowly (depending on the intensity H of the Joseph Effect), namely like the tail of a Weibull distribution. In practice, not accounting for the Joseph effect at the modeling stage can lead to overly optimistic performance predictions and thus to quality-of-service requirements that are impossible to guarantee in a realistic network scenario. This observation is of particular importance in the context of the widely used concept of *equivalent bandwidth* [10], used in a number of proposed call admission control schemes. At the same time, the presence of the Joseph Effect in measured traffic does not preclude economies of scale (i.e., statistical multiplexing gains) by multiplexing a large number of such sources (see [8]).

In view of our physical explanation that the Joseph Effect in aggregate LAN traffic is caused by the Noah Effect in the individual *ON/OFF* sources that generate the aggregate stream, understanding the impacts of the Noah Effect in simple *ON/OFF* source models on queueing performance becomes essential and is likely to provide valuable new insights into questions related to the design of efficient protocols and effective controls for realistic network traffic. In fact, work is already under way that provides such new insights (e.g., see [39, 3]). For example, investigating the queue length distribution for *ON/OFF* traffic that exhibits the Noah Effect (either directly or via a corresponding M/G/1 model), these authors show that the Noah Effect gives rise to an infinite mean waiting time, i.e., to queue length distributions that decrease much slower than a Weibull distribution (i.e., the corresponding distribution obtained when aggregating many such *ON/OFF* sources). Clearly, this is bad news from the point of view of trying to keep the traffic generated by individual sources isolated from other traffic as far into the network as possible: the resulting buffer requirements at each network node and the ensuing potential delays will be overwhelming. Similarly, traffic shaping at the source may not be feasible in practice due to the naturally occurring large *ON*-periods for these sources, which in turn would require huge buckets and thus give rise to unreasonably large delays. On the other hand, the results strongly suggest the idea of statistically multiplexing a large number of sources that exhibit the Noah Effect at the earliest possible stage in the network. By doing so, the theory ([35], [3], [17]) predicts smaller buffer requirements for the network elements and hence smaller packet delays. The simulation results presented in [39] also suggest a wide range of possibilities for protocols and controls for dealing with traffic scenarios that consist of sources with different combinations of infinite variance/finite variance *ON*- and *OFF*-periods. For example, protocol design should be expected to be sensitive to and take into account knowledge about network traffic such as the presence or absence of the Noah Effect in a "typical" traffic source. However, how to effectively design protocols that take such information into account remains largely an open issue. Similarly, call admission control algorithms and congestion control schemes that incorporate information about the presence or absence of the Noah Effect at the source level and the Joseph Effect at the aggregate level have yet to be proposed and investigated (however, for some recent work on admission control schemes that take source characteristics such as the Noah Effect into account, see [20]).

5 Conclusion

Traditional *ON/OFF* source models typically assume exponential or geometric distributions for their *ON*- and *OFF*-periods (or more generally, finite variance distributions). These models are widely used and are especially popular with queueing and performance analysts because of their analytic tractability. However, in recent years, it has been recognized that multiplexing a large number of sources with such distributions results in aggregate traffic that is inconsistent with traffic measurements from working networks. On the other hand, Jain and Routhier's packet train models [19] which arose directly from traffic measurement studies, were criticized because of their lack of a clear definition of a "train", their lack of suggestions for chosing the crucial model parameters, and their lack of a physical interpretation. Motivated by the desire to provide a physical explanation for the empirically observed self-similarity property in actual network traffic, we propose in this paper to expand the range of traditional traffic modeling at the level of individual sources to account for the Noah Effect, i.e., for the ability of individual sources to exhibit characteristics that cover a wide range of time scales ("high-variability sources"). By doing so, the criticisms for the *ON/OFF* source model as well as for the packet train model are deflected. Our results in Section 2 show that the superposition of many *ON/OFF* models, each of which exhibits the Noah Effect, results in aggregate packet streams that are consistent with measured LAN traffic and exhibits the same self-similar or fractal properties as can be observed in the data. Moreover, our statistical analysis in Section 3 confirms the presence of the Noah Effect in measured Ethernet LAN traffic at the source level, and demonstrates an appealing robustness property that renders the stated objections against packet train source models irrelevant.

By being able to (i) reduce the self-similarity phenomenon for aggregate LAN traffic to properties of the individual traffic components that make up the aggregate stream, (ii) expressing the essential difference between traditional and self-similar traffic modeling in the context of the well-known *ON/OFF* source models, and (iii) identifying the Noah Effect as the main point of departure from traditional to self-similar traffic modeling, we hope to facilitate the acceptance of self-similar traffic models as viable and practically relevant alternatives to traditional models. The benefits for doing so are immediate and include parsimonious and physically meaningful models for the seemingly very complex traffic dynamics in today's networks, and new insights into problems related to the performance and analysis of protocols and network controls. We have discussed in Section 4 some of the mounting evidence for the practical importance of the Noah and Joseph Effects for network engineering, and it is safe to expect that these empirically observed traffic characteristics will play an increasingly important role in the traffic modeling and network performance work for tomorrow's high-speed networks.

Finally, there remains the question about a plausible physical explanation of the empirically observed self-similar nature of WAN traffic. Recall that while the results in Section 4 provide compelling evidence in favor of explaining the self-similar nature of aggregate LAN traffic in terms of the Noah Effect of the individual source-destination pairs that make up the aggregate packet stream, they also point toward the need for a different physical explanation when it comes to understanding the empirically observed self-similar nature of WAN traffic. To this end, we recall a construction by Cox [5], also known as the *immigration death process* or *$M/G/\infty$ model* in queueing theory. In the context of WAN traffic, Cox's construction concerns network traffic at the application level and assumes that sessions (e.g., FTP, WWW) arrive in some random (e.g., Poisson or, more generally, renewal) fashion, transmit packets in some manner (e.g., constant rate) during their "life time" (connection or holding time), and then cease transmitting packets. Cox [5] shows that the aggregate packet stream (i.e., number of packets per time unit) generated by this process is (asymptotically second-order) self-similar, provided the session holding times during which packets are transmitted exhibit the Noah Effect. In fact, rigorous mathematical convergence results that lead to fractional Brownian motion limits (in the spirit of Theorem 1) and are directly applicable to Cox's construction and important generalizations thereof have been recently obtained by Kurtz [25]. At the same time, there already exists strong empirical evidence in favor of the Noah Effect at the application level in measured WAN traffic, e.g., see [38, 22, 6]. Further work on a rigorous physical explanation of WAN traffic self-similarity, together with a comprehensive statistical analysis of WAN traffic at the application level with emphasis on the Noah Effect exhibited by the individual applications (following the original work in [38]) are currently in progress. We expect that the resulting detailed insights into the nature of WAN traffic will prove especially useful in understanding of potential interactions between networks (e.g., controls, shaping mechanisms) and traffic (e.g., *Mbone*, TCP traffic).

6 Acknowledgment

W. Willinger and M. S. Taqqu were partially supported by the NSF grant NCR-9404931. M. S. Taqqu was also partially supported by the NSF grant DMS-9404093. Part of this research was done when M. S. Taqqu was visiting Academia Sinica in Taipei. He would like to thank the Institute of Statistical Science and Dr. Hwai-Chung Ho for their hospitality. This extended version of our *ACM SIGCOMM'95* paper [53] has benefited from many helpful and specific suggestions by Vern Paxson, discussions with Joe Abate and Ward Whitt concerning the inversion problem of Laplace transforms, and constructive criticisms from the *ACM SIGCOMM'95* and *IEEE/ACM Transactions on Networking* anonymous referees. We thank them all.

References

[1] T. Berners-Lee, R. Cailliau, A. Loutonen, H. F. Nielsen and A. Secret. The World-Wide Web. *Communications of the ACM*, Vol. 37, pp. 76–82, 1994.

[2] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation.* Cambridge University Press, 1987.

[3] F. Brichet, J. Roberts, A. Simonian, and D. Veitch. Heavy Traffic Analysis of a Fluid Queue fed by On/Off Sources with Long-Range Dependence. Preprint, 1995.

[4] D. R. Cox. *Renewal Theory.* Methuen & Co., London, 1967. Science Paperback Edition.

[5] D. R. Cox. Long-range dependence: a review. In H.A. David and H.T. David, editors, *Statistics: An Appraisal*, pages 55–74. Iowa State University Press, 1984.

[6] M. E. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proceedings of the 1996 ACM SIGMETRICS. International Conference on Measurement and Modeling of Computer Systems*, May 1996.

[7] P. Danzig, S. Jamin, R. Cáceres, D. Mitzel and D. Estrin. An Empirical Workload Model for Driving Wide-Area TCP/IP Network Simulations. *Internetworking: Research and Experience*, Vol. 3, pp. 1–26, 1992.

[8] N. G. Duffield. Economies of Scale in Queues with Sources having Power-Law Large Deviation Scalings. Preprint, 1994.

[9] N. G. Duffield and N. O'Connell. Large Deviations and Overflow Probabilities for the General Single-Server Queue, with Applications. *Mathematical Proceedings of the Cambridge Philosophical Society*, 118:363–375, 1995.

[10] A. I. Elwalid and D. Mitra. Effective Bandwidths of General Markovian Sources and Admission Control of High Speed Networks. *IEEE/ACM Transactions on Networking*, Vol. 1, pp. 329–343, 1993.

[11] H. Eriksson. Mbone: The Multicast Backbone. *Communications of the ACM*, Vol. 37, pp. 54–60, 1994.

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

[13] W. Feller. *An Introduction to Probability Theory and its Applications*, volume 2. Wiley, New York, 2nd edition, 1971.

[14] R. Gusella. A Characterization of the Variability of Packet Arrival Processes in Workstation Networks. Ph.D. dissertation, Univ. of California, Berkeley, 1990.

[15] R. Gusella. A Measurement Study of Diskless Workstation Traffic on an Ethernet. *IEEE Transactions on Communications*, Vol. 38, pp. 1557–1568, 1990.

[16] R. Gusella. Characterizing the Variability of Arrival Processes with Indexes of Dispersion. *IEEE Journal on Selected Areas in Communications*, Vol. 9, pp. 203–211, 1991.

[17] D. Heath, S. Resnick and G. Samorodnitsky. Heavy Tails and Long-Range Dependence in On/Off Processes and Associated Fluid Models. Preprint, School of ORIE, Cornell University, Ithaca, NY, 1996.

[18] B. M. Hill. A Simple General Approach to Inference about the Tail of a Distribution. *The Annals of Statistics*, Vol. 3, pp. 1163–1174, 1975.

[19] R. Jain and S. A. Routhier. Packet Trains: Measurements and a New Model for Computer Network Traffic. *IEEE Journal on Selected Areas in Communications*, Vol. 4, pp. 986–995, 1986.

[20] S. Jamin, P. B. Danzig, S. Shenker and L. Zhang. A Measurement-Based Admission Control Algorithm for Integrated Services Packet Networks. *ACM/SIGCOMM Computer Communications Review*, Vol. 25, pp. 2–13, 1995. Proceedings of the ACM/SIGCOMM'95, Cambridge, MA, August 1995.

[21] S. Karlin and H. M. Taylor. *A First Course in Stochastic Processes*. Academic Press, New York, second edition, 1975.

[22] S. M. Klivansky, A. Mukherjee, and C. Song. On Long-Range Dependence in NSFNET Traffic. Preprint, 1994.

[23] M. F. Kratz and S. I. Resnick. The QQ-Estimator and Heavy Tails. Preprint, 1995.

[24] K. R. Krishnan. The Hurst Parameter of non-Markovian On-Off Traffic Sources. Preprint, 1995.

[25] T. G. Kurtz. Limit Theorems for Workload Input Models. In F. P. Kelly, S. Zachary and I. Ziedins. editors, *Stochastic Networks: Theory and Applications*, Oxford University Press, Oxford, 1996 (to appear).

[26] W. E. Leland and T. J. Ott. Unix Process Behavior and Load Balancing among Loosely-Coupled Computers. In O. J. Boxma, J.—W. Cohen, and H. C. Tijms, editors, *Teletraffic Analysis and Computer Performance Evaluation*, pp. 191–208, Amsterdam, 1986, Elsevier Science Publishers B. V.

[27] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *ACM/SIGCOMM Computer Communications Review*, Vol. 23, pp. 183–193, 1993. Proceedings of the ACM/SIGCOMM'93, San Francisco, September 1993.

[28] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, Vol. 2, pp. 1–15, 1994.

[29] S. B. Lowen and M. C. Teich. Fractal Renewal Processes Generate $1/f$ Noise. *Physical Review E*, 47:992–1001, 1993.

[30] B. B. Mandelbrot. Self-Similar Error Clusters in Communication Systems and the Concept of Conditional Stationarity. *IEEE Transactions on Communications Technology*, Vol. 13, pp. 71–90, 1965.

[31] B. B. Mandelbrot. Long-Run Linearity, Locally Gaussian Processes, H-Spectra and Infinite Variances. *International Economic Review*, Vol. 10, pp. 82–113, 1969.

[32] B. B. Mandelbrot. *The Fractal Geometry of Nature*. Freeman, New York, 1983.

[33] B. B. Mandelbrot and J. R. Wallis. Some Long-Run Properties of Geophysical Records. *Water Resources Research*, 5:321–340, 1969.

[34] K. Meier-Hellstern, P. E. Wirth, Y.-L. Yan, and D. A. Hoeflin. Traffic Models for ISDN Data Users: Office Automation Application. In A. Jensen and V. B. Iversen, editors, *Teletraffic and Datatraffic in a Period of Change, Proc. of ITC13, Copenhagen*, pp. 167–172, Amsterdam, 1991, Elsevier Science Publishers B. V.

[35] I. Norros. A Storage Model with Self-Similar Input. *Queueing Systems*, Vol. 16, pp. 387–396, 1994.

[36] V. Paxson. Growth Trends in Wide-Area TCP Connections. *IEEE Network*, Vol. 8, pp. 8–17, 1994.

[37] V. Paxson and S. Floyd. Wide-area traffic: The Failure of Poisson Modeling. *Proceedings of the ACM/SIGCOMM'94*, pp. 257–268, 1994.

[38] V. Paxson and S. Floyd. Wide-area traffic: The Failure of Poisson Modeling. *IEEE/ACM Trans. on Networking*, Vol. 3, pp. 226–244, 1995.

[39] P. Pruthi and A. Erramilli. Heavy-Tailed On/Off Source Behavior and Self-Similar Traffic. In *Proceedings of the ICC '95*, pages 445–450, Seattle, WA, 1995.

[40] S. I. Resnick. Heavy Tail Modeling and Teletraffic Data. Preprint, School of ORIE, Cornell University, Ithaca, NY, 1995.

[41] S. I. Resnick and C. Starcia. Smoothing the Hill Estimator. Preprint, 1995.

[42] G. Samorodnitsky and M. S. Taqqu. *Stable Non-Gaussian Processes: Stochastic Models with Infinite Variance*. Chapman and Hall, New York, London, 1994.

[43] D. F. Swayne, D. Cook, and A. Buja. XGobi: Interactive Dynamic Graphics in the X Window System with a Link to S. *1991 Proceedings of the Section on Statistical Graphics*, pp. 1–8, 1991.

[44] M. S. Taqqu. A Bibliographical Guide to Self-Similar Processes and Long-Range Dependence. In E. Eberlein and M. S. Taqqu, editors, *Dependence in Probability and Statistics*, pages 137–162, Boston, 1986. Birkhäuser.

[45] M. S. Taqqu. Weak Convergence to Fractional Brownian Motion and to the Rosenblatt Process. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 31:287–302, 1975.

[46] M. S. Taqqu and J. Levy. Using Renewal Processes to Generate Long-Range Dependence and High Variability. In E. Eberlein and M. S. Taqqu, editors, *Dependence in Probability and Statistics*, pp. 73–89, Boston, 1986. Birkhäuser.

[47] M. S. Taqqu and V. Teverovsky and W. Willinger. Is network traffic self-similar or multifractal? *Fractals*, 1996 (To appear).

[48] M. S. Taqqu, W. Willinger and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review*, 1997 (To appear).

[49] J. L. Teugels. Renewal Theorems when the First or the Second Moment is Infinite. *The Annals of Mathematical Statistics*, 39, pp. 1210–1219, 1968.

[50] V. Teverovsky and M. S. Taqqu. Testing for Long-Range Dependence in the Presence of Shifting Means or a Slowly Declining Trend Using a Variance-Type Estimator. Preprint, 1995.

[51] J. W. Tukey and P. A. Tukey. Strips Displaying Empirical Distributions: I. Textured Dot Strips. *Bellcore Technical Memorandum*, 1990.

[52] W. Willinger, M. S. Taqqu, W. E. Leland, and V. Wilson. Self-Similarity in High-Speed Packet Traffic: Analysis and Modeling of Ethernet Traffic Measurements. *Statistical Science*, 10:67–85, 1995.

[53] 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. *Computer Communications Review*, 25, pp. 100–113, 1995. Proceedings of the ACM/SIGCOMM'95, Cambridge, MA, August 1995.

[54] W. Willinger, M. S. Taqqu and A. Erramilli. A Bibliographical Guide to Self-Similar Traffic and Performance Modeling for Modern High-Speed Networks. In F. P. Kelly, S. Zachary and I. Ziedins. editors, *Stochastic Networks: Theory and Applications*, Oxford University Press, Oxford, 1996 (to appear).