Preprint 1996. Appears on pages 27-53 in the book:
"A Practical Guide To Heavy Tails: Statistical Techniques and Applications"
Robert Adler, Raisa Feldman and Murad S. Taqqu, editors. Birkhauser, Boston, 1998

Self-Similarity and Heavy Tails: Structural Modeling of Network Traffic

Walter Willinger, Vern Paxson and Murad S. Taqqu [1] [2] [3]

Abstract

High-resolution traffic measurements from modern communications networks provide unique opportunities for developing and validating mathematical models for aggregate traffic. To exploit these opportunities, we emphasize the need for *structural* models that take into account specific physical features of the underlying communication network structure. This approach is in sharp contrast to the traditional *black box* modeling methodology from time series analysis that ignores, in general, specific physical structures. We demonstrate, in particular, how the proposed structural modeling approach provides a direct link between the observed self-similarity characteristic of measured aggregate network traffic, and the strong empirical evidence in favor of heavy-tailed, infinite variance phenomena at the level of individual network connections.

1. Introduction

Recent empirical studies of high-resolution traffic measurements from a variety of different working communication networks (e.g., see [LTWW1, LTWW2, PF1, PF2]) have provided ample evidence that actual network traffic is *self-similar* or *fractal* in nature, i.e., bursty over a wide range of time scales. This observation is in sharp contrast to commonly made modeling choices in today's traffic engineering theory and practice, where exponential assumptions still dominate and are only able to reproduce the bursty behavior of measured traffic either on a pre-specified time scale or over a very limited range of time scales. That an observer can easily distinguish between traffic patterns predicted by currently used traffic models and actual, measured traffic traces (data sets) from today's

---

[1] W. Willinger and M. S. Taqqu were partially supported by the NSF grant NCR-9404931. M. S. Taqqu was also supported by the NSF grant DMS-9404093 at Boston University. Vern Paxson was supported by the Director, Office of Energy Research, Scientific Computing Staff, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

[2] *AMS Subject classification:* 60K30, 60G18, 60E07.

[3] *Keywords:* Long-range dependence, infinite variance, limit theorems, Pareto, stable distribution, computer network traffic models, local area network, wide area network.

networks, challenges traditional approaches to traffic modeling. The purpose of this paper is to demonstrate how the self-similar nature of network traffic at the macroscopic level (i.e., the aggregation of traffic generated by all active hosts on the network) leads to new insights into the traffic dynamics at the microscopic level (i.e., the traffic patterns generated by the individual hosts).

To this end, we consider two of the most commonly encountered network environments, *local area networks (LANs)* and *wide area networks (WANs).* Instead of applying the traditional *black box* approach often used in time series analysis, we focus on *structural* models. These take into account specific features of the underlying network structure and hence provide a physical explanation for observed phenomena such as self-similarity.

LANs were introduced in the mid-1970's to interconnect data processing equipment (host computers, file servers, printers etc.) in office and R&D environments, or within university departments. One of the most popular LAN technologies is Ethernet (e.g., see [MB]), and measured Ethernet traces were used in the original studies by Leland et al. [LTWW1, LTWW2] to demonstrate the self-similar nature of LAN traffic. Here we show how LAN traffic self-similarity gives rise to structural models that can be reduced to simple *ON/OFF* sources (also known as packet trains [JR]) with the distinctive feature that their *ON*- and/or *OFF*-periods are heavy-tailed with infinite variance. Convergence results for such *ON/OFF* processes (and superpositions thereof) provide a direct link between self-similar characteristics at the macroscopic level and the heavy-tailed phenomena observed at the microscopic level, that is, between the aggregate traffic stream and the traffic patterns displayed by individual source-destination pairs. We summarize in this paper the main features of the structural modeling approach to LAN traffic developed in [WTSW, WTSW2], and validate the underlying assumptions against an additional set of Ethernet LAN source-destination pairs.

In contrast to LANs, WANs provide interconnectivity between users (e.g., host machines on different LANs) that reside, in general, in different geographic regions. The best-known WAN is the worldwide *Internet*, a global network connecting tens of millions of hosts and users. Evidence of WAN traffic self-similarity has been reported in recent studies by Paxson and Floyd [PF1, PF2], who, in an analysis of a number of different WAN traffic traces, showed the inadequacy of traditional exponential (Poisson) traffic models in describing many crucial aspects of WAN traffic behavior. Subsequently, a number of attempts have been made at providing structural models for WAN traffic, including the same reduction to simple *ON/OFF* models for individual source-destination pairs as mentioned earlier for LAN traffic (e.g., see [WTSW2, WTSW]), and a description of WAN traffic at the level of individual applications, e.g., TELNET, FTP, and HTTP (see [PF2, CB]). Here we summarize the different attempts, and report on recently obtained limit theorems as well as on new statistical evidence that support a structural modeling approach for WAN traffic connecting the empirically observed self-similarity at the macroscopic level directly with infinite variance phenomena at the microscopic level. In this case, microscopic refers to the level of individual applications.

Thus, for LANs as well as for WANs, self-similarity of the aggregate network traffic results directly from structural models that mimic traffic dynamics at lower network levels and identify at those levels traffic characteristics that can be validated against actual high-resolution traffic measurements. The resulting models are simple and parsimonious, provide engineering insights and possess a number of desirable properties that are robust under constantly changing network conditions. They are therefore of considerable practical relevance for a wide range of network engineering tasks.

The paper is organized as follows. In Section 2, we focus on LAN traffic, introduce some commonly used terminology, and follow essentially [WTSW2] in deriving and validating a structural modeling approach for aggregate LAN traffic. Section 3 outlines the derivation of a structural model for WAN traffic, and includes a relevant limit theorem for self-similar processes as well as preliminary but persuasive empirical evidence in favor of the proposed approach. In Section 4, we discuss future work and illustrate the practical importance of structural models for the traffic engineering of today's high-speed networks.

2. Self-Similarity and Heavy Tails in LAN Traffic

We describe an *ON/OFF* model for the traffic transmitted between two typical LAN hosts and validate it by analyzing Ethernet LAN traces at the level of individual source-destination pairs.

2.1 Self-similarity

There are a number of different, not equivalent, definitions of self-similarity. The standard one states that a continuous-time process $Y = \{Y(t), t \geq 0\}$ is *self-similar* (with self-similarity parameter H) if it satisfies the condition:

$$Y(t) \stackrel{d}{=} a^{-H} Y(at), \quad \forall t \geq 0, \quad \forall a > 0, \ 0 < H < 1, \tag{2.1}$$

where the equality is in the sense of finite-dimensional distributions. The canonical example of such a process is Fractional Brownian Motion (see [ST, Be]), Brownian Motion if $H = 1/2$. While a process Y satisfying (2.1) can never be stationary (stationary requires $Y(t) \stackrel{d}{=} Y(at)$), Y is typically assumed to have stationary increments.

A second definition of self-similarity, more appropriate in the context of standard time series theory, involves a stationary sequence $X = \{X(i), i \geq 1\}$. Let

$$X^{(m)}(k) = 1/m \sum_{i=(k-1)m+1}^{km} X(i), \qquad k = 1, 2, \ldots, \tag{2.2}$$

be the corresponding aggregated sequence with level of aggregation m, obtained by dividing the original series X into non-overlapping blocks of size m and averaging over each block. The index, k, labels the block. If X is the increment process of a self-similar process Y defined in (2.1), that is, $X(i) = Y(i+1) - Y(i)$, then for all integers m,

$$X \stackrel{d}{=} m^{1-H} X^{(m)}. \tag{2.3}$$

A stationary sequence $X = \{X(i), i \geq 1\}$ is called *exactly self-similar* if it satisfies (2.3) for all aggregation levels m. This second definition of self-similarity is closely related to the first, with $mX^{(m)}(\cdot)$ corresponding to $Y(a\cdot)$.

A stationary sequence $X(i), i \geq 1$ is said to be *asymptotically self-similar* if (2.3) holds as $m \to \infty$. Similarly, we call a covariance-stationary sequence $X(i), i \geq 1$ *exactly second-order self-similar* or *asymptotically second-order self-similar* if $m^{1-H}X^{(m)}$ has the same variance and autocorrelation as X, for all m, or as $m \to \infty$. Note that for a Gaussian process, the last two definitions (i.e., self-similarity and second-order self-similarity) are equivalent.

2.2 LANs, LAN Traffic and Ethernet LAN Traffic Measurements

The last ten years have seen a tremendous growth in the number of LANs, reflecting the need to link users together and to offer connectivity to common resources such as file servers and printers. The 10 megabits/sec (Mbps) Ethernet, a multi-access system for local computer networking with distributed control has been and remains the workhorse LAN technology. By using properly instrumented monitoring hardware, it is possible to record the time stamp and header information of every single (complete) packet that is put on the monitored Ethernet cable by any of the host stations. Such high-resolution Ethernet LAN traffic measurements over week-long periods have been reported in [FL], where a “typical” day results in about 20-30 million Ethernet packets, or about 2 gigabytes worth of data.

By treating the recorded Ethernet packets as black boxes and only using the time stamp information, Leland et al. [LTWW1, LTWW2] showed that measured aggregate Ethernet LAN traffic (i.e., number of packets or bytes sent over the Ethernet by all active host stations per time unit), with its mean subtracted, is consistent with second-order statistical self-similarity; that is, Ethernet LAN traffic measured over microseconds and seconds exhibits the same second-order statistics as Ethernet LAN traffic measured over minutes or over even larger time scales. Intuitively, this scale invariance of measured Ethernet LAN traffic manifests itself in the *absence* of a characteristic *burst length*; Ethernet traffic is bursty on all (or a wide range of) time scales, and plotting it over different time scales results in similar-looking pictures, all of which feature a distinctive “burst-within-burst” structure (for details, see [LTWW2]).

When trying to “explain” this empirically observed self-similarity in terms of simpler quantities, a structural modeling approach has been proposed in [LTWW2] and studied in greater detail in [WTSW2] (for related work, see also [LT, WTSW, HRS]). In short, by using the time stamp information as well as the Ethernet source and destination addresses contained in the recorded header information of each packet seen on the Ethernet, it is possible to separate the aggregate Ethernet LAN traffic into the individual components representing traffic flows between each active pair of host computers, or *source-destination pairs*. At the level of individual source-destination pairs, simple traffic models such as *ON/OFF sources* or *packet train models* have been very popular. Informally, these models assume that a source alternates between an "active" state (or *ON*-period) and "idle" state (or *OFF*-period). During *ON*-periods, packets are sent at a constant rate, and during *OFF*-periods, no packets are transmitted. The group of packets sent during an *ON*-period are termed a "train," and the lull between two trains (i.e., the *OFF*-period) is termed the "intertrain gap." Traditionally, the successive *ON*-periods as well as the successive *OFF*-periods are assumed to be independent and identically distributed (i.i.d.) and independent from each other. Thus, the only stochastic elements in describing *ON/OFF* sources are the distributions that govern the lengths of the *ON*- and *OFF*-periods, respectively. In Section 2.3 below, we provide the details of a limit theorem that states that the superposition of many such *ON/OFF* sources will capture the empirically observed self-similar nature of measured aggregate Ethernet LAN traffic, provided that the distribution of either the *ON*- or *OFF*-periods of an individual source-destination pair has infinite variance. The assumptions under which Theorem 2.1 below holds will be validated in Section 2.3 against measured traffic flows of individual Ethernet source-destination pairs.

2.3 A Limit Theorem for Aggregate Traffic

We consider first a single *ON/OFF* 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. (Such a process is sometimes referred to as an *alternating renewal* process.)

Suppose now that there are M such i.i.d. *ON/OFF* sources. Since each source m 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.

In considering the possible distributions, we are motivated by the empirically derived fractional Brownian motion model for measured aggregate cumulative packet traffic in Ethernet LANs in [LTWW2, No], or equivalently, by its increment process, the so-called fractional Gaussian noise model for aggregate traffic (i.e., number of packets per time unit). They model deviations from the mean value. Accordingly, we want to choose the *ON* and *OFF* 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 [ST]. The parameter $1/2 \leq H < 1$ is called the *Hurst parameter* or the *index of self-similarity.*

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

$$f_1(x),\quad F_1(x) = \int_0^x f_1(u)du,\quad \overline{F}_1(x) = 1 - F_1(x),\quad \mu_1,\quad \sigma_1^2$$

denote the probability density function, cumulative distribution function, complementary (or tail) distribution, mean length and variance of an *ON*-period, and let $f_2, F_2, \overline{F}_2, \mu_2, \sigma_2^2$ correspond to an *OFF*-period. Assume that as $x \to \infty$, both of the following statements hold:

$$\text{either}\quad \overline{F}_1(x) \sim \ell_1 x^{-\alpha_1} L_1(x) \text{ with } 1 < \alpha_1 < 2 \quad\text{or}\quad \sigma_1^2 < \infty,$$

and

$$\text{either}\quad \overline{F}_2(x) \sim \ell_2 x^{-\alpha_2} L_2(x) \text{ with } 1 < \alpha_2 < 2 \quad\text{or}\quad \sigma_2^2 < \infty,$$

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 $t > 0$. We also assume that either probability densities exist or that $F_j(0) = 0$ and F_j is non-arithmetic. 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. $\overline{F}_j(x) = K^{\alpha_j} x^{-\alpha_j}$ for $x \geq K > 0$, $1 < \alpha_j < 2$ and equal 0 for $x < K$, satisfying the first clause of the "either–or"; or it could be exponential, satisfying the second clause (finite variance). 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.

Before stating the main result, we introduce the following normalization factors and limiting constants that prove convenient in formulating a single limit theorem that holds under the different assumptions on F_1 and F_2. 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 $\alpha_{\min} = \alpha_1 = \alpha_2$,

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

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

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

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. Under the conditions stated above, the following holds:

Theorem 2.1 *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,

$$\mathcal{L}\lim_{T\to\infty} \mathcal{L}\lim_{M\to\infty} \frac{1}{T^H L^{1/2}(T) M^{1/2}} \left(W^*_M(Tt) - \frac{\mu_1 MTt}{\mu_1+\mu_2} \right) = \sigma_{\lim} B_H(t), \quad (2.4)$$

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

For a proof and for discussions of special cases and generalizations of Theorem 2.1, see [WTSW2]. Heuristically, Theorem 2.1 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 stated in [WTSW2], 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

$$\overline{F}_j(x) \sim \ell_j x^{-\alpha_j} L_j(x), \qquad \text{as } x \to \infty,\ 1 < \alpha_j < 2 \qquad (2.5)$$

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.

2.4 Statistical Analysis of Ethernet LAN Traffic at the Level of Individual Source-Destination Pairs

In view of Theorem 2.1, in order to validate the proposed structural modeling approach for Ethernet LAN traffic in terms of the superpositions of many *ON/OFF* source-destination traffic flows, one must check whether measured traffic data are consistent with (i) the *ON/OFF* traffic model assumption for

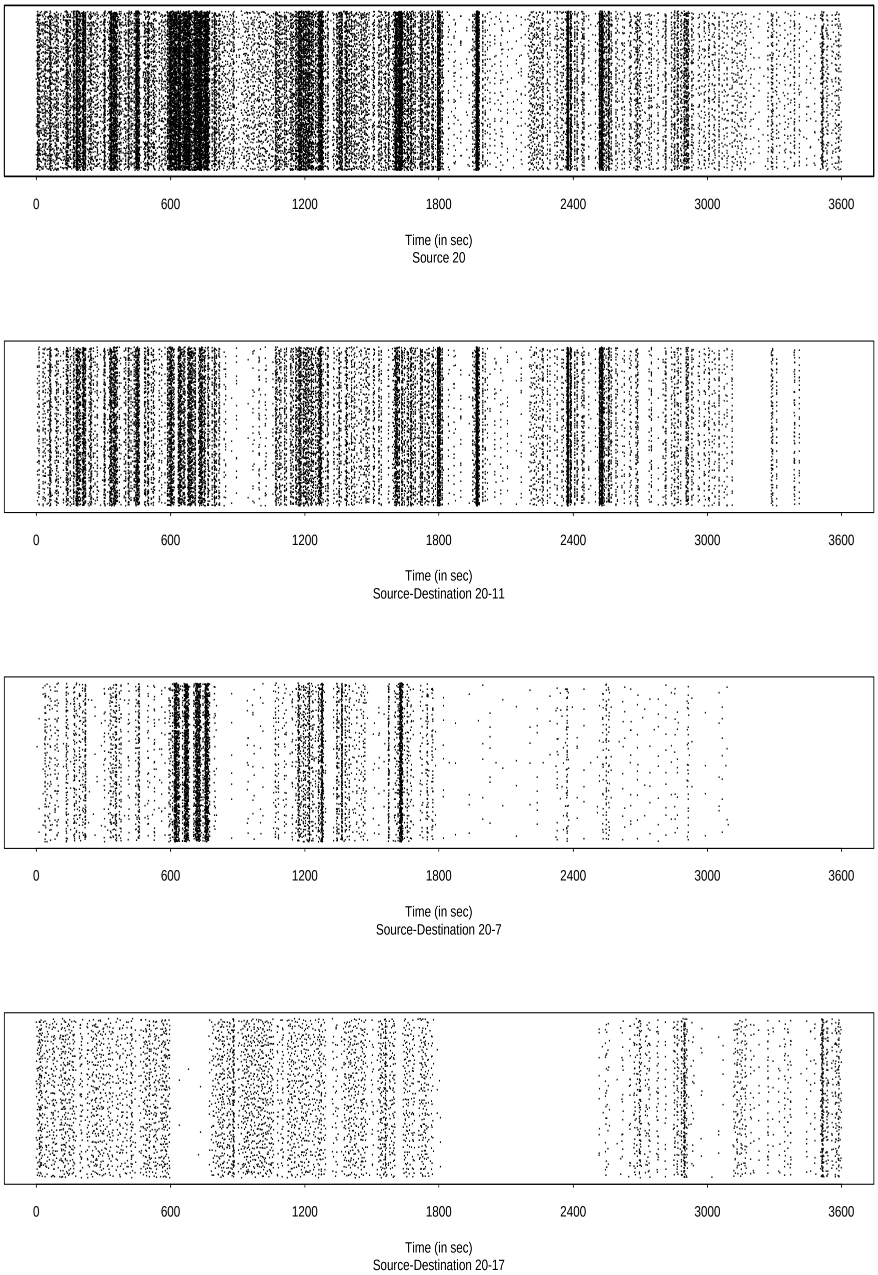

Figure 1: Textured plots of packet arrival times for (top to bottom) source 20 and source-destination pairs 20-11, 20-7, and 20-17.

individual sources or source-destination pairs and (ii) the crucially important heavy-tailed property (2.5) for the corresponding *ON*- or *OFF*- periods. To this end, we make use of a data set known as the "busy hour" of a trace of August 1989 Ethernet LAN traffic. This data set was presented, analyzed and shown 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) in [LTWW1].

We first partition this hour-long sequence of Ethernet packets generated by all active Ethernet hosts into their distinct Ethernet source and destination addresses. We find that during the given hour 105 Ethernet hosts sent or received packets over the network. With regard to individual source-destination pairs, only 748 out of $104 \cdot 105 = 10,920$ possible pairs (about 6.8%) were actually sending or receiving packets.

In order to assess the *ON/OFF* nature of traffic generated by an individual source or source-destination pair, we use *textured dot strip plots* or simply *textured plots*, originally introduced in [TT]. 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. Figure 1 shows four textured plots associated with source 20 (other sources result in similar plots), each point in the plots representing the time of a packet arrival. Source 20 contributed 2.67% to the overall number of packets and sent data to 13 different destinations. The top plot in Figure 1 represents the textured plot corresponding to the arrival times of all packets originating from source 20 (there are 37,582 packets), and the subsequent 3 panels result from applying the textured plot technique to the arrival times of all packets originating from source 20 and destined for hosts 11, 7, and 17, respectively. These three source-destination pairs are responsible for 20,152, 7,497, and 5,511 packets, respectively, and make up about 88% of all the packets generated by source 20. Figure 1 supports the observations that the traffic generated by a reasonably active individual Ethernet host (e.g., source 20, top panel) is itself bursty in nature, and that a "typical" individual source-destination pair (e.g., source-destination pair 20-11, second panel) exhibits an apparent bursty or *ON/OFF* structure. However, Figure 1 also shows that there is a clear ambiguity associated with identifying a "typical" burst length or *ON/OFF*-period in the traffic patterns generated by an individual Ethernet host or an individual Ethernet source-destination pair. We will show below that this ambiguity is a natural consequence of the heavy-tailed property (2.5) and can be exploited when inferring property (2.5) from a given set of data.

To determine whether or not a given data set is consistent with the heavy-tailed property (2.5) and, if in the affirmative, to estimate a range for the

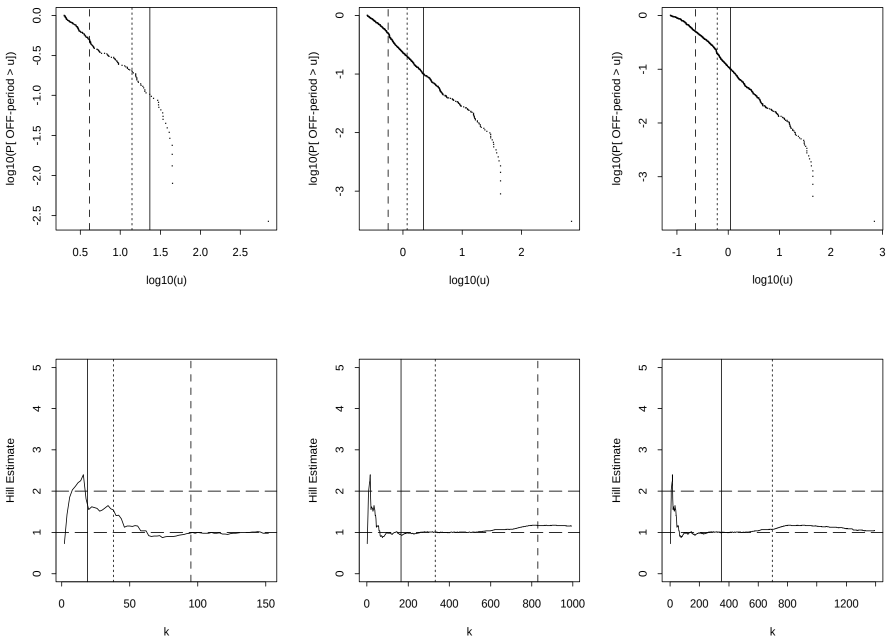

Figure 2: Robustness property of the *OFF*-periods (for source-destination pair 20-17), for threshold values (from left) $t = 2.0$ sec, .25 sec, .075 sec.

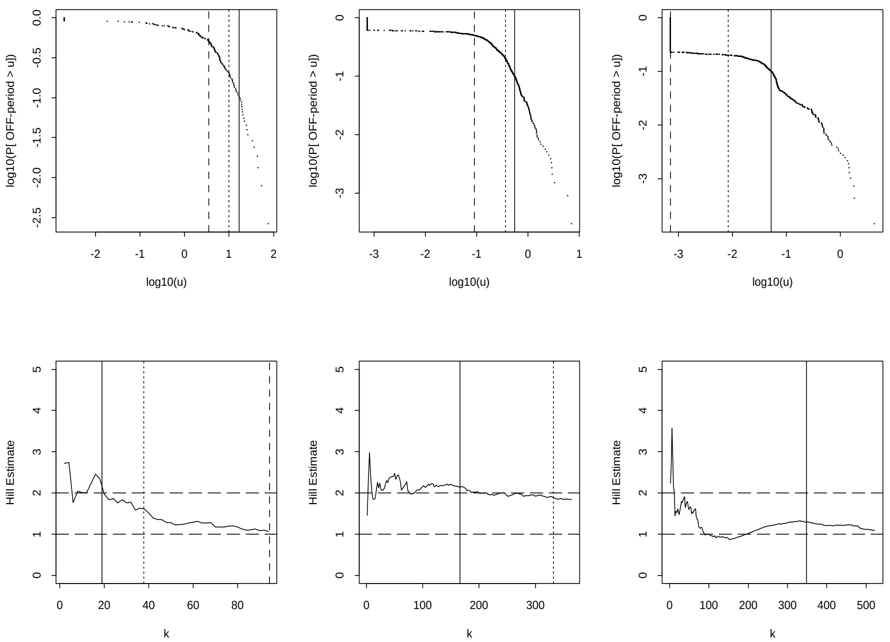

Figure 3: Robustness property of the *ON*-periods (for source-destination pair 20-17), for threshold values (from left) $t = 2.0$ sec, .25 sec, .075sec.

parameter α that characterizes the power law decay in (2.5), we make extensive use of *complementary distribution plots* (related to the *qq-plot* method discussed in [KR]) and *Hill plots* (derived from Hill's method, e.g., see [Hi, RS]). We assume that the reader is familiar with both methods, but would like to emphasize that although both methods are reasonably well understood in theory, they can perform quite erratically in practice. Since theoretical results for these procedures (e.g., confidence intervals for the Hill estimator) are known to hold only under conditions that often cannot be validated in practice (see for example [Re]), it is preferable to use data-intensive heuristics. As a result, we typically end up with strong empirical evidence on whether or not property (2.5) holds but without precise point estimates for α.

As an illustration, consider the traffic generated by the Ethernet source-destination pair 20-17 (Figure1, bottom panel); the traffic streams from other source-destination pairs yield similar results (see below). Given its packet arrival process, it is natural to define an *OFF*-period to be *any interval of length longer than t seconds that does not contain any packet arrival*; this in turn, defines the *ON*-periods unambiguously. Of course, the ambiguity experienced earlier when trying to "eyeball" *ON/OFF*-periods based on the textured plots in Figure 1 has now been replaced by the equivalent problem of selecting the "right" threshold value t in the formal definition of an *OFF*-period. In the packet train terminology, the problem is to decide in a coherent manner on the "appropriate" intertrain gap, i.e., on deciding when the "departure" of the previous train took place and when the "arrival" of the next one occurs. However, we show below that *as far as property (2.5) 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, property (2.5) is robust under a wide range of choices for the threshold value t. In the case of the *OFF*-periods, the reason is the well-known scaling property of distributions that satisfy the power tail condition (2.5): if the distribution of the random variable U satisfies (2.5) and t denotes a threshold value, then for sufficiently large u, t with $u > t$,

$$P(U > u \mid U > t) \sim (\frac{u}{t})^{-\alpha}, \quad 1 < \alpha < 2. \tag{2.6}$$

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. In the case of the Ethernet host pair 20-17, this is illustrated in Figure 2, where we show the complementary distribution plots (top row) and corresponding Hill plots (bottom row) for three different ways of defining *OFF*-periods. More specifically, we chose t-values that span 3 orders of magnitude, namely $t = 2.0$ sec (left column, 190 observations), $t = 0.25$ sec (middle column, 1,658 observations), and $t = 0.075$ sec (right column, 3,484 observations). The three vertical lines in the complementary distribution plots in Figure 2 indicate that 10%, 20% and 50% of all the data points are to the right of the solid, dotted and dashed line, respectively; in the

Hill plots, the solid, dotted and dashed lines indicate that 10%, 20% and 50% of the largest order statistics have been included in the computation of the Hill estimator. Figure 2 confirms the robustness of property (2.5) under the different choices of t, with an estimated α-value between 1.0 to 1.3.

To explain the robustness property of the *ON*-periods with respect to a wide range of threshold values t, recall that a t-*ON* period (i.e., an *ON*-period that was obtained using the threshold value t) typically consists of a number of s-*ON* and s-*OFF* periods where $s < t$. However, if the s-*ON/OFF* periods satisfy relation (2.5), then so does their sum. Subsequently, when interested in the power-law decay for the distribution of the *ON*-periods, fragmentation into smaller *ON/OFF*-periods as the threshold value t decreases should have minor impact and suggests that the *ON*-periods are generally robust under a wide range of t-values. This property is illustrated in Figure 3, where we consider the *ON*-periods corresponding to the three different choices of *OFF*-periods that were used in Figure 2 for the pair 20-17 (i.e., $t = 2.0$ sec, 0.25 sec, 0.075 sec). As can be seen, the estimate of the α-parameter that characterizes the power law decay in (2.5) for the *ON*-periods remains within the interval $[1, 2]$, even though t varies from seconds to 100 milliseconds to milliseconds.

Source-destination pairs other than 20-17, as well as individual Ethernet hosts, yield similar results and provide convincing evidence that property (2.5) for the *ON/OFF*-periods of the individual sources and source-destination pairs is robust under a wide range of choices of threshold values. The results of a full-fledged analysis of the 181 most active source-destination pairs (out of a total of 748 active pairs) that make up 93% of all the packets in this data set, are shown in Figure 4. Using the same thresholding procedure as in Figures 2 and 3, we checked each of the 181 source-destination pairs for the presence or absence of property (2.5) in their corresponding sequences of *ON*- and *OFF*-periods. We determined, for each of the 181 source-destination pairs, 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, according to whether the resulting α-estimates fall within the intervals $(0, .85), (.75, 1.35), (1.25, 1.75), (1.65, 2.25)$ or $(2.25, 2.75)$. These represent the 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 example, we classified the *ON* and *OFF* periods of source-destination pair 20-17, shown in Figures 2 and 3, as "somewhere in the middle of the interval (1,2)" and "around 1.0." For the vast majority of source-destination pairs, the categorization process worked well, especially for the *OFF*-periods. While not all *ON/OFF*-periods fitted this framework, the number of inconclusive cases was low, under 10%. In Figure 4 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

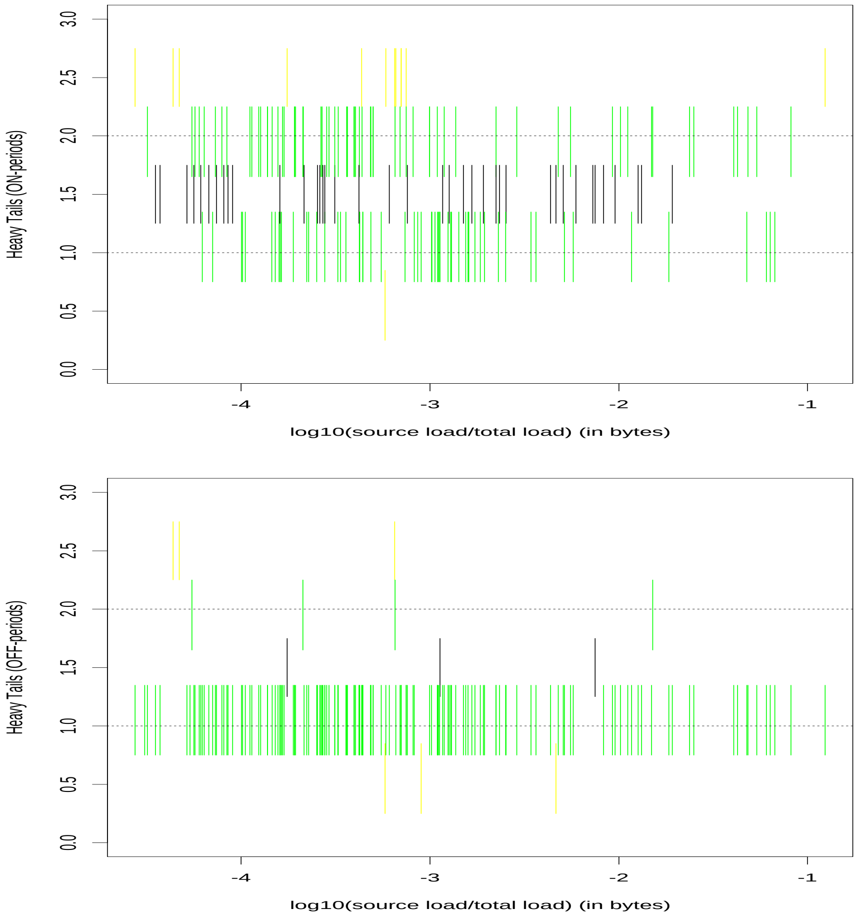

Figure 4: Summary plot of ranges for the α-estimates for the *ON*-periods (top) and *OFF*-periods (bottom) of the 181 most active source-destination pairs, as a function of their loads (in bytes, on log scale).

pretty much the whole interval $(1, 2)$. In comparison, the bottom plot in Figure 4 shows that in the case of the *OFF*-periods, α-estimates in the lower part of the interval $(1, 2)$ clearly dominate the picture. Thus, the evidence of infinite-variance *ON*-periods is not as consistently strong as that for the *OFF*-periods. Nevertheless, the mathematical results in Section 2.3 still apply because they allow the α-values for the *ON*- and *OFF*-periods to be different, or even one of the distributions to have finite variance.

In summary, Figure 4 provides strong statistical evidence in favor of our proposed physical explanation for the empirically observed self-similarity property of aggregate Ethernet LAN traffic. Our analysis shows that the data at the source-destination level are consistent with the *ON/OFF* model and that the distributions of the corresponding *ON/OFF*-periods tend to satisfy the heavy-tailed property (2.5). One may question the assumption made in Theorem 2.1 that W_M^* is the superposition of *independent ON/OFF* sources. Hosts that compete for available (finite) network resources and protocols that determine which host machine can send how many bytes at what times are bound to introduce dependencies among the individual source-destination pair; in fact, a careful examination of Figure 1 shows segments in the textured plots of source-destination pairs 20-11 and 20-7 with interchanging burst- and idle-period, indicating periods when there are obvious interactions between individual source-destination pairs. It is likely, however, that the conclusion of Theorem 2.1 will continue to hold if the sources are weakly dependent.

3. Self-Similarity and Heavy Tails in WAN Traffic

Research on structural models for wide-area network traffic is still at the preliminary stage. We focus on Internet-related traffic, consider a number of well-known applications (e.g., FTP, HTTP, TELNET), discuss structural modeling approaches and comment on their validation.

3.1 WANs, WAN Traffic and WAN Traffic Measurements

Wide-area networks (or *WANs*) arose in the 1970's as a means for interconnecting computers located in geographically distant locations. WANs are often a collection of distinct (independently administered) networks using possibly different interconnection hardware. The best known example is the Internet, which has experienced sustained explosive growth, increasing in size by 80–100% per year for well over a decade. Presently, the Internet encompasses some 100,000 different networks and 10,000,000 hosts.

WANs differ from LANs in a number of fundamental ways. They are much more heterogeneous, making it difficult to predict what traffic conditions one might encounter in any particular situation. Another crucial difference is that speed-of-light limitations have an immense impact on WAN engineering, because the *time constants* associated with obtaining feedback on network conditions are measured from tens of milliseconds to seconds, instead of microseconds as in the case of LANs. This makes it much more difficult for WAN applications to adapt to current network conditions, and can create *congestion*, that is, a performance degradation due to the overloading of some of the components in the end-to-end chain of networks between the two WAN hosts. If a component has no more buffer memory for temporarily storing the data packets it receives, it will drop them instead of forwarding them. Thus, the reliable transfer of data across a WAN requires a sophisticated "transport" protocol. This protocol must ensure that data packets are retransmitted when lost, and must avoid unnecessary retransmission which would further add to congestion.

For the Internet, the dominant transport protocol is the Transmission Control Protocol (TCP). The WAN traffic traces discussed below were obtained by recording the header of every TCP packet[4] at network links connecting various institutions to the Internet "backbone." TCP is not the only protocol in use on the Internet; indeed, a growing proportion of the traffic is "multicast" (often used for sending digitized audio and video), in which a single sender transmits to multiple receivers. These relatively new applications, however, have been less studied than the more well-established Internet applications that use TCP, so we confine ourselves here to the latter.

3.2 WAN Traffic at the Application Level

The profile of dominant Internet applications keeps changing over time. Currently, the most significant applications on the Internet are: file transfer (FTP); structured information retrieval (HTTP; the "World Wide Web"); remote login (TELNET); electronic mail; and Network News. Just four years ago, HTTP traffic was virtually nonexistent. There also exists no "typical" WAN application "mix": the dominant WAN applications vary greatly from site to site. For example, a recent study [Pa] found that in January 1991, the proportion of traffic volume (total bytes) due to the Network News application at two different California sites varied from 15% at one of them to 60% at the other. Thus in general one must be wary of assuming that a particular WAN link trace reflects "typical" traffic. One way of dealing with this difficulty is to partition the traffic according to the different applications. This is easy to do when examining a trace of WAN traffic, because the TCP packet header includes a "port" number that identifies different applications. Here we will concentrate exclusively on TELNET, FTP and HTTP.

FTP and HTTP are "bulk transfer" applications, whose main task is to move a predetermined amount of data from one Internet host to another. While in a LAN environment bulk transfer is relatively straightforward, in a WAN setting it becomes significantly more complex, because of transient congestion and the dynamics of TCP. As a result, useful traffic models associated with WAN bulk transfers are rare and ad-hoc, at best. Based on the recent empirical evidence of self-similar features in measured aggregate WAN traffic (see [PF1, PF2]), we propose here a structural modeling approach for WAN bulk transfer that connects WAN traffic characteristics at the macroscopic (i.e., aggregate) level to the microscopic (i.e., application) level by focusing on typical bulk transfer features such as arrival patterns, the amount of data transferred, or duration (which depends not only on the amount of data transferred, but also on the network conditions encountered during the transfer). The resulting structural models suggest that for some problems of practical interest, capturing the fine details of bulk transfer (that is, the precise manner in which packets are transmitted during a session, which is mainly determined by transient congestion within the WAN and the dynamics of TCP) can be avoided; note however that these

[4] Some of the analysis is based on data sets that recorded only the "connection-begun" and "connection-finished" packets that delimit each TCP connection.

fine details are of great importance when studying, for example, how network controls affect traffic.

In contrast to FTP and HTTP, TELNET is an "interactive" application. The packets sent from the host initiating the connection to the receiving host are determined by the pattern of keystrokes made by the TELNET user. As such, these patterns can be expected to be fairly robust in the presence of widely varying networking conditions, and our structural modeling approach for TELNET traffic will focus on identifying "typical" features in these patterns.

3.3 Some Limit Results for Aggregate WAN Traffic

A natural modeling approach for aggregate WAN traffic is based on the idea of "separation of time scale." That is, there exist two distinct processes of interest: the start times of *sessions* (where a session consist of one or more related network connections), and the arrival process of *packets* within a session. One can observe empirically this separation when plotting Figure 1 for "typical" source-destination pairs in a WAN environment (not shown here): transmission starts at some random point in time ("start of a session"), packets are transmitted (in some bursty fashion) for some time, and then transmission stops ("end of session") until the next session begins. While session arrivals can, in general, be identified unambiguously, defining a "typical" burst of packets within a session is as ambiguous as for a host pair in a LAN environment (see Section 2). In the following, we present some known approaches for traffic modeling that explicitly mimic this two-stage procedure. We then comment on their relevance for capturing empirically observed WAN traffic characteristics.

We first recall a construction due to Cox [Co], also known as an *immigration death process* or *$M/G/\infty$ queueing model.* Cox's construction has been suggested for modeling modern communication network traffic in [LTWW2], and was considered explicitly for WAN traffic in [PF2]. In the context of WAN traffic, the Cox construction assumes that sessions (e.g., FTP, HTTP, TELNET) arrive according to a Poisson process, transmit packets deterministically at a constant rate (i.e., in a "fluid" fashion) during their "lifetime" or session length, and then cease transmitting packets. Note that once the Poisson nature of session arrivals is taken for granted (for further discussion of Poisson session arrivals, see Section 3.4), the only stochastic element left unspecified is the distribution of session lengths or durations. We shall choose it so as to capture in a parsimonious manner the empirically observed long-range dependence property of aggregate WAN traffic (see [PF2]), which in turn is related to statistical self-similarity. More precisely, working in discrete time, let X_n denote the the number of customers in the system at time n in the $M/G/\infty$ model, or equivalently, the total number of packets generated by all the sessions that are active at time n (assuming that packets are transmitted one per time unit during the lifetime of a session). Let $(f_n)_{n\geq 1}$, $F(n) = \sum_{k\leq n} f(k)$, $\overline{F} = 1 - F$ and μ be, respectively, the probability mass function, cumulative distribution function, tail distribution and mean of the session length and assume that as $n \to \infty$, F satisfies the heavy-tailed property (2.5), that is,

$$\overline{F}(n) \sim n^{-\alpha} L(n), \quad \text{as } n \to \infty,\ 1 < \alpha < 2 . \tag{3.7}$$

Under these conditions, Cox [Co] (see also [PF2]) obtained the following result:

Theorem 3.2 *The aggregate packet process $X = (X_n : n = 0,\ 1,\ 2, \ldots)$ exhibits long-range dependence. More precisely, denoting by $r(k)$ the autocorrelation function of X, we have*

$$r(k) = \mu^{-1} \sum_{n=k}^{\infty} \overline{F}(n) \sim Ck^{1-\alpha} L(k),\ \text{as } k \to \infty, \tag{3.8}$$

for some constant $C > 0$. Moreover, the degree of long-range dependence (i.e., Hurst parameter) is given by $H = (3 - \alpha)/2$.

The main ingredient is the heavy-tailed property (3.7) of the session durations. Intuitively, this property implies that the length of a "typical" session is highly variable (*infinite variance*), i.e., exhibits fluctuations over a wide range of time scales. This basic characteristic at the application level manifests itself at the network level through property (3.8). This property implies that the aggregate traffic process X is *second-order asymptotically self-similar*, that is, when viewed over sufficiently large time scales, the second-order statistical properties of X remain essentially unchanged, and the traffic looks "similar" over a wide range of time scales.

While appealing in its simplicity, Cox's construction suffers from a number of shortcomings that limits its immediate applicability to WAN traffic modeling. First, the Poisson nature of session arrivals often turns out to be too restrictive in practice. Second, and more importantly, the applications that currently contribute major portions to WAN traffic, namely FTP and HTTP, are known (see [PF2, CB]) to transmit their packets *not* at a constant rate, but in a highly bursty manner, mainly as a result of transient network congestion and TCP dynamics. Concerning the first shortcoming, replacing the Poisson assumption in Cox's construction by a general renewal structure for the start time of sessions is straightforward and introduces greater flexibility, since it requires only that the session interarrival times are i.i.d. With regard to relaxing the assumption of a fixed and constant packet rate for the duration of an entire session, we refer to recent work by Kurtz [Ku], which provides new insights into this difficult problem. Briefly, Kurtz considers a large number of sessions, each of which has a random starting time and duration. Associated with each active session is a cumulative input process, that is, a nondecreasing stochastic process $Y = (Y(t), t \geq 0)$ such that $Y(t)$ represents the cumulative number of packets contributed during the first t time units during an active session; the total length of time τ that a session is active is often modeled separately from

Y. For example, $(Y(t),\ 0 \le t \le \tau)$ with $Y(t) = t$ means constant rate, as in Cox's construction; a $(Y(t),\ 0 \le t \le \tau)$ that has partly slope 1 and partly slope 0, describes *ON/OFF* behavior as in Section 2, but restricted to within a session; and a $(Y(t),\ 0 \le t \le \tau)$ that is piecewise linear with different slopes (including slope 0) is a natural candidate for capturing actual TCP dynamics.

Kurtz's main results show that the same limiting regime holds for an appropriately normalized version of the aggregate packet process under very different assumptions on the fine structure of packet arrivals within sessions (assuming that the starting time of sessions is Poisson): in one case, constant rate is assumed; in the other case, Y is only required to have stationary, ergodic increments. In both cases, the limiting process is fractional Brownian motion, the only Gaussian process with stationary increments that is (exactly) self-similar. Moreover, as in Cox's construction, the essential ingredient for the self-similarity of the limiting process is the heavy-tailed property (3.7) of τ, i.e., the infinite variance assumption on the duration of the individual sessions. For further details and proofs, see [Ku].

3.4 Statistical Analysis of WAN Traffic at the Application Level

Validating the structural modeling approaches to WAN traffic, suggested by Cox's construction (see Theorem 3.2) or Kurtz's, requires checking whether or not measured WAN traffic at the application level is consistent with (i) Poisson (or more general, renewal) session arrival instants and (ii) session lifetimes that satisfy the fundamental heavy-tailed property (3.7). Detailed information concerning the arrival pattern of packets within a session is important for assessing whether a statistically self-similar limiting process is appropriate for accurately describing actual WAN traffic. To this end, we report on empirical studies and provide empirical evidence in favor of the proposed structural models in Section 3.3, especially of Kurtz's construction with its flexible intra-session packet arrival structure. In terms of WAN applications, we focus on measured FTP, HTTP and TELNET traffic traces; while FTP traffic still consumes a major part of available WAN capacity, HTTP traffic continues to increase in volume and has begun to replace FTP as the dominant WAN traffic contributor. TELNET, on the other hand, is a service qualitatively different from these two, with much less demand for bandwidth, but generating a high volume of packets, often one per user keystroke.

First, with regard to the stochastic properties of network session arrivals, measurement studies show that, not surprisingly, the arrival instants of network sessions exhibit a clear diurnal cycle. For example, TELNET sessions peak during afternoon hours and reach a low in early morning hours, almost identical in nature with the calling pattern observed in traditional telephony ([DMRW]). While this observation rules out a homogeneous Poisson model for session arrivals, Paxson and Floyd [PF2] have shown that the arrivals of both TELNET and FTP sessions are well-modeled by nonhomogeneous Poisson processes, with rates that are constant over an hour, but are allowed to change from one hour to the next. The arrivals of HTTP sessions have not been studied in this regard. The difficulty here is determining, from a WAN link trace, when an HTTP session begins. Because of details in how the applications are structured, this determination is straight-forward for FTP and TELNET.

Next we use measured FTP and HTTP traffic to illustrate the empirical evidence in favor of the heavy-tailed property (3.7) of session lifetimes extracted from WAN link traces. We start by postulating that session lifetimes are proportional to "session sizes," where a session size is characterized by the total number of bytes transmitted during the session. (Clearly, matching lifetime and size of a session is an over-simplification, and this aspect will be discussed in more detail below.) For FTP traffic extracted from a mid-1996 30-day WAN traffic trace, Figure 5 shows a log-log complementary distribution plot of the FTP session sizes, for all sessions (left, 56,421 observations) and for only those transferring more than two megabytes (right, 5,886 observations). This distribution closely matches those studied in [Pa], which found the distribution bodies to be well-modeled using the log-normal family of distributions. But the upper tail of the distribution (right plot), containing about 10% of all the data points, provides a very good fit with assumption (3.7), with an estimated α value around 1.1 (the straight line corresponds to a least-squares fit yielding $\alpha = 1.13$). Note that this is a case where the graphical technique (i.e., least-squares fit) by itself is compelling, mainly because of the large number of observations in the right tail; as can be expected from the appearance of the log-log complementary distribution plots in Figure 5, properly performed Hill estimation simply confirms the findings obtained via the graphical method.

The patterns of WAN HTTP *session* lifetimes and sizes have not been characterized, because of the above-mentioned difficulty in extracting this data from a trace of traffic on a WAN link. However, using HTTP *connection* measurements (where a HTTP session is typically made up of many HTTP connections), we obtain strong empirical evidence in favor of the heavy-tailed property (3.7) for HTTP session size and – making the same assumption as above for FTP sessions – for HTTP session lifetimes. Indeed, for a 1996 24-hour measurement period, Figure 6 depicts log-log complementary distribution plots for all HTTP connection sizes (left, 226,386 observations) and for only those transmitting more than 10 kilobytes (right, 32,630 observations). In this case, the upper tail of the distribution contains 14% of the data points, is fully consistent with property (3.7), and yields an α estimate of about 1.35. Figure 7 shows the same information for HTTP connection lifetimes. Here, the upper tail consists of all connections that lasted for more than 10 seconds. The tail contains 28,469 data points (13%), and results in an α estimate of about 1.0. Both figures are clear indications of the heavy-tailed nature of HTTP session lifetimes, but they also illustrate that equating size with duration is not wholly accurate.

To accurately couple session size with session duration or lifetime requires a detailed understanding of the nature of the packet arrival process within a

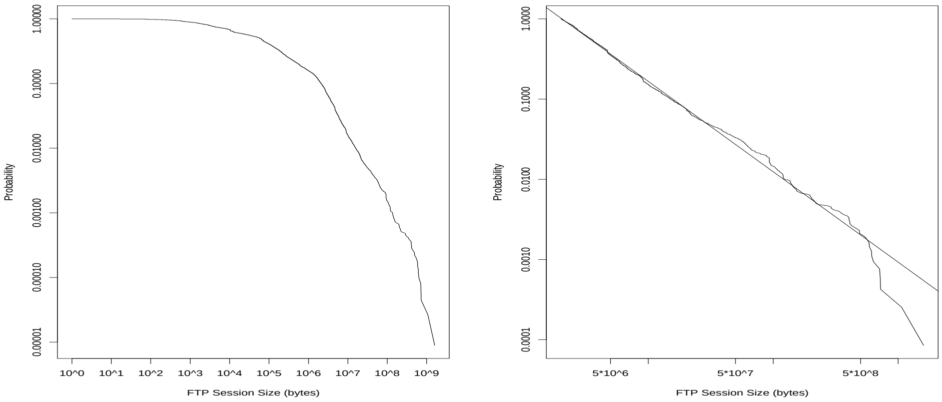

Figure 5: Log-log complementary plot of FTP session sizes: full data set (left), upper tail (right).

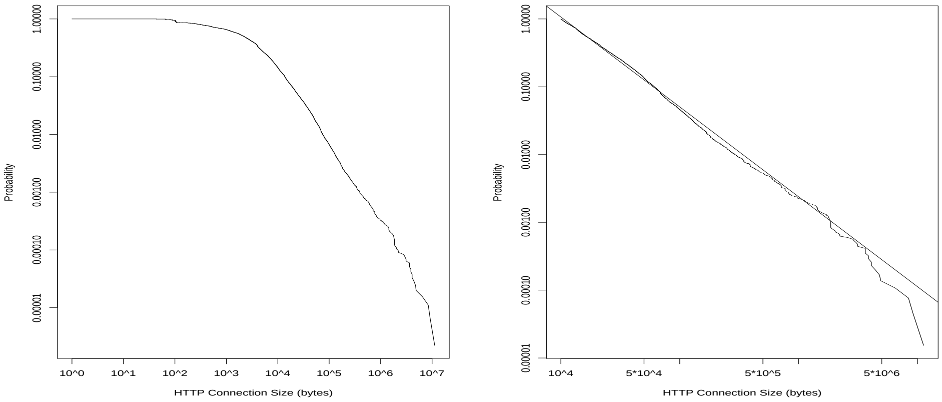

Figure 6: Log-log complementary plot of HTTP connection sizes: full data set (left), upper tail (right).

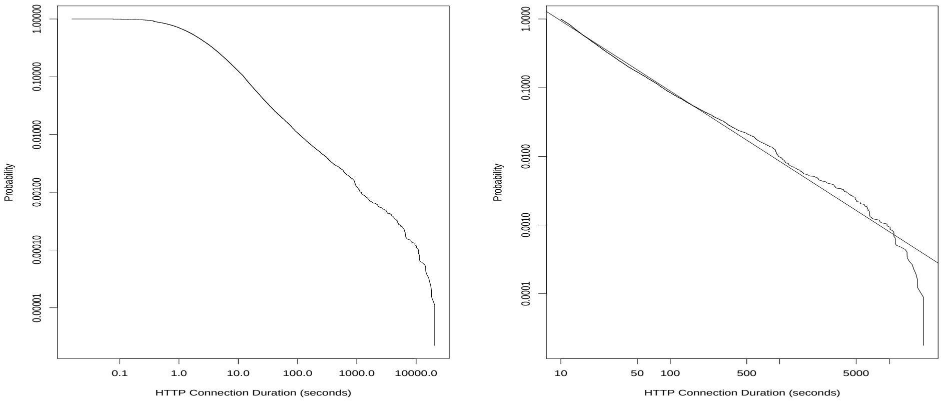

Figure 7: Log-log complementary plot of HTTP connection durations: full data set (left), upper tail (right).

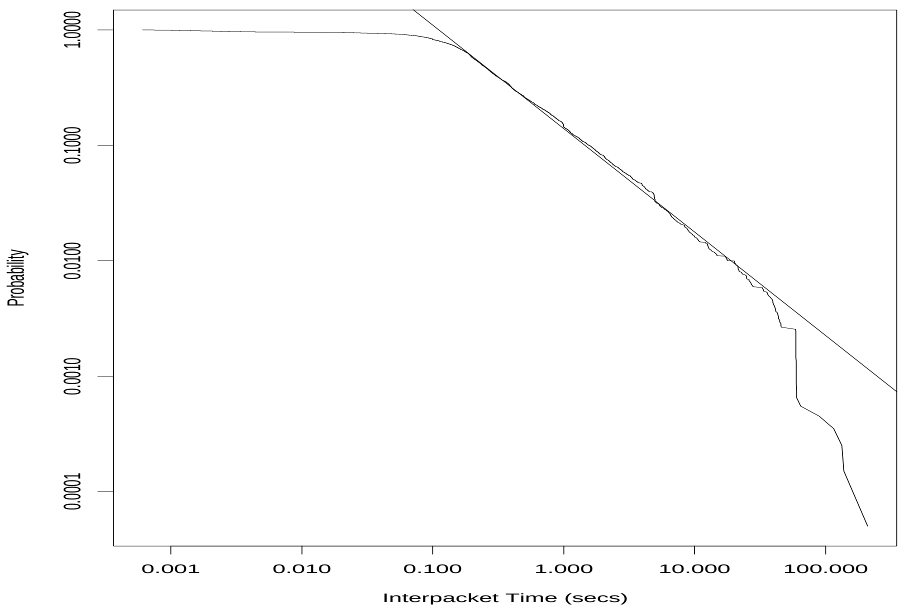

Figure 8: Log-log complementary plot of the full distribution of packet interarrival times for TELNET sessions.

session; as mentioned earlier, such an understanding is also crucial for validating our proposed structural modeling approaches for WAN traffic. To illustrate, consider TELNET traffic, for which extensive packet level measurements exist. For example, a study of WAN traffic at three different sites by Danzig et al. [DJCME] found that the distribution of TELNET packet interarrival times within a single TELNET session was invariant across the three sites. In a subsequent WAN study, Paxson and Floyd [PF2] confirmed this finding, observing the same distribution for data captured at a fourth site three years later [PF2]. What is striking about this distribution is that unlike the exponential models widely assumed for user keystrokes, it is unambiguously Pareto, with an α estimate around, or even slightly below, 1.0. Figure 8 shows a log-log complementary plot of the distribution; the straight line corresponds to a least-squares fit to the upper 75% of the distribution, giving $\alpha = 0.90$. Adopting the conservative position of neglecting the evidence of infinite mean interarrival times (for a discussion about self-similarity and distributions with infinite mean, see [Ma1]), the results of these traffic studies show that for TELNET, the assumption – required for Cox's construction – of a constant packet arrival rate during a session is a large stretch, irrespective of the possible session lifetime characteristics. On the other hand, Kurtz's construction applies directly and suggests, for example, modeling the intra-session packet dynamics via a simple *ON/OFF* process, with the *OFF*-periods corresponding to the packet interarrival times and the *ON*-periods to the transmission times of the individual TELNET packets.

Based on extensive analyses of FTP traffic traces by Paxson and Floyd [PF2] and HTTP traffic measurements by Crovella and Bestavros [CB], similar arguments apply in the context of structural models for FTP and HTTP traffic: the packet arrival instants within FTP and HTTP sessions show highly complex structures and seriously question the assumption of a constant packet rate. On the other hand, Kurtz's construction offers the ability to explicitly model the observed very general intra-session packet arrival behavior with, for example, versions of *ON/OFF* process that allow for different packet arrival rates during the different *ON*-periods, and thus mimic actual TCP dynamics in a real network setting. The two constructions let us choose between capturing only large-scale behavior, but retaining a simpler model (i.e., Cox's), or delving into finer-scale behavior, too, at the cost of a more complex model (Kurtz's). Both structural models are capable of capturing the empirically observed self-similar features in measured WAN traffic traces and explaining it in terms of infinite variance phenomena exhibited by the major applications contributing to aggregate WAN traffic.

4. Conclusions

Previous studies of modern communication network traffic (e.g., LAN and WAN traffic) have conclusively ruled out the possibility of modeling network *packet* arrivals using traditional Poisson processes [JR, Gu, FL, DJCME] – there is simply far too much correlation among packet arrivals for any hope of consistency with the Poisson assumption of independence. One major source of correlation is the prevalence of packets being sent in batches. Another is the correlations between batches introduced by the network itself, e.g., TCP's rate adaptation.

In this paper, we demonstrate how the recently observed self-similarity features of modern network packet traffic [LTWW1, PF1, CB] result in structural traffic models that (i) have a physical meaning in the network context, (ii) highlight the predominance of heavy tails in the packet arrival patterns generated by the individual source-destination pairs or by the major applications that make up the aggregate packet traffic, and (iii) provide fundamental insight into how individual network connections behave. For LAN traffic, the proposed structural models are based on a construction by Willinger, Taqqu, Sherman, and Wilson, which is a modification of one proposed by Mandelbrot [Ma2] in an economic setting. For WAN traffic, the proposed models are based on two related constructions by Cox and Kurtz; all three constructions relate (exactly or asymptotically second-order) self-similarity at the macroscopic or network level directly to heavy-tailed behavior with infinite variance at the microscopic or connection/packet level. While Cox's construction focuses largely on global features at the level of individual network connections and offers little flexibility for modeling the fine structure of packet arrivals within a connection, the constructions of Willinger *et al.* and Kurtz offer the possibility to model packet level dynamics in greater detail. Despite this difference, all three constructions result in self-similar traffic dynamics when the aggregate packet traffic is viewed on large time scales and suggest that for a number of network engineering problems, the fine structure of the packet arrival instants within individual network connections is of minor importance, provided the global connection characteristics exhibit heavy-tails. Clearly, however, this fine structure is bound to play an important role when considering aggregate packet traffic over small time scales and when trying to understand the impact of network protocols on the nature of the traffic through, for example, TCP dynamics. In this context, an important unanswered question concerns the meaning of "large" and "small" time scales from a practical perspective. A closely related issue which requires further study is the adequacy of self-similar limiting processes for describing actual network traffic that necessarily consists of only finitely many connections which, in turn, transmit packets at different rates and which can only be observed over a finite range of time scales. Another important but difficult aspect of traffic modeling that is left for future work concerns the finite link capacities in a network. The models for aggregate network traffic on a particular link, described in this paper, do not include the fact that the link has finite capacity. Yet it is exactly this finite capacity that drives the TCP dynamics and couples the different simultaneous connections sharing the link.

Finally, the question remains: Where do these heavy tails come from? For possible answers, we refer to empirical studies reported, for example, in [PF2, CB, CTB, WTSW2], that relate the observed heavy tails at the connection level directly to heavy-tailed phenomena observed in the sizes of the documents that reside on a typical file or WWW server, and to heavy-tailed characteristics exhibited by human-computer interactions. We must leave the explanation of these latter phenomena to human factors experts and researchers in the cognitive sciences.

References

[Be] J. Beran, *Statistics for Long-Memory Processes* Chapman & Hall, New York, 1994.

[CB] M. E. Crovella and A. Bestavros, Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes, in: *Proc. ACM/SIGMETRICS'96, May 1996.*

[CTB] M. E. Crovella and M. S. Taqqu and A. Bestavros, Heavy-tailed Probability Distributions in the World Wide Web, in: *This volume.*

[Co] D. R. Cox, Long-Range Dependence: A Review, in: *Statistics: An Appraisal*, ed. H. A. David and H. T. David, Iowa State University Press, 1984, 55–74.

[DJCME] 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* **3** (1992), 1–26.

[DMRW] D. E. Duffy, A. A. Mcintosh, M. Rosenstein and W. Willinger, Statistical Analysis of CCSN/SS7 Traffic Data from Working CCS Subnetworks, *IEEE Journal on Selected Areas in Communications* **12** (1994), 544–551.

[FL] H. Fowler and W. E. Leland, Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management, *IEEE Journal on Selected Areas in Communications* **9** (1991), 1139–1149.

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

[HRS] D. Heath, S. I. Resnick and G. Samorodnitsky, Heavy Tails and Long-Range Dependence in On/Off Processes and Associated Fluid Models, *Mathematics of Operations Research* (1997) (to appear)

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

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

[KR] M. Kratz and S. I. Resnick, The qq-Estimator and Heavy Tails, *Stochastic Models* **12** (1996), 699–724.

[Ku] T. G. Kurtz, Limit Theorems for Workload Input Models, in: *Stochastic Networks: Theory and Applications*, ed. F. P. Kelly, S. Zachary and I. Ziedins, Clarendon Press, Oxford, UK, 1996, 339–366.

[LTWW1] 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, 1993*, *Computer Communications Review* **23** (1993), 183–193.

[LTWW2] 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* **2** (1994), 1–15.

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

[Ma1] B. B. Mandelbrot, Self-Similar Error Clusters in Communications Systems and the Concept of Conditional Stationarity, *IEEE Transactions on Communications Technology* **COM-13** (1965), 71–90.

[Ma2] B. B. Mandelbrot, Long-Run Linearity, Locally Gaussian Processes, H-Spectra and Infinite Variances, *International Economic Review* **10** (1969), 82–113.

[MB] R. M. Metcalfe and D. R. Boggs, Ethernet: Distributed Packet Switching for Local Computer Networks, *Communications of the ACM* **19** (1976), 395–404.

[No] I. Norros, A Storage Model with Self-Similar Input, *Queueing Systems And Their Applications* **16** (1994), 387–396.

[Pa] V. Paxson, Empirically-Derived Analytic Models of Wide-Area TCP Connections, *IEEE/ACM Transactions on Networking* **2** (1994), 316–336.

[PF1] V. Paxson and S. Floyd, Wide-Area Traffic: The Failure of Poisson Modeling, in: *Proc. ACM/SIGCOMM'94, London, 1994*, *Computer Communications Review* **24** (1994), 257–268.

[PF2] V. Paxson and S. Floyd, Wide-Area Traffic: The Failure of Poisson Modeling, *IEEE/ACM Transactions on Networking* **3** (1995), 226–244.

[Re] S. I. Resnick, Heavy Tail Modeling and Teletraffic Data, *Annals of Statistics* **25** (1997), 1805–1869. With discussions and rejoinder.

[RS] S. I. Resnick and C. Starica, Consistency of Hill's Estimator for Dependent Data, *Journal of Applied Probability* **32** (1995), 139–167.

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

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

[WTSW] 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*, *Computer Communications Review* **25** (1995), 100–113.

[WTSW2] 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 (Extended Version), *IEEE/ACM Transactions on Networking*, **5** (1997), 71–86.

AT&T Labs-Research, Murray Hill, NJ 07974
Email: walter@research.att.com

Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94720
Email: vern@ee.lbl.gov

Department of Mathematics, Boston University, Boston, MA 02215-2411
Email: murad@math.bu.edu