Looking behind and beyond self-similarity: On scaling phenomena in measured WAN traffic

A. Feldmann, A.C. Gilbert, W. Willinger and T.G. Kurtz

October 1997

To appear in: *Proceedings of the 35th Allerton Conference on Communication, Control and Computing,*
Allerton House, Monticello, IL, September 29 – October 1, 1997

Looking behind and beyond self-similarity: On scaling phenomena in measured WAN traffic*

A. Feldmann, A.C. Gilbert[†], W. Willinger [‡] and T.G. Kurtz

October, 1997

Abstract

In this paper, we report on some preliminary results from an in-depth, wavelet-based analysis of a set of high-quality, packet-level traffic measurements, collected over the last 6-7 years from a number of different working wide-area networks (WANs). We first validate and confirm an earlier finding, originally due to Paxson and Floyd [10], that actual WAN traffic is consistent with statistical self-similarity for sufficiently large time scales. Second, we illustrate that TCP traffic characteristics have undergone major changes within the last 3-4 years, because of the increasing popularity of the Web (WWW) and its emergence as the major contributor to WAN traffic. Finally, we provide empirical evidence that actual WAN traffic traces also exhibit scaling properties over small time scales, but that the small-time scaling phenomenon is distinctly different from the observed large-time scaling property. We relate this newly observed characteristic of WAN traffic to the effects that the dominant network protocols (e.g., TCP) and controls have on the flow of packets across the network, and discuss the potential that *multifractals* have for providing a structural modeling approach for WAN traffic that captures in a compact and parsimonious manner the observed scaling phenomena at large as well as small time scales.

1 Introduction

Self-similar traffic modeling has, from the beginning, emphasized the need for *physical-based* or *structural* approaches for understanding and describing actual network traffic dynamics. Structural models, as discussed in [13], attempt to implicitly take into account the complex hierarchical structure of modern computer communications networks and the intertwined networking mechanisms that determine the nature of the traffic which these networks carry. These structural models are different from the "black box" or "operational" traffic models which were popular in the past.

The main objective of this paper is to improve our current understanding of the dynamic nature of traffic carried over wide-area networks (WANs) such as the Internet and to outline new structural modeling approaches for WAN traffic that are consistent with actual measurements. In particular, we (i) identify distinct scaling regions in aggregate packet-level WAN traces, one for small time scales (a few hundreds of milliseconds and

*A. Feldmann, A. Gilbert, and W. Willinger are with AT&T Labs-Research, Florham Park, NJ 07932-0971, email: {anja, agilbert, walter}@research.att.com. T. Kurtz is with the Center for the Mathematical Sciences, University of Wisconsin at Madison, Madison, WI 53717, email: kurtz@math.wisc.edu.

†This Research was supported by the NSF Grant DMS-9705665 at Yale University and AT&T Labs-Research, Florham Park, NJ.

‡This research was partially supported by the NSF Grant NCR-9628067 at the University of California at Santa Cruz.

below) and one for large time scales (beyond a few hundreds of milliseconds); (ii) demonstrate the changing nature of WAN traffic characteristics at the transport (i.e., TCP/IP) level over time; (iii) provide empirical evidence for the presence of application-layer WAN traffic characteristics, which allows us to adopt a structural modeling approach that explains the empirically observed large-time scaling phenomenon in aggregate WAN traffic; and (iv) discuss the feasibility of a structural modeling approach for WAN traffic that is capable of capturing the large-time as well as small-time scaling properties in a compact and parsimonious manner. To accomplish these goals, we rely on a unique set of high-quality, packet-level WAN traffic measurements, collected over the past six to seven years from a number of different working WANs. Our in-depth analysis of these WAN traffic traces is based on a wavelet-based technique recently suggested and developed by Abry and Veitch [1]. This technique provides a natural and effective tool for investigating scaling properties that may be present in a large data set.

There have been significant advances in the recent past in structural modeling approaches for local area network (LAN) traffic as opposed to WAN traffic. In fact, the empirically observed self-similar or fractal nature of LAN traffic at the aggregate or macroscopic level (i.e., generated by all active hosts on the network; see for example [9]) has recently resulted in structural models that are based on self-similar processes, have a physical meaning in the LAN context, and provide fundamental and new insights into how individual network connections behave. However, WANs differ from LANs in a number of fundamental ways, and these differences make structural modeling of WAN traffic more challenging. For example, because of the problem of latency, WANs often rely on sophisticated transport protocols (e.g., Internet uses predominantly TCP) which, in turn, are bound to introduce additional structure to the flow of packets; such a structure is less obvious in LAN traffic because feedbacks in LANs are essentially instantaneous. Despite these differences, empirical evidence for (asymptotic) self-similarity, i.e., large-time scaling phenomena, also exists for WAN traffic and was first reported by Paxson and Floyd [10]; see also [14]. In Section 2 below, we use a collection of more recent WAN traces and a wavelet-based scaling analysis to validate this previously observed large-time scaling property of WAN traffic. In addition, we identify a new scaling phenomenon for small time scales (below a few hundreds of milliseconds) that has gone unnoticed in previous studies and is distinctly different from that for large time scales. In Section 3, we provide new evidence that structural modeling approaches for WAN traffic (e.g., Cox's construction; see [3]) remain partly valid for today's WAN traffic, even though over the last three to four years, WWW has become the main WAN application and typically makes up a major portion of modern WAN traffic. At the same time, we also demonstrate in Section 3 that mainly due to the Web, the traffic characteristics at the transport layer have undergone significant changes between the pre-WWW days and now, and that these changes clearly reveal the limitations of structural modeling a la Cox. To this end, we confirm here with our latest collection of WAN traffic traces a recent finding by Feldmann [5] who observed self-similar features not only at the packet/bytes level but also at the TCP level. In light of this empirical evidence, we revisit the question of structural modeling approaches for modern WAN traffic and discuss two alternatives. One approach which builds on Cox's approach is based on a construction proposed by Kurtz [8] and allows for very flexible "within-session" traffic patterns. A second, more radical approach to understanding and describing the actual dynamics of modern WAN traffic is based on *multifractals* (for details, see for example, [4, 7] and references therein). We conclude in Section 4 with an intuitive explanation for why multifractals might be an appropriate mathematical technique for gaining a better understanding of the dynamic nature of modern high-speed network traffic.

2 Scaling Properties of Measured WAN Traffic

2.1 MRA and a Wavelet-Based Scaling Analysis

Wavelets provide a tool for time-frequency localization; the wavelet transform divides data into different frequency components and analyzes each component with a resolution matched to its scale. We can use the coefficients of a wavelet decomposition to directly study the scale (or frequency) dependent properties of the data. In particular, applying a recently proposed technique by Abry and Veitch [1], we can use wavelets to investigate the scaling structure in the spectra of self-similar processes which provide statistical models for describing certain aspects of modern high-speed network traffic. Multiresolution analysis (MRA) gives us a natural framework for understanding wavelet bases and their transforms.

The key feature of an MRA is that we can write an approximation of a signal X_t at scale j (with resolution 2^j) as the sum of a coarser approximation at scale $j+1$ (with resolution 2^{j+1}) and the difference between these two approximations. We may iterate this procedure, writing the approximation at scale $j+1$ as a sum of a coarser approximation and the difference. This procedure can be implemented using a pyramidal filter-bank algorithm which has computational cost $O(N)$ for data of length N. We write the *wavelet decomposition* of a signal X as

$$X = \sum_{k\in\mathbf{Z}} \langle X, \phi_{0,k}\rangle\, \phi_{0,k} + \sum_{j\le 0}\sum_{k\in\mathbf{Z}} \langle X, \psi_{j,k}\rangle\, \psi_{j,k} = \sum_{j\in\mathbf{Z}}\sum_{k\in\mathbf{Z}} \langle X, \psi_{j,k}\rangle\, \psi_{j,k}.$$

We call the inner products $\langle X, \psi_{j,k}\rangle$ of X with the rescaled and translated copies of the wavelet ψ the *wavelet coefficients* $d_{j,k}$ of X. The set of all wavelet coefficients is generally referred to as the *discrete wavelet transform (DWT)* of the signal X. The coefficient $|d_{j,k}|^2$ measures the amount of energy in a signal X about the time $t_0 = 2^j k$ and about the frequency $2^{-j}\lambda_0$, where λ_0 is a reference frequency which depends on the wavelet ψ. Abry and Veitch show that the time average of $|d_{j,k}|^2$ at each scale j is a useful spectral estimator. In fact, if E_j denotes the average of $|d_{j,k}|^2$ at each scale, $E_j = N_j^{-1}\sum_k |d_{j,k}|^2$ (N_j is the number of available wavelet coefficients at scale j), then E_j is a measure of the energy that lies within a given bandwidth 2^{-j} around frequency $2^{-j}\lambda_0$.

Consider now a signal X_t generated by a finite variance, wide-sense stationary, *long-range dependent* process with *Hurst parameter* $H \in (1/2, 1)$. In the frequency domain, long-range dependence is characterized by a spectral density $f(\lambda)$ that exhibits a power-law near the origin; i.e., $f(\lambda)$ has the form

$$f(\lambda) \sim c_f |\lambda|^{1-2H} \quad \text{as } \lambda \to 0, \tag{1}$$

where $c_f = 1/\pi\sigma^2 c_\rho \Gamma(2H-1)\sin(\pi(1-H))$ and $\sigma^2 = \mathrm{Var}(X_t)$. Long-range dependence plays an important role in the study of self-similar processes, where a wide-sense stationary process X is called *exactly self-similar* if for all integers $m > 0$,

$$X = m^{1-H} X^{(m)}. \tag{2}$$

Here, the equality is understood in the sense of finite-dimensional distributions, and the aggregated processes $X^{(m)}$ are defined by $X^{(m)}(k) = m^{-1}\sum_{i=(k-1)m+1}^{i=km} X_i,\ k \ge 1$. *Asymptotic self-similarity* is defined similarly but we only require that the above equality holds in the limit as $m \to \infty$. With this definition, X is long-range dependent if and only if X is (asymptotically) self-similar, and methods for studying the dependence structure of X can thus be exploited for investigating the scaling phenomenon expressed in Equation

(2). To illustrate, we return to the wavelet-based technique proposed by Abry and Veitch [1] and assume that the signal X is such that Equation (1) holds for all frequencies. Then the expectation of E_j is given by

$$\mathbf{E}[E_j] = \int f(\lambda) 2^j |\hat{\psi}(2^j \lambda)|^2 \, d\lambda = c_f |2^{-j}\lambda_0|^{1-2H} \int |\lambda|^{1-2H} |\hat{\psi}(\lambda)|^2 \, d\lambda,$$

where $\hat{\psi}(\lambda)$ is the Fourier transform of $\psi(t)$. Observe that the multiplicative bias of E_j is simply a multiplicative constant which is independent of scale j and results in an effective scaling analysis of a given signal X by plotting $\log_2 E_j$ against scale j and identifying scaling regions, breakpoints and non-scaling behavior. It also yields an unbiased estimator for the Hurst parameter H (for details, see [1]). To insure that the integral $\int |\lambda|^{1-2H} |\hat{\psi}(\lambda)|^2 \, d\lambda$ exists and is finite we must choose a wavelet ψ which has $M > H - 1$ vanishing moments. In other words, the wavelet ψ must be orthogonal to all polynomials of degree less than M; i.e., $\int t^l \psi(t) \, dt = 0 \quad \text{for } l = 0, \ldots, M-1$. The number of vanishing moments controls the order of the zero in the Fourier transform of ψ about $\lambda = 0$. If M is large enough, the behavior of $|\hat{\psi}(\lambda)|^2$ will balance the singularity of the long-range dependent spectrum.

2.2 Measured WAN Traces Now and Then

In [1], Abry and Veitch use their wavelet-based technique to validate previously obtained results showing that measured Ethernet LAN traffic is consistent with long-range dependence and to estimate the corresponding values of the Hurst parameter. In this paper, we use their methodology primarily for investigating the scaling behavior of measured WAN traffic traces; that is, for identifying regions where the scaling property (2) holds for a range of time scales, for detecting changes in scaling behavior, and for finding ranges of time scales with more complex scaling patterns than those captured by Equation (2).

Packet Traces	Year	Time Stamp Accuracy	% WWW (packets)
Bellcore	Jan.1990	1msec	0%
LBL	Jan.1994	1msec	3%
Bellcore	Dec.1994	10μsec	10%
MH (AT&T-Labs)	Feb.1997	1msec	27%
FP (AT&T-Labs)	Aug.1997	1msec	32%
TCP Connection Traces			% WWW (connections)
lbl1	Fall 1993	1msec	0%
lbl2	Fall 1993	1msec	0%
cmu	Jun.1995	1msec	22%
mh (AT&T-Labs)	Feb.1997	1msec	80%
fp (AT&T-Labs)	Aug.1997	1msec	87%

Table 1: Packet-level and TCP connection-level WAN traces used in our study.

The five hour-long packet-level WAN traces used in our study are summarized in the top part of Table 1. The Bellcore'90 trace pre-dates WWW, and TELNET and FTP (as well as NNPT) comprised the main applications at that time. Also note that in 1990, the NSFNET backbone was T1-based; i.e., running at 1.5 Mbps. Both the LBL'94 and the Bellcore'94 WAN traces contain some Web traffic (close to 10% for the Bellcore trace) and represent measurements of Internet traffic as the Internet is going through drastic changes (e.g., emergence of WWW, transition to a 45 Mbps backbone). WWW is the main WAN application and makes up a major portion of the traffic in both the Murray

Hill (MH) AT&T-Labs'97 and Florham Park (FP) AT&T-Labs'97 traces and by that time, the privatization of the Internet is well on its way.

Selecting three representative traces which span seven years of WAN traffic, with 0%, 10%, and 30% or more Web traffic, respectively, Figure 1 (left column) shows the scaling analyses for the time series representing the number of packets per ten milliseconds for the Bellcore'90, Bellcore'94, and FP'97 WAN traces; the corresponding analysis for the number of bytes per 10 milliseconds time series is given in the right column in Figure 1. We have plotted the energy E_j as a function of the resolution level 2^j, on a $\log_2 - \log_2$-scale; i.e., $\log_2 E_j$ as a function of scale j. To clarify the time scales which correspond to scales j, we have included the time (in seconds) on the top axis of each graph. The scaling analysis of a traffic trace which is asymptotically self-similar (or, equivalently, exhibits long-range dependence) will, for large times/scales, result in a linear relationship between $\log_2 E_j$ and the scale j. If the trace is exactly self-similar (i.e., the spectral density characterization of long-range dependence in Equation (1) holds for all frequencies), a plot of $\log_2 E_j$ vs. j will show a linear relationship for all scales. As can be seen, all plots show an approximate linear relationship between $\log_2 E_j$ and scale j for sufficiently large time scales, i.e., the corresponding WAN traces are fully consistent with asymptotic self-similarity or long-range dependence, irrespective of where the traces were collected and whether or not they contained no, some or mostly Web traffic. However, while pre-Web traffic (i.e., Figure 1, top row) seems, in fact, to be consistent with exact self-similarity (note that the linear relationship holds throughout the whole range of observed time scales, from tens of milliseconds to seconds and minutes and beyond), the more recently collected WAN traces exhibit in addition a non-trivial scaling behavior for small time scales (typically below a few hundreds of milliseconds) that is distinctly different from the large-time scaling behavior. Moreover, the "change-point" shows up very clearly as a pronounced "knee" in the graphs in the middle and bottom rows in Figure 1, approximately at scale $j = 5$ corresponding to about 500 milliseconds, irrespective of whether we consider the packet counts or byte counts.

While our finding of the large-time scaling is not new and simply confirms earlier studies by Paxson and Floyd [10] (see also Willinger et al. [14]) who first observed the asymptotically self-similar or long-range dependent nature of measured WAN traffic, our empirical observation of the existence of a pronounced small-time scaling behavior that's clearly different from the large-time scaling phenomenon is new and appears to have gone unnoticed in earlier empirical studies of measured WAN traffic that relied exclusively on techniques such as the variance-time analysis and/or the R/S method. Subsequently, the results of our wavelet-based scaling analysis of a historical collection of WAN traces imply that to fully understand and accurately describe modern WAN traffic, exactly self-similar processes are not sufficient and need to be augmented or replaced by traffic models that allow for a richer scaling behavior.

3 Physical Explanations for Scaling in WAN Traffic

3.1 Large-Time Scaling Phenomenon: Application Layer

In their WAN traffic study, Paxson and Floyd [10] also looked into plausible physical explanations for the observed asymptotically self-similar nature (i.e., large-time scaling phenomenon) of measured WAN traffic. Focusing on the applications, they report on an extensive analysis of FTP and TELNET traffic traces which shows Poisson session arrivals and heavy-tailed session durations or sizes with infinite variance. In view of these

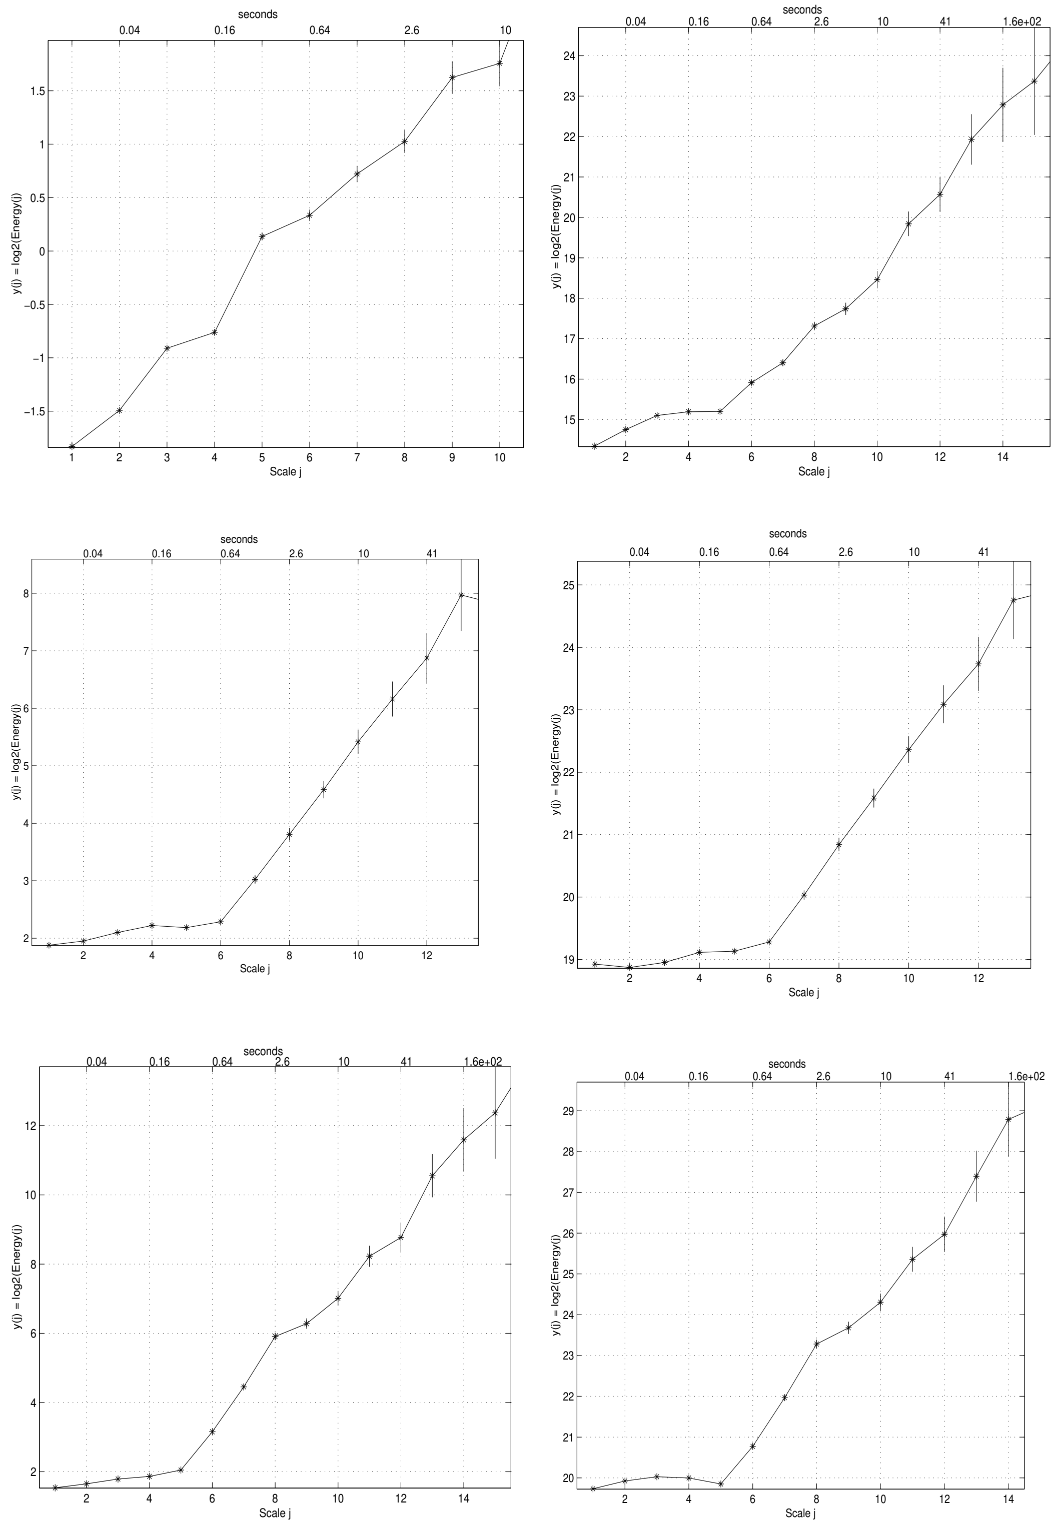

Figure 1: Scaling analysis of packet-level WAN traces (left column: packets/10 msec; right column: bytes/10 msec): Bellcore'90 (top), Bellcore'94 (middle), FP'97 (bottom).

empirically observed application-layer traffic characteristics, a construction due to Cox [3], also known as an *immigration death process* or *$M/G/\infty$ queueing model*, provides a structural modeling approach for pre-Web WAN traffic that is (i) mathematically rigorous, (ii) consistent with measured packet-level WAN traffic at the macroscopic (i.e., aggregated over all individual connections) as well as microscopic (i.e., for individual connections) level, and (iii) highlights the intimate connection between high variability (i.e., heavy-tailed phenomenon) at the microscopic level and (asymptotic) self-similarity (i.e., large-time scaling phenomenon) at the macroscopic level. For details, see for example [13]. In a longer version of this paper [6], we use a plausible, measurable substitute for Web session information, and show that Cox's construction remains valid as a structural modeling approach for WAN traffic, even when the latter is dominated by the Web – and as long as we are only concerned with application and packet level information and ignore other layers in the networking hierarchy, for example TCP layer or IP layer (see Section 3.2 below).

3.2 Large-Time Scaling Phenomenon: TCP Layer

To date, structural modeling approaches for WAN traffic (e.g., Cox's construction) have focused almost exclusively on explaining the observed large-time scaling phenomenon in packet-level WAN traffic in terms of the phenomenon of high variability or infinite variance at the application level. These approaches reflect a source-centered view of traffic modeling, where sources are assumed to simply inject (at a constant rate) as many packets into the network as is required by a user session, and they essentially ignore the impact that the applications (e.g., downloading several Web pages) and networks (e.g., TCP dynamics) have on the flow of packets that individual sources attempt to transmit over the network. In this subsection, we focus exclusively on the transport layer and study the characteristics of measured WAN traffic at the TCP connection level; that is, we investigate the statistical properties of the process representing the number of TCP connection arrivals per time unit. We illustrate our analysis of WAN traffic at the TCP connection level with four hour-long traces whose pertinent statistics are summarized in the lower part of Table 1. Note that lbl2'93 contains an hour worth of all TCP connections from LBL 7 in [10], while lbl1'93 consists of only those TCP connections in lbl1'93 that correspond to successful TELNET or FTP control sessions. As in Section 2.2, we can again study the changing nature of WAN traffic, this time at the level of TCP connection arrivals, over a period where WANs have undergone drastic changes.

Performing the same wavelet-based scaling analysis as in Section 3.1, Figure 2 shows the scaling properties for the time series representing the number of TCP connection arrivals per second, for all five connection-level WAN traces. Starting with the top left plot in Figure 2, we observe a trivial (i.e., a value of the Hurst parameter of about 0.5) scaling behavior for the lbl1'93 data set across all scales. However, when moving from the lbl1'93 to the lbl2'93 data set, i.e., when considering all TCP-connections arrivals (including those spawned by FTP control connections and SMTP or email), we see in Figure 2 (top right plot) a drastic change, namely the appearance of a non-trivial scaling region for the large time scales. The difference between the two plots in the top row of Figure 2 can be attributed to the facts that SMTP-triggered TCP connections dominate and that FTP-control connections typically spawn a number of TCP connections within the FTP session. Moving on to the cmu'1995 trace and analyzing only the HTTP-related portion of the TCP connections, Figure 2 (middle row) shows yet another change in the scaling behavior; Web traffic at the level of TCP connection arrivals no longer exhibits trivial scaling at the small scales but shows across all scales the same non-trivial scaling

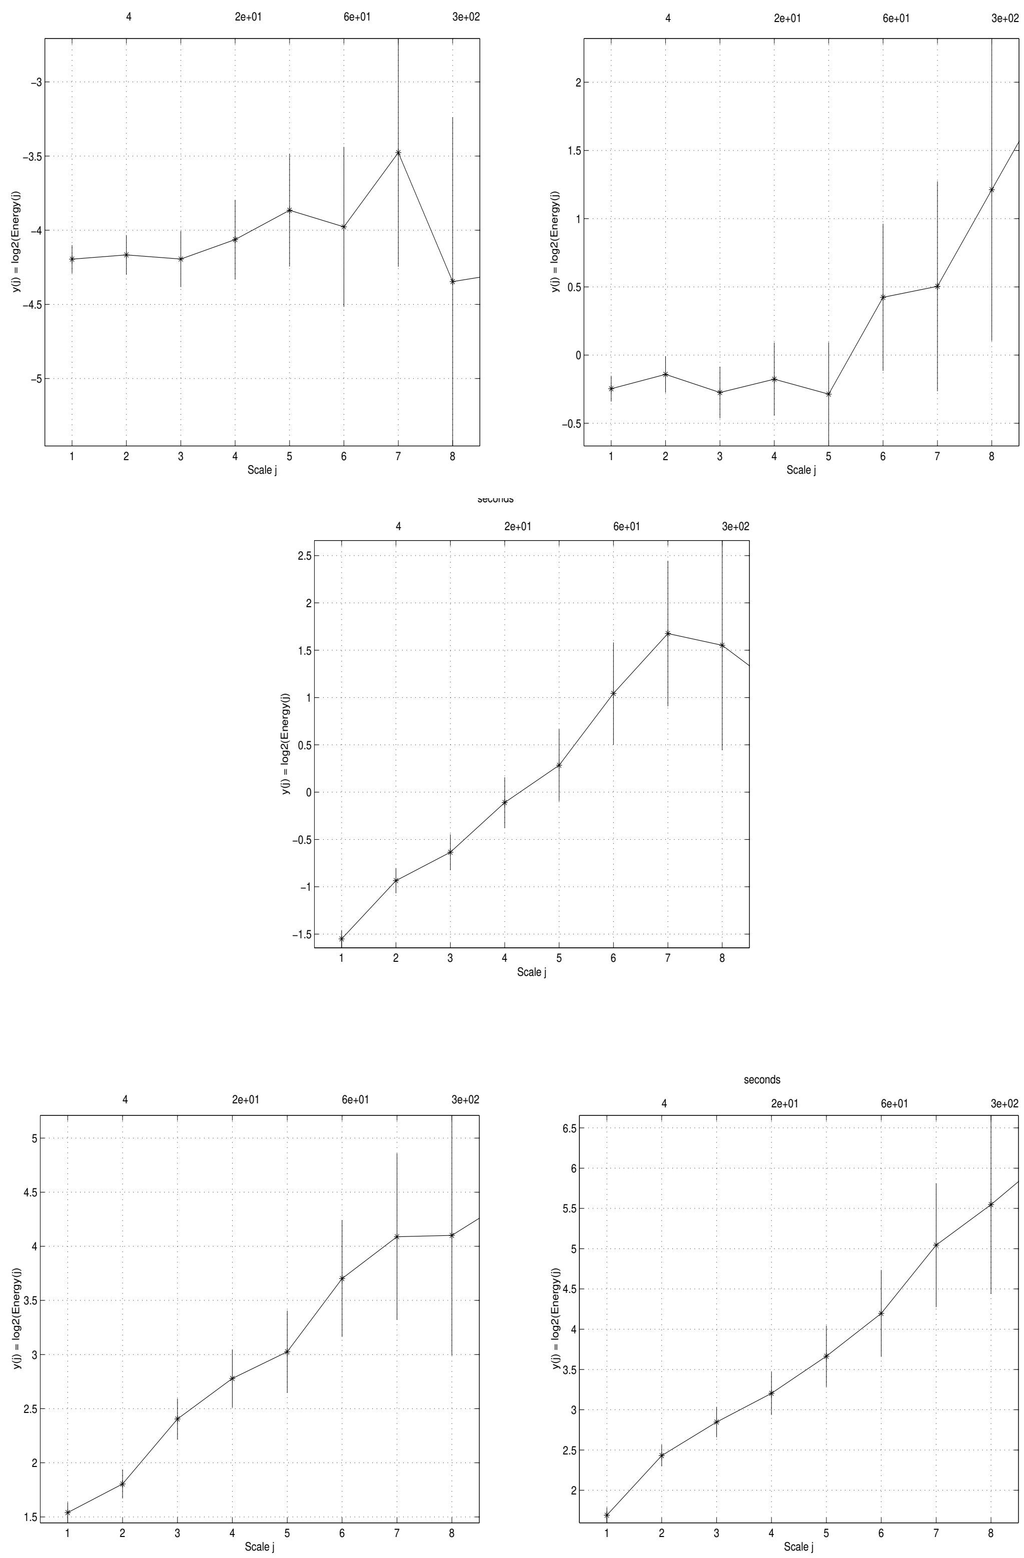

Figure 2: Scaling analysis of TCP connection-level WAN traces (number of TCP connection arrivals per second): lbl1'93 (top left); lbl2'93 (top right); cmu'95 (middle row); mh'97 (bottom left), fp'97 (bottom right).

that we have seen in the pre-Web lbl2'93 trace for large time scales; that is, the traffic representing the number of HTTP-related TCP connection arrivals per second is (exactly) self-similar for time scales of one second and beyond. As can be seen from the remaining two plots in Figure 2 (i.e., bottom row), this observation is confirmed and becomes even more pronounced as we move to the latest connection-level WAN traces and perform the same scaling analysis for the HTTP-related portion of TCP connections. These results confirm earlier findings by Feldmann [5] and shows in addition the gradual appearance over the last few years of non-trivial scaling in WAN traffic at the transport layer, from trivial scaling to non-trivial large-time scaling to non-trivial scaling for all scales. It will be interesting to see, how this empirically observed scaling phenomenon at the transport layer will be affected as a growing portion of Internet traffic will be "multicast", where a single sender transmits to multiple receivers, as other new transport protocols come along, or as the current characteristics of a "typical" Web-page change.

As far as structural modeling approaches for WAN traffic is concerned, the observed self-similar nature of WAN traffic at the TCP connection level clearly demonstrates the limitations of constructions a la Cox that allow for no other "within-session" structure than sending packets at constant rate. In particular, Cox-like constructions are unable to account for the observed TCP connection level characteristics and are, therefore, inadequate to capture small-time scaling phenomena as observed in Section 2.2. However, recent work by Kurtz [8] provides new insights into the problem of allowing for a more flexible "within-session" structure. His construction explicitly allows for the fragmentation of an application-layer session into one or more transport-layer connections (e.g., an Web-session being segmented into a number of TCP connection, interspersed with "idle periods" representing periods of user inactivity or "think time"; see [2]). From the perspective of structural modeling approaches for WAN traffic, Kurtz's construction is an improvement over Cox's construction because it can accommodate a number of relatively complex "within-session" structures. However, at this time, it remains unclear how one would account for the observed self-similar nature of the TCP connection arrival counts or how such a structure (if it exists) would impact the limiting result.

3.3 Small-Time Scaling Properties: Preliminary Observations

We have seen in Section 2.2 that modern-day packet-level WAN traffic exhibits small-time scaling features that are distinctly different from the observed large-time scaling phenomenon. Phenomenological or structural modeling approaches for WAN traffic, such as Cox's or Kurtz's constructions, successfully explain the latter phenomenon in terms of the observed high variability of WAN traffic at the application level, but have failed so far to capture the former. As an illustration of the insights gained from measuring and analyzing high-time resolution WAN packet traces and the ensuing difficulties in trying to understand and describe in a compact manner the small-time scaling behavior in packet-level WAN traffic, we show in Figure 3 the results of a wavelet-based scaling analysis for a WAN trace collected in July'97 from a working commercial network. For this WAN traffic trace, the guaranteed time stamp accuracy is a few micro-seconds, and allows for a scaling analysis down to the millisecond scale. Somewhat surprisingly, the plot shows yet further scaling structure on the very fine time scales, distinctly different from the observed scaling properties over intermediate time scales (see Figure 1), and still different from the well understood large-time scaling feature. This observation should serve as a reminder that even though "black box" models such as fractional ARIMA(p, d, q) processes can often be tweaked so as to reproduce a certain prescribed scaling behavior, they do so without providing any insights into the origins of the fitted scaling behavior. Moreover,

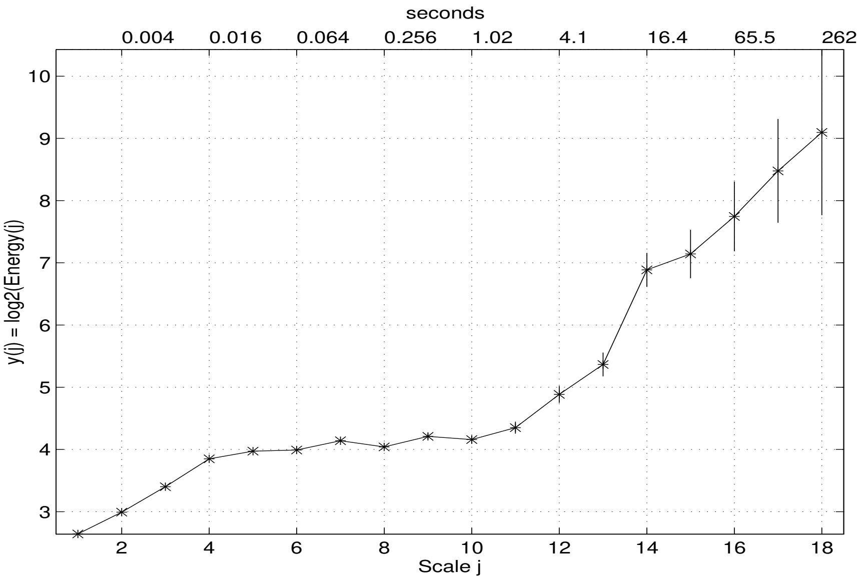

Figure 3: Scaling analysis for a high-time resolution WAN trace (time stamp accuracy: of a few microseconds).

if a more detailed scaling analysis reveals yet further scaling features (as in the case in Figure 3), these "operational" modeling approaches simply result in replacing one black box by another.

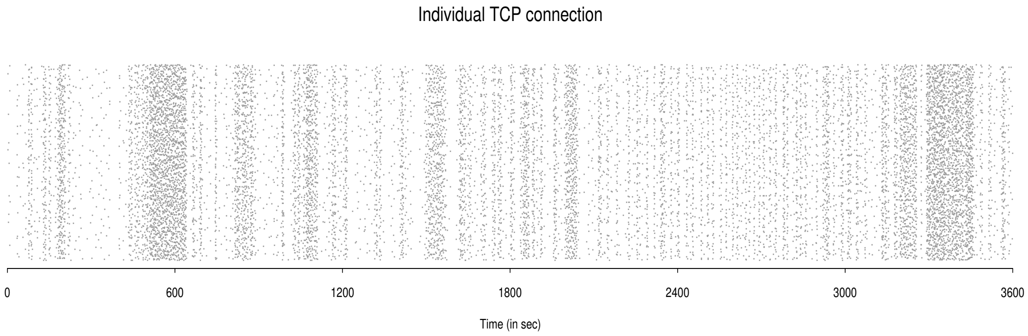

Figure 4: Textured plot for an individual TCP connection. The connection lasted for an entire hour, and consists of 14,305 packets.

From a networking perspective, the presence of non-trivial scaling behavior at small time scales comes as no surprise and is intimately related to the intricacies of the different protocols that rule the flow of traffic at each layer in the networking hierarchy. These protocols interact in non-trivial ways with one another, and are responsible for creating complex interactions between the network on the one side and the sources on the other side. Figure 3.3 illustrates the richness in structure that exists within a single TCP connection via a textured plot and brings out very clearly the high variability of "within-TCP-connection" activity levels and behaviors. The challenge for structural modeling approaches for WAN traffic is to (i) understand the underlying mechanisms that give rise to such rich structures at the transport layer and possibly at the other layers in the networking hierarchy as well, (ii) describe the relevant aspects of the observed structure in a compact and parsimonious manner, and (iii) provide mathematical constructions that reflect the underlying mechanisms and, at he same time, continue to capture the well-understood large-time scaling behavior.

4 Conclusions and Outlook

Currently considered structural modeling approaches for WAN traffic, especially two constructions due to Cox and Kurtz, respectively, are well-suited to explain the large-time scaling features in terms of the well-understood application layer traffic characteristic. However, in their current form, neither of these approaches are capable of explaining the small-time scaling properties, nor are they able to reproduce the self-similar nature of traffic at the TCP layer in an intrinsic and natural manner. The most promising of the current approaches appears to be Kurtz's construction, mainly because in its general form, it allows for the possibility of very general "within-session" packet arrival patterns, and because to date, the limiting regime (i.e., aggregation over many sessions) is only known for a very special cases, all of which can only account for large-time scaling properties. The challenge is to incorporate a more "network-centered" view of traffic modeling into Kurtz's construction (as opposed to further pursuing the traditional "source-centered" approach; see Section 3.2), so as to better account for the impact that the network has on the sources.

One plausible way of making Kurtz's construction more network-oriented – without getting into the details of how the different protocol at the different layers in the networking hierarchy work and interact with one another – is to view WANs or other networks, together with their protocols and controls, as defining the deterministic mechanisms and rules of a process that fragments units of information at one layer into smaller units at the next layer etc. That is, networks act as *cascades* by fragmenting, for example, a Web session into a number of TCP connections (determined by the HTTP protocol) which, in turn, are further fragmented into several HTTP requests (the mechanism for this is provided by the HTTP protocol) which, in turn, are fragmented into IP flows, and so on. Note that during this fragmentation process, the total number of information units (in our example, this would be the total number of bytes transmitted during the Web-session) remains essentially preserved and hence, the cascade analogy is appropriate. Cascades are an important paradigm in the theory of *multifractals* as they can account for many multifractal or self-similar measures, that is, measures with a nontrivial multifractal structure (for an overview, definitions and properties of multifractals, see for example Gupta and Waymire [7] and Evertsz and Mandelbrot [4]). Although the multifractal approach to understanding and describing network traffic appears to be a radical departure from the more conventional structural traffic modeling approaches pursued to date, there exists a close link between Kurtz's construction and random cascade models: simply allow the "within-session" traffic patterns (for details, consult [8]) to be multifractal measures, generated via appropriately defined random cascade models. In addition, there already exists empirical evidence in favor of non-trivial multifractal properties of measured WAN TCP traffic – see the recent work by Riedi and Levy-Vehel [11] and the discussion in [12]. Although multifractals and random cascade models are new traffic modeling paradigms, the motivation for studying them in the network context comes directly from the desire for physically meaningful descriptions of empirically observed phenomena in measured network traffic. In view of the already existing empirical evidence, these new modeling paradigms are more then speculation and offer unique opportunities for studying challenging new mathematical and statistical problems. At the same time they point towards new ways of understanding and describing real-world network traffic, resulting in new engineering insights that can be expected to be of practical relevance for a wide range of network engineering tasks.

Acknowledgments

Many of the traces considered in this paper were collected using the *tcpdump* packet capture tool developed by V. Jacobbson, C. Leres and S. McCanne and available via anonymous ftp to `ftp.ee.lbl.gov`. To extract TCP connection information from the traces, we relied on V. Jacobson's *tcp-conn* and/or V. Paxson's *tcp-reduce* tools, both of which are available from `http://ita.ee.lbl.gov/index.html`.

The Bellcore traces were collected by D. V. Wilson, the LBL/lbl traces were gathered by V. Paxson and are available from `http://ita.ee.lbl.gov/index.html`, and we thank S. Alexander and S. Gao from AT&T Labs-Research for making the MH and FP traffic collection possible. We also acknowledge the help of many of our colleagues at AT&T Labs, especially of J. Friedmanns, with the data collection effort.

References

[1] P. Abry and D. Veitch. Wavelet analysis of long-range dependent traffic. To appear in IEEE Transactions on Information Theory.

[2] M. E. Crovella and A. Bestavros. Self-similarity in world wide web traffic - evidence and possible causes. *Proceedings of Sigmetrics'96*, pp. 160–169, 1996.

[3] 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, Ames, Iowa, 1984.

[4] C. J. G. Evertsz and B. B. Mandelbrot. Multifractal measures. In H.-O. Peitgen, H. Jurgens and D. Saupe, editors, *Chaos and Fractals: New Frontiers in Science*, Springer-Verlag, New York, 1992.

[5] A. Feldmann. Modelling characteristics of TCP connections. Preprint, 1996.

[6] A. Feldmann, A. C. Gilbert, W. Willinger and T. G. Kurtz. The changing nature of network traffic: Scaling phenomena. Preprint, 1997.

[7] V. K. Gupta and E. C. Waymire. A statistical analysis of mesoscale rainfall as a random cascade. *Journal of Applied Meteorology* **32**, pp. 251–267, 1993.

[8] T. G. Kurtz. Limit theorems for workload input models. In F. P. Kelly, S. Zachary, and I. Ziedins, editors, *Stochastic Networks: Theory and Applications.* Clarendon Press, Oxford, 1996.

[9] 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**, pp. 1–15, 1994.

[10] V. Paxson and S. Floyd. Wide area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking* **3**, pp. 226–244, 1995.

[11] R. H. Riedi and J. Levy Vehel. Multifractal properties of TCP traffic: A numerical study. Preprint, 1997.

[12] M. S. Taqqu, V. Teverovsky and W. Willinger. Is network traffic self-similar or multifractal? *Fractals* **5**, pp. 63-73, 1997.

[13] W. Willinger, V. Paxson, and M. S. Taqqu. Self-similarity and heavy tails: Structural modeling of network Traffic. To appear in R. Adler, R. Feldman, and M. S. Taqqu, editors, *A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy Tailed Distributions*, Birkhauser Verlag, Boston.

[14] 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. *IEEE/ACM Transactions on Networking* **5**, pp. 71–86, 1997.