Data networks as cascades:
Investigating the multifractal nature of Internet WAN traffic

A. Feldmann, A. C. Gilbert[*] and W. Willinger[†]

AT&T Labs – Research, Florham Park, NJ

{anja,agilbert,walter}@research.att.com

Abstract

In apparent contrast to the well-documented self-similar (i.e., monofractal) scaling behavior of measured LAN traffic, recent studies have suggested that measured TCP/IP and ATM WAN traffic exhibits more complex scaling behavior, consistent with multifractals. To bring multifractals into the realm of networking, this paper provides a simple construction based on cascades (also known as multiplicative processes) that is motivated by the protocol hierarchy of IP data networks. The cascade framework allows for a plausible physical explanation of the observed multifractal scaling behavior of data traffic and suggests that the underlying multiplicative structure is a traffic invariant for WAN traffic that co-exists with self-similarity. In particular, cascades allow us to refine the previously observed self-similar nature of data traffic to account for local irregularities in WAN traffic that are typically associated with networking mechanisms operating on small time scales, such as TCP flow control.

To validate our approach, we show that recent measurements of Internet WAN traffic from both an ISP and a corporate environment are consistent with the proposed cascade paradigm and hence with multifractality. We rely on wavelet-based time-scale analysis techniques to visualize and to infer the scaling behavior of the traces, both globally and locally. We also discuss and illustrate with some examples how this cascade-based approach to describing data network traffic suggests novel ways for dealing with networking problems and helps in building intuition and physical understanding about the possible implications of multifractality on issues related to network performance analysis.

1 Introduction

The empirically observed self-similar or fractal nature of aggregate network traffic [17, 25] is caused by the high-variability of the individual connections that make up the aggregate traffic; in the LAN context, see [31]; for WANs, see [25, 5, 16, 9, 30]. More precisely, aggregate packet-level network traffic is (asymptotically) self-similar, i.e., exhibits fractal-like scaling behavior over time scales on the order of a few hundreds of milliseconds and larger, if and only if the durations (in seconds) or sizes (in bytes) of the individual sessions or connections that generate the aggregate traffic have a heavy-tailed distribution with infinite variance, i.e., range from extremely short (small) to extremely long (large). This ability to explain self-similarity of aggregate traffic streams has de-mystified fractal traffic modeling and has opened up new opportunities for queueing and performance analysis; in particular, it has provided new insights into how self-similarity (through the underlying heavy-tailed connections) can impact network performance, both qualitatively and quantitatively. It has also led to the realization that the self-similarity property of the aggregate traffic does *not* seem to depend on the connections' local traffic characteristics, i.e., on how the individual packets within a connection are sent over the network.

Yet, because of the predominant protocols and end-to-end congestion control mechanisms that exist in today's Internet and that determine the flow of packets at the different layers in the TCP/IP protocol hierarchy, networking researchers have argued that to provide a complete description of network traffic, these local traffic characteristics should not be ignored and have asked the question "Where does the impact of the network show up?" In this paper, we use a number of different high time-resolution packet-level traffic traces collected from both a corporate and Internet Service Provider (ISP) WAN environment to demonstrate that (i) the impact of the network on the traffic shows up when studying network traffic over small time scales, from a few hundreds of milliseconds and downward; (ii) the empirically observed local traffic characteristics are consistent with *multifractal* scaling; (iii) there is a plausible physical "explanation" for the multifractal nature of measured Internet WAN traffic over small time scales; and (iv) the multifractal finding suggests a class of parsimonious models that provide a more complete and accurate description of actual data traffic than is available to date and hence allows for a systematic investigation of a wide range of queuing/networking-related performance issues.

While multifractals are new to the networking area, they have been applied in the past to such diverse fields as the statistical theory of turbulence, the study of strange attractors of certain dynamical systems, and more recently, to physically based rain and cloud modeling; see for example [8, 14] and references therein. In the present context, multifractals

[*]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.

Appeared in: *Computer Communication Review*, Vol. 28. No. 4, Proceedings of the ACM/SIGCOMM'98, September 1998, Vancouver, Canada pp. 42–55, 1998.

extend and refine in a natural way the previously observed fractal or self-similar behavior in measured network traffic. Indeed, while self-similarity or, more generally, monofractal scaling, is characterized by a single scaling law that holds globally in time and essentially involves only one parameter, the Hurst parameter, multifractals allow for time-dependent scaling laws and hence offer great flexibility in describing irregular phenomena that are localized in time. The latter are typically caused by network-specific mechanisms that operate on small time scales and—depending on the state of the network—can have a more or less severe impact on the packet dynamics within individual connections.

From a networking perspective, the special appeal of multifractals lies in their close connection to certain *multiplicative processes* or *cascade models*. Motivated by the explicit hierarchical structure of modern data networks, it is plausible to view WANs or other networks, together with their protocols and controls, as specifying the mechanisms and rules of a process that fragments units of information at one layer in the networking hierarchy into smaller units at the next layer, etc. Such a fragmentation mechanism is called a cascade; it preserves the mass of the initial set at each stage of the construction, the rules for fragmentation make up what is commonly referred to as the *generator* of the cascade, and the limiting object or multiplicatively generated multifractal is a mathematical construct that describes the highly irregular way the connection's total mass (i.e., number of bytes or packets) has been redistributed during this fragmentation procedure over the lifetime of the connection. To validate this hypothesis of networks acting as cascades, we develop and use a set of wavelet-based analysis and inference tools that are tailor-made for the multiplicatively generated class of multifractals considered in this paper. We provide empirical evidence that measured WAN traffic conforms to the proposed cascade model, is consistent with the intricate local irregularities exhibited by the corresponding multiplicatively generated multifractal, and cannot be completely described by self-similar (i.e., monofractal) or other strictly second-order traffic processes.

Here, by a strictly second-order process, we mean a complete description of traffic in terms of its first- and second-order statistical characteristics, i.e., its marginal distribution and autocorrelation function (or equivalently, its spectral density). E.g., a Gaussian marginal distribution and an autocorrelation function of the form $r(k) = 2^{-1}(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}), k \geq 1, 0 < H < 1$ completely describes the self-similar processes known as *fractional Gaussian noises*; similarly, a Poisson process is fully characterized by requiring the marginal distribution to be Poisson and the autocorrelation function to be identically zero. On the other hand, using exclusively first- and second-order statistical characteristics to specify an asymptotically self-similar process with non-Gaussian marginals results only in an incomplete description of the process—higher-order statistical properties (e.g., expressions of the form $\mathbf{E}[X_k X_l X_m], k \neq l \neq m$) have to be specified to provide a complete statistical description of the process. In fact, we will show that the presence of non-trivial higher-order statistics in a traffic process is closely related to a non-degenerate multifractal scaling behavior. In this sense, multifractals offer great promise for providing a sufficiently complete description of network traffic in cases when a specification in terms of purely second-order statistics is inadequate and may lead to erroneous or misleading conclusions about expected implications for network performance.

Empirical evidence in support of within-connection or local traffic characteristics in measured WAN traces that can be traced to the protocol architecture of IP networks has been reported in the original comprehensive analysis of WAN traffic by Paxson and Floyd [25], and more recently, in work by Feldmann et al. [9]. The work by Paxson and Floyd [25] is closest in spirit to our present study and concerns some aspects of the local traffic structure of individual connections (e.g., TELNET and FTP). Technically, our paper is related to the works by Abry and Veitch [1] and Feldmann et al. [9] in the sense that we also rely crucially on wavelet-based techniques. However, we pursue the wavelet-based analysis of network traffic one step further and develop and illustrate tools that can be used for statistical inference problems related to cascade models and their limiting multifractals. Finally, our work is closely related to that of Riedi and Levy-Vehel (e.g., see [26, 18] for TCP/IP traces; for ATM WAN traces, see Mannersalo and Norros [19]), who originally advocated the use of multifractals for network traffic modeling; though, for an earlier discussion, see also [28]. In contrast to Riedi and Levy-Vehel's work, this paper attempts to present, motivate and explain multifractals in the networking context and qualitatively discusses the relevance and impact of multifractal scaling in measured data traffic on network performance-related problems.

The remaining part of the paper is structured as follows. In Section 2, we use measured WAN traces to motivate the use of multifractals as plausible models for WAN traffic; we introduce wavelets as our main mathematical technique, discuss the notions of global vs. local scaling, and give an intuitive definition of monofractals and multifractals. Section 3 provides the mathematical framework for our proposed cascade-based approach to modeling the multifractal nature of WAN traffic and presents the main results of the corresponding wavelet-based analysis. In Section 4 we present empirical evidence in favor of our assumption that IP networks act as cascades and we discuss a workload model for data traffic that captures both the multifractal (i.e., small time scaling properties) as well as the asymptotically self-similar (i.e., large time scaling properties) nature of measured WAN traffic. We conclude in Section 5 by illustrating with some examples the potential impact and relevance of our findings for network performance analysis and traffic management.

Short description of data sets: Throughout this paper we use the following high-quality data sets (i.e., packet drops reported by *tcpdump* were negligible and other causes for drops have been identified to be negligible as well; high time stamp accuracy of about 10-100 μsec). The trace DIAL1 was gathered from an FDDI ring (with typical utilization levels of 5-10%) that connects about 420 modems to the rest of the Internet. Although we collect every packet seen on the FDDI ring, DIAL1 contains (bidirectional) *modem user traffic only*. It was collected on July 23, 1997 between 19:02 and 23:43 and consists of a total of 12,870,502 packets and 4.212 Gbytes. A 1-hour segment of this trace (from 22:00 to 23:00), referred to as DIAL2, contains 2,752,779 packets (a total of 8,719,659 packets were seen on the FDDI ring during this period). The trace DIAL3 was collected in the same location as DIAL1, on July 22, 1997 between 22:00 and 23:00, and contains modem user as well as non-modem user traffic totaling 8,910,014 packets. A second dataset was gathered off a T3 backbone link of the same ISP; the trace BACKB was collected on December 7, 1997 between 21:27 and 21:49 and consists of a total of 9,919,939 packets and 2.617 Gbytes. Finally, a third, non-ISP related dataset, consisting of 3,903,350 packets and 1.131 Gbytes, was collected off an Ethernet connecting AT&T Labs-Research at Florham Park, NJ to the Internet via a fractional T3 connection (3Mbps). The trace ATTLAB1 was collected on October 19, 1997, between 12:15 and 20:05; the 1-hour segment (16:00 to 17:00) is referred to as ATTLAB2, and the 17:00-18:00 hour segment by ATTLAB3.

2 Wavelets and the nature of WAN traffic

In this section we introduce wavelets as our main mathematical technique for detecting and identifying global and local scaling of measured network traffic. We explain the notions of monofractal and multifractal on an intuitive level, and relate them to the concept of self-similarity.

2.1 The discrete wavelet transform

The ability of wavelets to "localize" a signal in both time and scale makes them an attractive mathematical tool for many applications in the physical and engineering sciences; see [15] for an introduction to wavelets, and [6] for a more mathematical treatment of the subject. Wavelets provide the mathematical framework in a multiresolution analysis (MRA) that formalizes the notion of coarse and fine approximations and gives meaning to the increment in information needed to pass from one level of approximation to another. The key feature of an MRA is that we can write an approximation, X_j, of a signal X, at scale j (with resolution 2^j) as the sum of a coarser approximation X_{j+1} at scale $j+1$ (with resolution 2^{j+1}) and the "detail" $D_{j+1} = X_j - X_{j+1}$; i.e., the difference between these two approximations. We may iterate this procedure, writing the approximation at scale $j+1$, X_{j+1}, as a sum of a coarser approximation X_{j+2} and the difference $D_{j+2} = X_{j+1} - D_{j+2}$, and so on: $X_j = X_{j+1} + D_{j+1} = X_{j+2} + D_{j+2} + D_{j+1} \ldots$.

More formally, an MRA guarantees the existence of a *scaling function* ϕ (which is used to express the approximation) and a *wavelet* ψ (which is essential for the definition of the details) such that a signal X can be written as

$$\begin{aligned} X &= \sum_{k\in\mathbf{Z}} \langle X, \phi_{0,k}\rangle \phi_{0,k} + \sum_{j\le 0}\sum_{k\in\mathbf{Z}} \langle X, \psi_{j,k}\rangle \psi_{j,k} \qquad (1) \\ &= \sum_{j\in\mathbf{Z}}\sum_{k\in\mathbf{Z}} \langle X, \psi_{j,k}\rangle \psi_{j,k}, \end{aligned}$$

where $\phi_{j,k}(t) = 2^{-j/2}\phi(2^{-j}t-k)$ and $\psi_{j,k}(t) = 2^{-j/2}\psi(2^{-j}t-k)$ are the shifted and dilated versions of the scaling function and the wavelet, respectively. For example, the wavelet ψ given by $\psi(t) = 1$ if $t \in [0, 1/2)$, $\psi(t) = -1$ if $t \in [1/2, 1)$ and $\psi(t) = 0$ otherwise, is known as the *Haar wavelet*, and the corresponding scaling function ϕ is given by $\phi(t) = 1$ if $t \in [0, 1)$ and 0 otherwise. The representation (1) is called the *wavelet decomposition* of the signal X, and $d_{j,k} = \langle X, \psi_{j,k}\rangle$, the inner product of X with $\psi_{j,k}$ is commonly referred to as the *wavelet coefficient* at scale j and time $2^j k$. The quantity $|d_{j,k}|^2$ measures the amount of energy in the 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. The set of all wavelet coefficients $\{d_{j,k} : j \in \mathbf{Z}, k \in \mathbf{Z}\}$ is called the *discrete wavelet transform* (DWT) of the signal X and its key feature is that it contains the same information as the signal X; i.e., it allows us to reconstruct X completely from its wavelet coefficients by setting $X(t) = \sum_{j\in\mathbf{Z}}\sum_{k\in\mathbf{Z}} d_{j,k}\psi_{j,k}(t)$.

Intuitively, the discrete wavelet transform divides a signal into different frequency components and analyzes each component with a resolution matched to its scale. Letting $k(t_0, j)$ specify those wavelet coefficients at scale j that are influenced by the value of the signal X at time t_0; i.e., $d_{j,k(t_0,j)}$ is in the "cone of influence associated with the point t_0," we can use the wavelet coefficients to study directly either scale- or time-dependent properties of a given signal X. For example, by fixing a given scale j and studying X at that scale across time, we can obtain information about the scaling behavior of X, as a function of j. On the other hand, fixing a point t_0 in time and investigating the wavelet coefficients $\{d_{j,k(t_0,j)} : j \le 0\}$ across finer and finer scales results in powerful techniques for investigating the nature of local irregularities or singularities in the signal, as a function of t_0. While the former method results in scaling properties that hold globally (across the whole signal), the latter technique captures the idea behind the notion of "the wavelet transform as a mathematical microscope" (e.g., see Arneodo [2]), provides (local) information about the fine structure of the signal at a given point in time, and thus opens up new ways for studying the intrinsic nature of "bursts" in measured network traffic.

2.2 DWT and scale-localization: Self-similarity

We first illustrate that wavelets with their built-in scale-localization ability provide an ideal mathematical tool for investigating the scaling behavior of self-similar processes across all (a wide range of) time scales.[1] Abry and Veitch [1] have shown that if X is a self-similar process with Hurst parameter $H \in (1/2, 1)$, then the expectation of the energy E_j that lies within a given bandwidth 2^{-j} around frequency $2^{-j}\lambda_0$ is given by

$$\mathbf{E}[E_j] = \mathbf{E}\left[\frac{1}{N_j}\sum_k |d_{j,k}|^2\right] = c|2^{-j}\lambda_0|^{1-2H} \qquad (2)$$

where c is a prefactor that does not depend on j, and where N_j denotes the number of wavelet coefficients at scale j. By plotting $\log_2 E_j$ against scale j (where $j = 1$ is the finest scale and $j = N > 1$ is the coarsest) and identifying scaling regions, breakpoints and non-scaling behavior, we have an unbiased scaling analysis of a given signal X that is simple, computationally efficient and informative. For example, the scaling analysis of a signal which is exactly self-similar will yield a linear plot of $\log_2 E_j$ vs. j for all scales; for a fractional Gaussian noise trace with $H = 0.7$ and for a Poisson trace (i.e., $H = 0.5$), the corresponding scaling plots are shown in Figure 1 (left). On the other hand, for an asymptotically self-similar signal a linear relationship between $\log_2 E_j$ and scale j will be apparent only for large times or scales.

Figure 1 shows the scaling analysis for five different traffic traces: August'89 Bellcore Ethernet LAN trace (left), an 1994 LBL WAN trace (left), the WAN trace ATTLAB2 (right), and WAN traces DIAL3 and BACKB (middle). The LAN trace shows an approximate linear relationship for a wide range of scales, with an estimated Hurst parameter of about 0.8 (consistent with previously reported estimates, e.g., [17, 1]), but with some deviations from linearity at the very small scales. All the WAN traces show a scaling behavior that is, in general, more complex than that of the LAN trace: well-defined large-time scaling regions

[1] For the global scaling analysis presented in this subsection, we use the Daubechies wavelets [6].

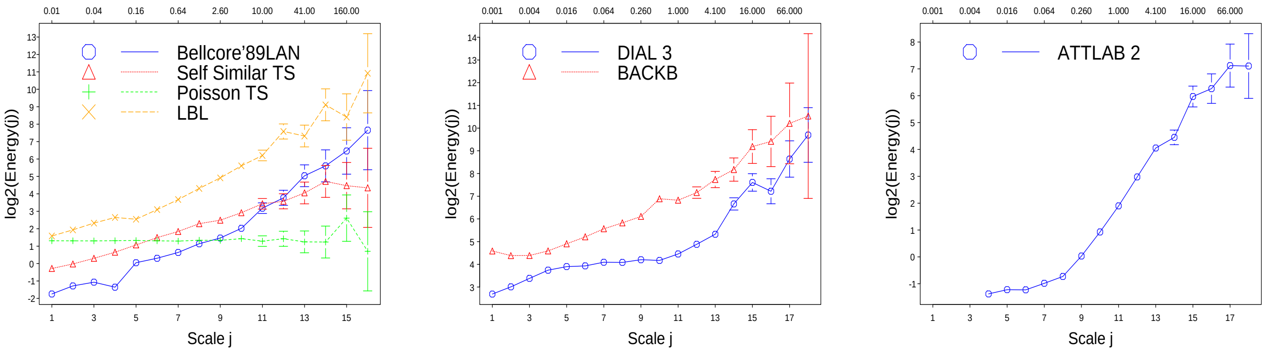

Figure 1: Global scaling analysis of packet-level LAN and WAN and test traces: exactly self-similar trace with $H = 0.7$, Poisson trace, i.e., $H = 0.5$, Bellcore'89 LAN, and LBL'94 WAN (left); DIAL3 (middle) and BACKB (middle); and ATTLAB2 (right). Note the different labeling at the bottom and the top of the plots: scale j (bottom, $j = 1$ is finest scale); actual time (top, in seconds). The vertical bars at each scale represent 95% confidence intervals for $\log_2(E_j)$.

where the relationship is roughly linear—confirming the previously reported asymptotically self-similar nature of WAN traffic (e.g., see [25] for the 1994 LBL WAN trace); apparent breakpoints at scales on the order of a few hundreds of milliseconds; and complex small-time scaling behavior that is distinctly different from the large-time scaling features. For a more comprehensive study of the global scaling properties of measured WAN traffic, see [9].

2.3 From self-similarity to multifractals

Figure 1 gives a concise picture of our current understanding of WAN traffic dynamics: measured WAN traffic is consistent with asymptotic self-similarity or large time scaling and exhibits small time scaling features that are very different from those observed over large time scales. To provide an adequate and complete description of WAN traffic, it is therefore necessary to get a handle on those small time scaling features. To this end, results by Erramilli et al. [7] suggest that networking mechanisms operating on small time scales are a possible explanation for the observed small time scaling behavior in measured WAN traffic. Such mechanisms can cause the traffic to exhibit pronounced local variations and irregularities.

To quantify these local variations in the traffic at a particular point in time t_0, we turn to the traffic rate process, the number of packets or bytes in an interval $[t_0, t_0 + \delta t]$ of length δt at t_0. *We say that the traffic has a local scaling exponent* $\alpha(t_0)$ *at time* t_0 *if the traffic rate process behaves like* $(\delta t)^{\alpha(t_0)}$ *as* $\delta t \to 0$. Note that $\alpha(t_0) > 1$ corresponds to instants with low intensity levels or small local variations, while $\alpha(t_0) < 1$ is found in regions with high levels of burstiness. Informally, signals with $\alpha(t_0) = H$ at all instants t_0 are called *monofractal* (and include exactly self-similar processes) while signals with nonconstant scaling exponent $\alpha(t_0)$ are called *multifractal*.

Unfortunately, to obtain detailed information about the local variations of traffic at a particular point in time, traditional statistical inference techniques—including the scaling analysis presented in Section 2.2—are inadequate because they are global in nature; i.e., they provide information that holds across the whole trace. Instead, we rely here on the ability of wavelets to serve as "mathematical microscope" with which we can zoom in and examine the variations in a trace at a particular point in time. Because the DWT yields a complete reconstruction of a given signal, it can be used to recover the local irregularities in the traffic and, in particular, to estimate the local scaling exponents. Roughly speaking, if the signal or trace X has a local scaling exponent $\alpha(t_0)$ at t_0, then for large negative j-values (small scales), the wavelet coefficients affected by $X(t_0)$ behave like $d_{j,k(t_0,j)} \approx 2^{j(\alpha(t_0)+1/2)}$ [6], where for two functions f and g, $f(j) \approx g(j)$ means that $\lim_{j \to -\infty} f(j)/g(j) = \text{const},$.

To illustrate this local scaling property in measured WAN traffic, we employ a naive wavelet-based heuristic for a crude estimation of the scaling exponent associated with each point in the trace.[2] Then depending on the value of the scaling exponent, we pick a gray scale and plot the corresponding observation in the chosen shade of gray. The darker the shade of gray, the smaller the scaling exponent or the "burstier" the signal at that point in time; lighter shades of gray correspond to instants with larger scaling exponents or "lull" periods in the signal. Note that in theory, self-similar or monofractal scaling should result in one shade of gray throughout the entire trace, but in practice some variability in the gray-shading has to be expected. To get a sense for how much variability can be expected, the top plot in Figure 2 shows the results of applying our scaling heuristic to an exactly self-similar trace and serves as an example against which we can calibrate deviations from monofractal scaling. For example, the remaining plots in Figure 2 depict the local scaling behavior for a segment of the Bellcore'89 LAN trace averaged over 10 milliseconds, a segment of the WAN trace DIAL2 averaged over 500 milliseconds, and finally a segment of the WAN trace DIAL3 at the 1 millisecond time scale. Visually, the first two plots show a similar behavior, both in terms of the predominant shades of gray as well with regard to the relatively smooth transitions from one gray scale to another. More importantly, the two plots suggest that differences in the local scaling exponents (as expressed by the different shades of gray) in the LAN trace are well within the natural variability associated with the limited local scaling behavior of an exactly self-similar trace. In contrast, the 1 millisecond WAN trace (bottom plot) shows clear signs of multifractal scaling behavior: instants with dark shades of gray across the whole trace, abrupt transitions from dark- to light-shaded periods and vice versa, and a much less smooth overall texture than the top two plots. Note however that when averaging WAN traffic over 500 millisecond intervals, it becomes more LAN-like or self-similar, though still with a

[2] To pick out the bursty regions, we threshold the wavelet coefficients of the signal, keeping only those with magnitudes exceeding a given value. Then we calculate the local scaling exponents of the reconstituted signal via a linear regression of $\log d_{j,k(t_0,j)}$ versus j.

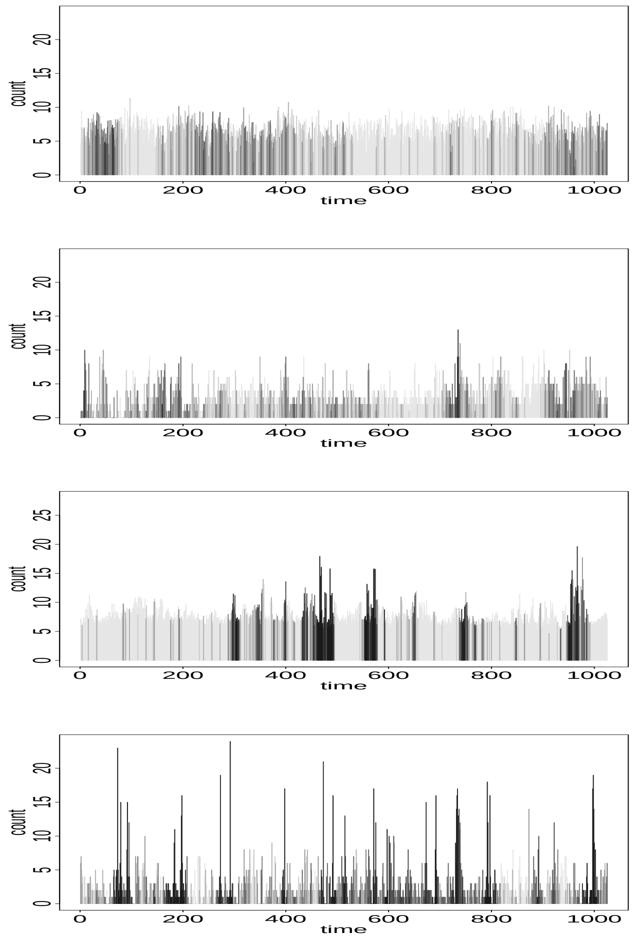

Figure 2: Local scaling analysis of packet-level data traffic; different shades of gray indicate different magnitudes of the local scaling exponents at the different point in the traffic trace (black for small scaling exponents or "bursty" instants), light for large scaling exponents or "lull" periods). From top to bottom: (exactly) self-similar traffic, Bellcore'89 LAN trace, WAN trace DIAL2 averaged over 500 msec, WAN trace DIAL3 at the 1 msec time scale.

few instances where the local scaling exponents exceed the variability associated with a monofractal trace. Thus, the bottom two plots visualize the previously observed asymptotically self-similarity property of measured WAN traffic: the small time scaling properties appear to be consistent with multifractal scaling behavior, and when aggregating over larger time scales, the large time scaling features start to conform to monofractality.

2.4 DWT and time-localization: Multifractals

Figure 2 depicts visually the distinctly different local scaling behavior of measured WAN traffic, of an exactly self-similar trace, and of measured LAN traffic. Motivated by this visually appealing heuristic for qualitatively assessing the local scaling behavior of a given trace, our goal is to develop a quantitative approach that is more rigorous and allows us to draw statistically sound conclusions about the local scaling behavior (e.g., whether some scaling exponents occur more frequently than others, and if so, which one).

To this end and to build intuition, we first assume that there is only one scaling exponent α for the entire trace; i.e., the trace is monofractal. Then the wavelet coefficients all behave like $d_{j,k} \approx 2^{j(\alpha+1/2)}$ as j tends to $-\infty$. In this case evaluating the so-called wavelet-based *partition function* $S(q,j)$, defined by summing across each level j the qth moments (with $q \geq 0$) of the absolute value of the normalized wavelet coefficients $\tilde{d}_{j,k} = 2^{-j/2} d_{j,k}$; i.e., setting

$$S(q,j) = \sum_k |\tilde{d}_{j,k}|^q, \tag{3}$$

we obtain $S(q,j) \approx 2^{-j} 2^{j\alpha q} = 2^{-j(1-\alpha q)}$. Note that for $q = 0$, $S(0,j) = N_j$, the number of wavelet coefficients at scale j, and for $q = 2$, $S(2,j)$ denotes the energy E_j at scale j considered in the global scaling analysis in Section 2.2 (up to a normalization factor). Intuitively, for $q > 2$, the function $S(q,j)$ takes into account the effects of higher-order statistics that may be present in a trace and hence may be contained in the DWT of the trace. Moreover, because wavelet coefficients tend to decorrelate quickly within a given scale as well across different scales (for the specific case of fractional Brownian motion see [29] and for more general settings see [21]), it can be expected that hardly any information about possibly strong correlations within the trace is lost by defining the partition function $S(q,j)$ as in (3).

Next, to examine the scaling behavior of $S(q,j)$ as the time scale or resolution level becomes finer and finer (i.e., $j \to -\infty$), we consider the corresponding wavelet-based *structure function* $\tau(q)$ defined as the scaling exponent of $S(q,j)$, as $j \to -\infty$; that is,

$$\tau(q) = \lim_{j \to -\infty} \frac{\log S(q,j)}{j \log 2}. \tag{4}$$

In other words, we check whether or not the partition function behaves like $S(q,j) \approx 2^{j\tau(q)}$ as we look at finer and finer time scales (i.e., $j \to -\infty$). For the example at hand, it is easy to see that $\tau(q) = \alpha q - 1$; i.e., the structure function of a monofractal signal is linear in q. In particular, if the trace is self-similar with Hurst parameter H, then $\tau(q) = Hq - 1$ and H can be easily inferred from the structure function.

For a slightly more complicated example that shows that $\tau(q)$ indeed contains information about the frequency of occurrence of different local scaling exponents, assume now that $100\gamma\%(0 < \gamma < 1)$ of the trace has a local scaling exponent α_1 and the other $100(1-\gamma)\%$ of the trace scales with an exponent $\alpha_2 \neq \alpha_1$. Then, for large negative j-values (fine resolution levels), $100\gamma\%$ of the wavelet coefficients $d_{j,k}$ scale like $2^{j(\alpha_1+1/2)}$ and the other $100(1-\gamma)\%$ like $2^{j(\alpha_2+1/2)}$; the actual location of these two types of coefficients within level j is not crucial. A simple calculation shows that in this case, the partition function behaves like

$$S(q,j) = 2^{-j(\gamma - \alpha_1 q)} + 2^{-j((1-\gamma)-\alpha_2 q)},$$

and that the structure function $\tau(q)$ is determined by the relative strengths of the local scaling exponents α_1 and α_2; in fact, identifying the leading term in the limiting expression $\lim_{j\to-\infty} \log S(q,j)/(j \log 2)$, we obtain

$$\tau(q) = \min(\alpha_1 q - \gamma, \alpha_2 q - (1-\gamma)).$$

In other words, because the trace contains more than one local scaling exponent, $\tau(q)$ is no longer linear in q but is instead a concave function of q. For this example, the structure function is in fact piecewise linear, following one of the linear functions in the expression for $\tau(q)$ for some values of q and then following the other linear function for the larger q values; furthermore, the location of the breakpoint reflects the composition of the scaling exponents in the trace.

These simple examples can easily be generalized to account for a finite number of different scaling exponents α_i in the trace, where a "histogram" $f(\alpha_i)$ measures the number of instants in the trace which have local scaling exponent α_i. For example, in the previous case of two different scaling exponents, we have $f(\alpha_1) = \gamma$ and $f(\alpha_2) = 1 - \gamma$. In turn, this motivates the precise relationship that exists between the "histogram" $f(\alpha)$, commonly referred to as the *multifractal spectrum* of the signal, and the partition function $\tau(q)$:

$$\tau(q) = \min_{\alpha}(\alpha q - f(\alpha)). \tag{5}$$

The above examples also allow us to properly interpret the $\tau(q)$ function derived from the DWT of a given signal. A more or less linear $\tau(q)$ function is consistent with monofractal scaling and rules out multifractality. On the other hand, the more concave the shape of $\tau(q)$, the wider the range of local scaling exponents found in the signal; in particular, a concave shape of the structure function is consistent with multifractality.

To illustrate the time-localization ability of wavelets to infer mono- or multifractal scaling, Figure 3 shows the results of applying the (Haar wavelet-based) DWT structure function method to a number of WAN and test traces. For each trace, we picked 10 milliseconds as the finest resolution level (i.e., $j = -18$; for the shorter self-similar trace, the finest resolution level corresponds to $j = -15$) and examined the scaling behavior of the partition function $S(q, j)$ over a range of fine resolution levels, i.e., for j-values bigger than -18.[3] The four left plots in Figure 3 show the logarithm of the modified partition function $\log \tilde{S}(q, j)$ against j, for different q-values ranging from $q = 0,\ 4,\ 8,\ 12,\ 16,\ 20$, for the Bellcore'89 LAN trace (top left), an exactly self-similar trace ($H = 0.7$, top middle), and the WAN trace DIAL3 (bottom left) and ATTLAB3 (bottom middle). All partition function plots suggest the presence of well-defined fine-time scaling regions (right-hand side of each plots, ranging over 10 or more of the finest time scales) where reading off the slopes of the different lines, i.e., determining the value of the structure function $\tau(q)$ at different q's, appears to be relatively insensitive to the particular choice of the upper cutoff scale (i.e., coarse time scales or small negative j-values), beyond which different scaling regimes seem to exist. To illustrate how to get the structure functions from the corresponding partition function plots, consider for example, $\tau(8)$ for the Bellcore'89 LAN trace (top right plot); $\tau(8)$ is obtained by estimating the slope of the line labeled "$q = 8$" in the top left partition function plot over scales (x-axis) ranging from, say, $j = -7$ to $j = -18$. The resulting structure functions $\tau(q)$ are depicted in the two right plots in Figure 3. The top right plot shows the $\tau(q)$ functions for the Bellcore'89 LAN trace, the self-similar trace and a Poisson trace and illustrates that all three traces result in linear $\tau(q)$ functions of the form $\tau(q) = Hq - 1$, and are hence fully consistent with monofractal scaling behavior; in fact, one can easily read off the Hurst parameters for each of these three traces from their structure function plots: $H \approx 0.8$ for the LAN trace, $H \approx 0.7$ for the self-similar trace, and $H \approx 0.5$ for the Poisson trace. In contrast, the bottom right plot in Figure 3 shows the $\tau(q)$ functions for the three WAN traces DIAL3, BACKB and ATTLAB3, all of which show indications of nonlinear, i.e., concave shapes that are inconsistent with monofractal behavior and suggest multifractal structure over small time scales. These results confirm our earlier observations that for strictly second-order models such as an exactly self-similar Gaussian process or a Poisson process, multifractal analysis should result in a trivial (i.e., linear) structure function, while the presence of higher-order statistics should be reflected in a more or less pronounced nonlinear structure function. To this end, the LAN trace seems to be adequately described by a purely second-order process, while the WAN traces are not.

3 Structural modeling of WAN traffic: Why multifractal?

In this section, we move beyond the empirical evidence that measured WAN traffic is consistent with multifractal scaling behavior and turn our attention to the question "Why is WAN traffic multifractal?" We will answer this question in two stages. First, we address the above question by claiming that "WAN traffic is multifractal because certain multiplicative cascades lurk in the background." In a second stage, we will investigate in Section 4 the problem of associating multiplicative structure in measured WAN traffic with certain layers in the TCP/IP protocol hierarchy.

3.1 Cascades and multifractals

Informally we say, following Evertsz and Mandelbrot [8], that a process that fragments a set into smaller and smaller components according to some rule, and at the same time fragments the measure or mass associated with these components according to some other rule is a *multiplicative process* or *cascade*. The more formal mathematical construction of a cascade starts with an initial mass M distributed uniformly over the unit interval $I = [0, 1)$. We assume for convenience a dyadic partitioning of I, and in a first stage of the cascade construction, we divide I into the two subintervals $I(0) = [0, 1/2)$ and $I(1) = [1/2, 1)$ and assign mass xM to $I(0)$ and mass yM to $I(1)$. The *multipliers* x and y are chosen according to a particular rule that characterizes the type of cascade and will be specified shortly. Iterating this construction process, we divide each parent interval into its two dyadic subintervals, choose multipliers x and y in agreement with the specified rule, and assign the appropriate mass to the left and right subinterval, respectively. To simplify notation, we denote the dyadic intervals of resolution size 2^{-l} that are generated at the l-th stage of this cascade construction by $I(j_1, \ldots, j_l) = [\ \sum_{k=1}^{l} j_k 2^{-k}, \sum_{k=1}^{l} j_k 2^{-k} + 2^{-l})$, with $j_k \in \{0, 1\}$ and $l = 1, 2, \ldots$.

Cascade models have been especially popular in the statistical theory of turbulence (see references in [8]), and more recently, in the hydrologic and atmospheric sciences [13]. In the networking context, cascades are motivated by the TCP/IP protocol hierarchy and give rise to the conjecture that IP networks act as cascades. Intuitively, this conjecture can be substantiated by considering for example the dynamics of a typical Web session: user clicks result in requests, requests give rise to connections, connections are made up of flows, and flows consist of individual packets. Note that during this fragmentation process, the total number of bytes transmitted during the Web-session is roughly preserved (or grows slightly, due to headers, acknowledgment packets and

[3] Instead of using the partition defined in (3), we relied in our analysis on the numerically more efficient modified partition function $\tilde{S}(q, j) = \sum_{\max} |\tilde{d}_{j,k}|^q$, where the sum is taken over the local maxima of the absolute value of the qth moment of the normalized wavelet coefficients $\tilde{d}_{j,k}$.

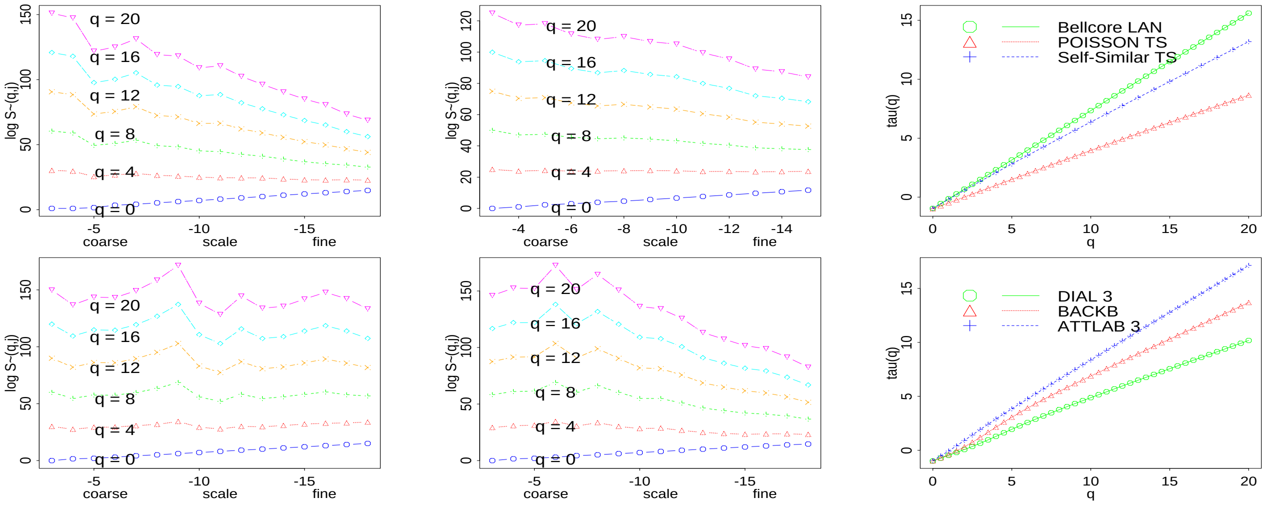

Figure 3: DWT partition function analysis (left and middle) of packet-level LAN, test, and WAN traces; Bellcore'89 LAN (top left), self-similar trace (top middle), DIAL3 (bottom left), and ATTLAB3 (bottom middle). DWT structure function analysis (on the right) of packet-level LAN, test, and WAN traces.

overhead). To satisfy this approximate mass preservation property and to also allow for some degree of randomness in the way mass gets redistributed in the process of this construction, we consider in the following a *semi-random rule* that assigns mass MW to the interval $I(0)$ and mass $M(1-W)$ to $I(1)$, where the "generator" W is a random variable with mean $1/2$, takes on values in $(0,1)$, and is symmetric about its mean. To iterate this procedure, we consider a sequence of random variables $W(j_1,\dots,j_l), l \geq 1$, with a dependence structure given by

$$W(j_1,\dots,j_{l-1},1) = 1 - W(j_1,\dots,j_{l-1},0). \qquad (6)$$

and where, because of the properties of the generator, the random variables $W(j_1,\dots,j_{l-1},0)$ and $W(j_1,\dots,j_{l-1},1) = 1-W(j_1,\dots,j_{l-1},0)$ are identically distributed as W. This construction gives rise to a *semi-random cascade*[4] and generates a collection of measures μ_l (think of the total number of packets or bytes per interval, where l defines the time scale) such that

$$\mu_l\big(I(j_1,\dots,j_l)\big) = MW(j_1)W(j_1,j_2)\cdots W(j_1,\dots,j_l). \quad (7)$$

Note that because of this multiplicative property, the μ_l's or, in our case, the traffic rate processes at fine time scales (i.e., large l) have perforce approximately lognormal marginals. The limiting object generated by a semi-random cascade can be shown to define a genuine multifractal; see [14, 11] and references therein.

3.2 Wavelet analysis of semi-random cascades

For the remaining part of the paper, we will focus exclusively on these semi-random cascades and variations thereof, where the generator W is allowed to change at each stage of the cascade construction in a way to be specified shortly. We summarize here the main results of a (Haar) wavelet-based global and local scaling analysis applied to this class of semi-random cascades (for further details and proofs, see [11]). In effect, we show below that the DWT of semi-random cascades gives rise to a set of analysis and inference tools that allows us to detect and identify the global and local scaling properties of multifractal objects generated by the semi-random cascade.

To start, consider a semi-random cascade with fixed generator W; i.e., W has mean $1/2$, takes on values in $(0,1)$ and is symmetric about its mean. If X denotes the limiting multifractal generated by this semi-random cascade, then X has global linear scaling; that is (see Section 2.2, left-hand side of Eq. (2)), the logarithm of the expected value of the energy E_l in X around level l in the cascade construction of X depends linearly on l (plotted from large l, fine scales, to small l, coarse scales) and has the form

$$\log_2 \mathbf{E}[E_l] = (1 + \log_2 \mathbf{E}[W^2])l + \log_2 \mathbf{E}[(2W-1)^2].$$

Note that the slope $1+\log_2 \mathbf{E}[W^2]$ depends only on $\mathbf{E}[W^2]$, the second moment of the generator. Thus, if we want a non-linear global scaling behavior for X, we need to change the second moment of the generator W at each level (or within a range of levels) in our cascade construction. One way to achieve this is to let $W(j_1,\dots,j_l)$ be equal in distribution to $\lambda_l W + 1/2(1-\lambda_l)$ where $\lambda_l^2 = \mathbf{Var}(W(j_1,\dots,j_l))/\mathbf{Var}(W) \leq 1$. The limiting object X resulting from a semi-random cascade construction with this type of variable generator can be shown to exhibit non-linear global scaling behavior. In particular, if the factors λ_l^2 associated with the generators at each stage in the cascade construction increase (decrease) monotonically as we go to finer and finer time scales, then the slope of the global scaling analysis increases (decreases) monotonically as we move from fine to coarse scales.

Turning to the local scaling analysis of a multifractal X generated by a semi-random cascade with fixed generator W, the DWT structure function $\tau(q)$ defined in (4), that is, the scaling exponent of the partition function $S(q,j)$ given by (3), can be computed as

$$\tau(q) = -1 - \log_2 \mathbf{E}[W^q], \quad q > 0.$$

Moreover, the multifractal spectrum $f(\alpha)$ of X can be obtained from $\tau(q)$ by setting

$$f(\alpha) = \min_q (q\alpha - \tau(q)).$$

[4] Semi-random cascade are also called *conservative cascades* (B.B. Mandelbrot, personal communication, 1998).

The same results hold if the fixed generator W is replaced by a variable generator of the type considered above (i.e., $W(j_1, \ldots, j_l) = \lambda_l W + 1/2(1-\lambda_l)$), with a more complicated expression for the DWT structure function associated with the underlying limiting multifractal. These results make rigorous the arguments given in Section 2.4 for the class of multifractals generated by semi-random cascades. To apply these findings in practice, we must check whether or not a given signal conforms to a semi-random cascade construction.

3.3 Aggregate WAN traffic and semi-random cascades

To check whether or not a semi-random rule for redistributing mass (which we will equate with number of packets per chosen time interval) from a parent to its child on the left and to the one on its right is consistent with a given traffic trace, we describe in the following a procedure that essentially inverts the cascade construction process in order to gain information about its generation mechanism.

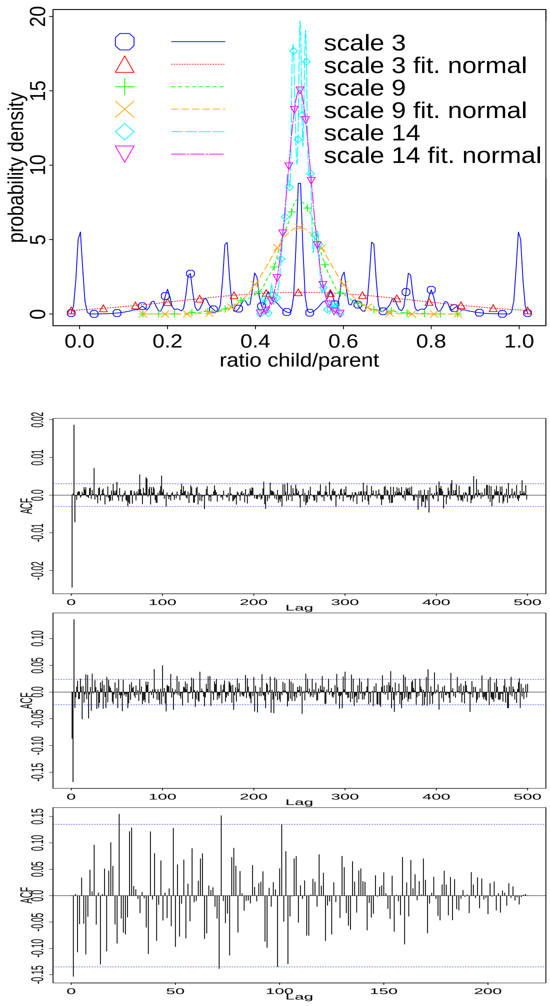

Figure 4: Inferring the generator of a semi-random cascade: density of the ratios for DIAL2 for levels 3, 9, and 14 in the inverse cascade procedure (top); autocorrelations of the ratios for DIAL2 for levels 3, 9, and 14 (the horizontal lines define the 95% confidence band for the autocorrelations to be statistically significant).

The procedure consists of fixing a fine time scale Δ (say, 1 or 10 msec intervals) and summing over non-overlapping blocks of children of size 2, thereby calculating the "parent" time series representing the number of packets per time unit of size 2Δ, and iterating. In effect we obtain the time series representing the total number of packets that each parent redistributed to its two children. We then check the properties of the empirical distribution of the ratios *number of packets in the left (child) interval* divided by *number of packets in the corresponding parent interval* and use this information to judge the appropriateness of a semi-random rule and to infer the underlying generator W for the semi-random cascade construction.

Figure 4 shows the results of applying this procedure to dataset DIAL2. The top plot depicts the empirical probability density functions of the ratios for a number of selected stages in the process of inverting the cascade construction, together with their fitted truncated normal distributions; for the same stages, the subsequent three plots give the empirical autocorrelation functions for the corresponding ratios. Figure 4 suggests that across the different levels in the construction, the empirically observed properties of the ratios agree reasonably well with a semi-random rule: the density plots conform to a generator W that is symmetric around $1/2$ (e.g., a truncated normal), and the autocorrelation plots indicate no significant dependence across a given level, except for some small yet statistically significant small-lag correlations. Moreover, the plot in Figure 5 labeled DIAL2 which gives the empirical standard deviation of the ratios (note log-scale on y-axis), as a function of scale (i.e., level in the cascade construction), suggests a variable generator W with a standard deviation that increases monotonically when moving from coarse scales to finer scales. Similar conclusions about the appropriateness of an underlying semi-random cascade construction with a variable generator hold for the other WAN traces, but only the plots of the empirical standard deviation of the corresponding ratios are shown in Figure 5.

There is nothing that prevents us from applying this procedure to traces that are not generated via an underlying semi-random cascade construction. Thus, the question arises how to use this procedure to identify such cases and how to distinguish them from those that are indeed consistent with an underlying semi-random cascade construction. To this end, we applied the procedure of inverting the cascade construction to the dataset BACKB, and then generated a synthetic trace according to a semi-random cascade construction; as our generator W, we picked a truncated normal on $(0, 1)$, symmetric around $1/2$, and we changed the variability of W according to the plot of the empirical standard deviation of the ratios labeled BACKB in Figure 5. The resulting trace agrees favorably with the data, not only with respect to the traditionally and often exclusively used measures (i.e., their first- and second-order statistics; not shown here), but also with respect to their structure functions $\tau(q)$ which capture the multifractal scaling properties of the data and the synthetic trace, respectively; see the plots labeled *trace data* and *cascade trace* in Figure 5. In stark contrast, when performing the same experiment for a trace generated from a Poisson process (for a plot of the corresponding empirical standard deviation of the ratios resulting from the inverse cascade procedure applied to a Poisson trace, see line labeled *Poisson* in Figure 5 (top)), the difference shows up very clearly when analyzing the multifractal scaling behavior of the trace data and the synthetic trace—see the $\tau(q)$-functions labeled *Poisson* and *cascade Poisson* in Figure 5.

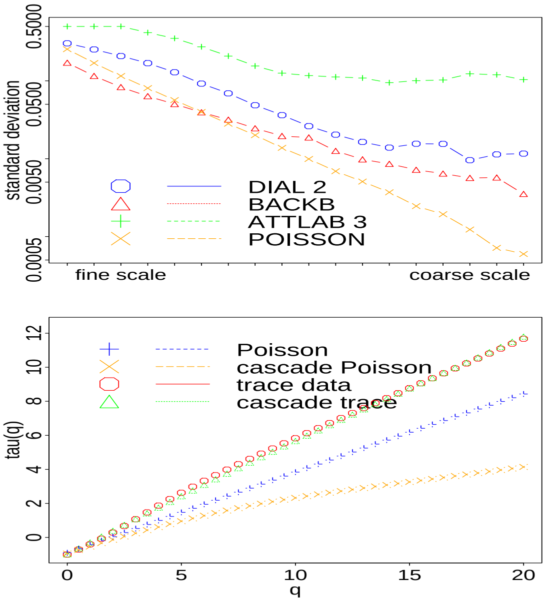

Figure 5: Variability of the ratios, as a function of the level in the inverse cascade procedure for DIAL2, BACKB, ATTLAB3, and a Poisson trace (top). Structure functions $\tau(q)$ for BACKB (at 10 millisecond resolution) and a Poisson trace, and the corresponding cascade traces.

Similar results hold when the Poisson process is replaced by, for example, a self-similar trace.

4 Networks as cascades

The results in Section 3.3 raise a fundamental networking question: "Why does WAN traffic appear to be multiplicatively generated?" We do not answer this question directly, but instead identify below levels in the protocol hierarchy where the multiplicative structure is more apparent than at other levels. Thus, by isolating those aspects of WAN traffic that conform to a multiplicative structure from those that do not, this study represents a first step in explaining how and why data networks act as cascades.

4.1 Additive vs. multiplicative

The mathematical results and physical explanations of the observed self-similarity or monofractal nature of measured traffic traces state explicitly that self-similarity is an *additive* property of network traffic. That is, self-similarity arises from aggregating many *ON/OFF*-streams [31] or from superposing many renewal-type connections (appropriately "thinned" so that the resulting connection arrival process is Poisson), provided the individual *ON/OFF*-periods or connection durations/sizes exhibit extreme variability (i.e., are heavy-tailed with infinite variance) [16, 30]. As such, self-similarity is plausibly the result of user behavior (e.g., dynamics of web browsing), application-specific features (e.g., layout of web pages), and the inherent properties of the objects (e.g., sizes of text, picture, video, audio files) that are sent across the network. In particular, these findings imply that the precise nature of the local traffic structure within individual *ON*-periods or connections is not essential for self-similarity of the aggregate traffic stream. Being additive in nature, aggregate network traffic will be approximately normal when viewed over sufficiently large time scales, provided certain weak conditions on the individual traffic streams (e.g., finite variance of their rates) hold for the central limit theorem to apply.[5]

On the other hand, the observed multifractal nature of network traffic over small time scales and the empirical evidence presented in Section 3.3 in support of an underlying cascade mechanism implies that over those fine time scales, network traffic is *multiplicatively* generated. In other words, at the microscopic level, the traffic rate process (e.g., number of packets per small time unit) has an approximate lognormal shape because it is the product of a large number of more or less independent "multipliers" (see (7)).[6]

These observations give rise to the questions of where in the network the multiplicative structure can be found most easily, how the multipliers can be explained in a networking context, and why network traffic might be multiplicatively generated. Leaving the last two questions for future work, we tackle in this section the first question and identify layers in the protocol hierarchy of IP network where the multiplicative structure seems to dominate. To contrast, we also illustrate the cases, usually found at the higher levels in the protocol hierarchy, where the additive aspects of network traffic start to show up and ultimately to dominate the multiplicative structure.

4.2 IP flows and packets within IP flows

When we try to identify aspects of network traffic that are not affected in a major way by additive components such as its global connection-level characteristics, we have to look at layers in the TCP/IP protocol hierarchy where the influence of user behavior and of application-related peculiarities is no longer dominant but where the network, through its end-to-end congestion control algorithms and other mechanisms, determines the flow of packets across the network. An obvious candidate for this purpose is the TCP layer where we can study global traffic characteristics related to the additive nature of network traffic (e.g., the arrival process of TCP of connections, distribution of the connection durations or sizes) as well as the local traffic characteristics of individual TCP connections (e.g., the packet arrival patterns within individual TCP connections). To this end, we define a *port-to-port flow* as *consisting of all packets flowing in either direction between two IP hosts that use the same source and destination port numbers and that are separated in time by less than 60 seconds* (see [4, 10]). Notice that port-to-port flows are reasonable substitutes for TCP connections because most packets within a TCP connection are part of a single port-to-port flow; in fact, only those packets within a single TCP connection with idle times longer than 60 seconds will be split among multiple port-to-port flows. One advantage of using port-to-port flows rather than TCP connections is that the former are also applicable to non-TCP traffic such as UDP.

Applying this definition to trace DIAL2 results in 362,371 port-to-port flows, with an average number of packets (bytes)

[5] Additive structure leading to normal distributions is discussed in [22], where a central limit theorem argument for justifying the normal distribution of the height of people is paraphrased by the popular saying "... foot bone 'tached to leg bone, leg bone 'tached to the knee bone, knee bone 'tached to the thigh bone, thigh bone 'tached to the hip bone," etc.

[6] Multiplicative structure giving rise to lognormal distributions is exemplified by adapting Franklin's proverb (see [22]) to the networking setting as follows: "for the want of a packet a flow is needed, for the want of a flow a TCP connection is needed, for the want of a TCP connection an application is needed," etc.

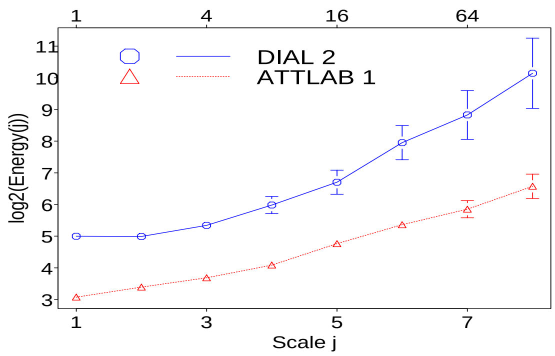

Figure 6: Global scaling analysis for number of port-to-port flows per second for trace DIAL2 and ATTLAB1. Note the different labeling at the bottom and top of the plot: scale j (bottom, $j = 1$ is finest scale); actual time (top, in seconds).

of 35.5 (12,480) and a median of 9 (1466). The global scaling analysis plots for the trace consisting of the number of port-to-port flow arrivals per second for this time series as well as for the WAN trace ATTLAB1 are given in Figure 6 and illustrate the self-similar scaling property of network traffic at the TCP layer over time scales larger than a few seconds. This finding also confirms the additive nature of network traffic at the level of global port-to-port characteristics: the time series of the total number of port-to-port flows is generated by summing over all "sessions" each of which contains a heavy-tailed number of flows (for details, see [9]).

Small time scaling properties that suggest multifractal behavior at the level of individual port-to-port flows for trace DIAL2 are shown in Figure 7. In particular, we consider the top 1000 port-to-port flows which make up more than 40% of the packets or bytes of the overall traffic; the largest flow consists of 216,959 packets, the smallest flow considered has about 1,000 packets. We then apply the inverse cascade procedure described in detail in Section 3.4 to each one of the 1000 flows and focus in Figure 7 on the properties of the logarithm of the empirical standard deviation of the ratios (as a function of scale, where we go from fine to coarse scales; see top plot in Figure 5) associated with the different flows. For each flow, (flow index on the x-axis), we fit a least squares line to the plot of the (logarithm of the) empirical standard deviation of the ratios and plot in Figure 7 (top plot) the slope (lower part) and the resulting R^2-value (upper part) where the latter serves as an indicator for the quality of the linear fit (i.e., R^2-values close to 1 indicate near-perfect fit, low values of R^2 indicate that a linear fit is ill-advised). The top plot in Figure 7 shows that slope-values around -0.2 dominate and that a clear majority of linear fits result in large R^2-values, which we take as strong indication of the appropriateness of a linear behavior (with slope around -0.2) of the (logarithm of the) empirical standard deviation of the ratios of a "typical" port-to-port flow. Picking a port-to-port flow with a generator W^* that shows this "typical" behavior for the corresponding empirical standard deviation of the ratios, the bottom plot in Figure 7 shows the $\tau(q)$ functions resulting from the local scaling analysis of the actual port-to-port flow trace that gave rise to this variable generator, and of four synthetically generated realizations of a semi-random cascade with W^* as its variable generator. The respective structure functions agree reasonably well, which we take as empirical evidence in favor of a multiplicative mechanism that governs the highly irregular traffic rate process within this port-to-port flow and that

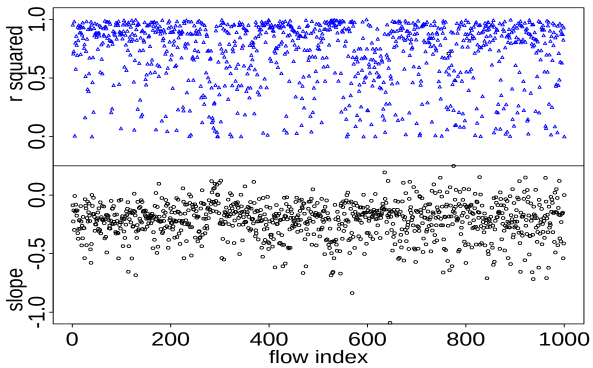

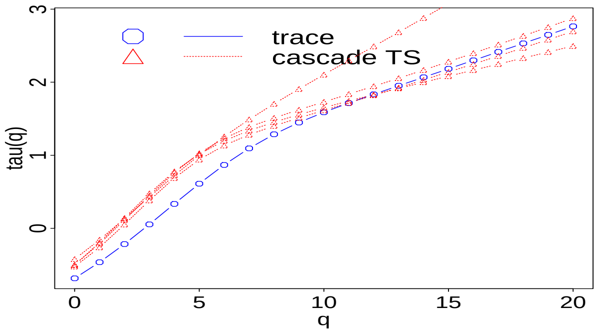

Figure 7: Port-to-port flows and their multifractal scaling analysis (trace DIAL2): Slopes (lower part of top plot) and R^2-values (upper part of top plot) obtained from fitting a least squares line to the logarithm of the empirical standard deviation of the ratios associated with the semi-random generator inferred from each of the top 1000 port-to-port flows (x-axis gives index of port-to-port flow); DWT structure function $\tau(q)$ for a "typical" port-to-port flow trace and four realizations of a semi-random cascade generated via the variable generator corresponding to the "typical" flow (bottom).

gives rise to multifractal scaling behavior of the flow's local traffic dynamics.

Similar results hold (not shown here) when analyzing the other WAN traces or when replacing port-to-port flows by flows that aggregate packets at a slightly higher level in the network hierarchy (e.g., host-application flows defined as consisting of all packets flowing in either direction between two IP hosts and separated in time by less than 60 seconds). Intuitively, replacing port-to-port flows by, for example, host-application flows should not make a significant difference in identifying the multiplicative and additive aspects of data traffic. This replacement simply shifts some higher-layer or user-specific activities to the within-flow packet dynamics. The results of our analysis using host-application flows (not shown here) confirm this intuition and suggest that the empirically observed multiplicative structure within flows offers great promise for gaining an in-depth understanding of the origins and the implications of the multifractal nature of measured WAN traffic.

4.3 Sessions and packets within sessions

For a more dramatic shift of user- and/or application-related activities to local packet traffic characteristics, we consider next network traffic aggregated over user sessions. Before defining what we mean by a "user session", recall that the problem with studying traffic at the level of user sessions is that determining from a packet-level WAN trace what constitutes a user session and when it starts and ends is, in general, difficult. This problem is apparent when we try to extract information regarding Web sessions and it only becomes worse when we attempt to define user sessions that are usually a mixture of many different applications, running either in parallel or sequentially. However, we can avoid these difficulties when we use WAN traces collected from certain ISPs by combining the packet-level WAN trace information with user information contained in other data bases maintained by the ISPs that provide details about every modem call made to the particular ISP, including time of arrival of call, duration (in seconds), size (in bytes) and source IP address. Thus, in an ISP environment, by equating user sessions with modem calls, we avoid the above difficulties and can study in detail modem call-related ISP WAN traffic at the session level. Note that this approach is similar in spirit to the one pursued in [9], where modem calls were used as approximate substitute for Web sessions. By considering the ISP modem traffic trace DIAL1 and studying the global session characteristics such as session arrivals and session duration/size, our findings coincide with the ones reported in [9] and allow us to view sessions as arriving in a more or less Poisson fashion, carrying with them a "workload" (e.g., duration in seconds or size in bytes) that follows a heavy-tailed distribution with infinite variance.

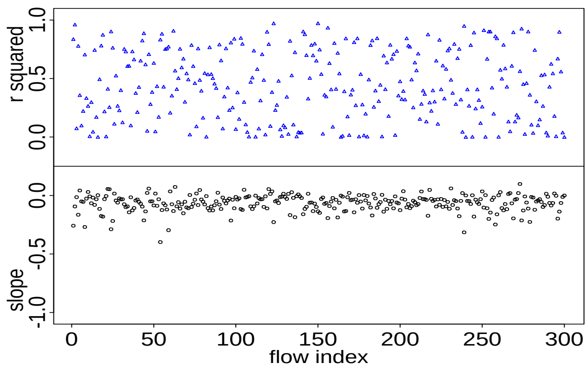

Figure 8: The inverse cascade procedure for user sessions (trace DIAL1): Slopes (lower part of top plot) and R^2-values (upper part of top plot) obtained from fitting a least squares line to the (logarithm of the) empirical standard deviation of the ratios associated with the semi-random generator inferred from each of the top 300 sessions (x-axis gives index of session).

Next, we focus on individual sessions and find that a total of 3,529 different sessions generated the modem call-specific traffic, with a median number of packets per session of about 1K and a mean of about 3.5K packets. To check the multiplicative structure of the within-session packet rate processes, we consider the top 300 sessions which contribute about 56% of the packets (bytes) to the overall trace; the largest user session consisted of 236,071 packets, while the smallest of the 300 sessions contained about 9,000 packets. Then we apply the inverse cascade procedure of Section 3.4 to each of the 300 sessions and obtain the variability plots of the semi-random cascade generators inferred from the 300 session traces. Figure 8 shows the same information as the top plot in Figure 7, with port-to-port flows replaced by user sessions. Note however, that in contrast to the port-to-port flow case, the resulting R^2-values are "all over the picture" (top part of the plot); this suggests that, in general, fitting a least squares line to the (logarithm of the) empirical standard deviation of the ratios associated with an inferred semi-random generator is ill-advised and will perforce yield a very shallow slope of about -0.1 (lower part of the plot). Moreover, in this case, picking a session with a generator W that shows a "typical" behavior for the corresponding empirical standard deviation of the ratios seems hopeless and is not recommended. Instead, we illustrate below with three different examples of user sessions how the interplay between the additive and multiplicative aspects of network traffic can affect the multifractal scaling behavior of within-session packet traffic.

To this end, we consider a user session (consisting of 25,979 packets and 1,100 port-to-port flows, with WWW being the predominant application) whose inverse cascade process yields a plot of the (logarithm of the) empirical standard deviation that decreases quickly for small scales and then shows a more gradual decrease on the larger scales; we denote the corresponding variable semi-random cascade generator by W_1. Intuitively, this behavior can be explained by the fact that the within-session traffic rate process not only reflects within-port-to-port packet dynamics but also application-specific variability due to, for example, the mix of applications within a session or the user activity within a Web session. As our second example, we pick another Web-dominated user session (containing 11,660 packets, twice as long in duration as the first example, with 725 flows) whose associated empirical standard deviation plot of the ratios decreases initially (small scales) but then starts to increase again as we move to larger scales, and finally decrease again for the largest scales; let W_2 denote the corresponding variable generator for the inferred semi-random cascade associated with this user-session. Finally, as a third example, we consider a user-session with a more or less linearly decreasing plot of the (logarithm of the) empirical standard deviation of the ratios (let W_3 denote the corresponding generator), suggesting a port-to-port like within-session structure, with a relative low level of activity at the higher protocol layers. This session contains 5,666 packets or 82 flows and appears to be a gaming application.

To support this intuition about the dynamics of within-session traffic rate processes, Figure 9 uses textured plots (top) to depict the within-session dynamics at the level of individual port-to-port flows (each points corresponds to the arrival of a port-to-port flow, a visible horizontal line indicates the duration of the corresponding flow), for the three generators W_1 (left), W_2 (middle) and W_3 (right). Also depicted in Figure 9 (bottom) are the structure functions corresponding to the traces whose generators are W_1, W_2, and W_3, respectively, and of four independent realizations generated by the corresponding semi-random cascade. On the one hand, Figure 9 confirms our intuition that the session with the low level of higher-layer activities (right) conforms to an underlying multiplicative within-session structure and agrees with the multifractal behavior observed earlier for within port-to-port traffic. On the other hand, the left two plots of Web-related user-sessions shows that high levels of higher-layer activities can either be implicitly accounted for by a semi-random cascade construction with a variable generator (middle) or cannot be adequately captured by such a process. Thus, while the multiplicative structure of network traffic can be clearly identified and isolated at the level of individual port-to-port flows, this becomes less clear when

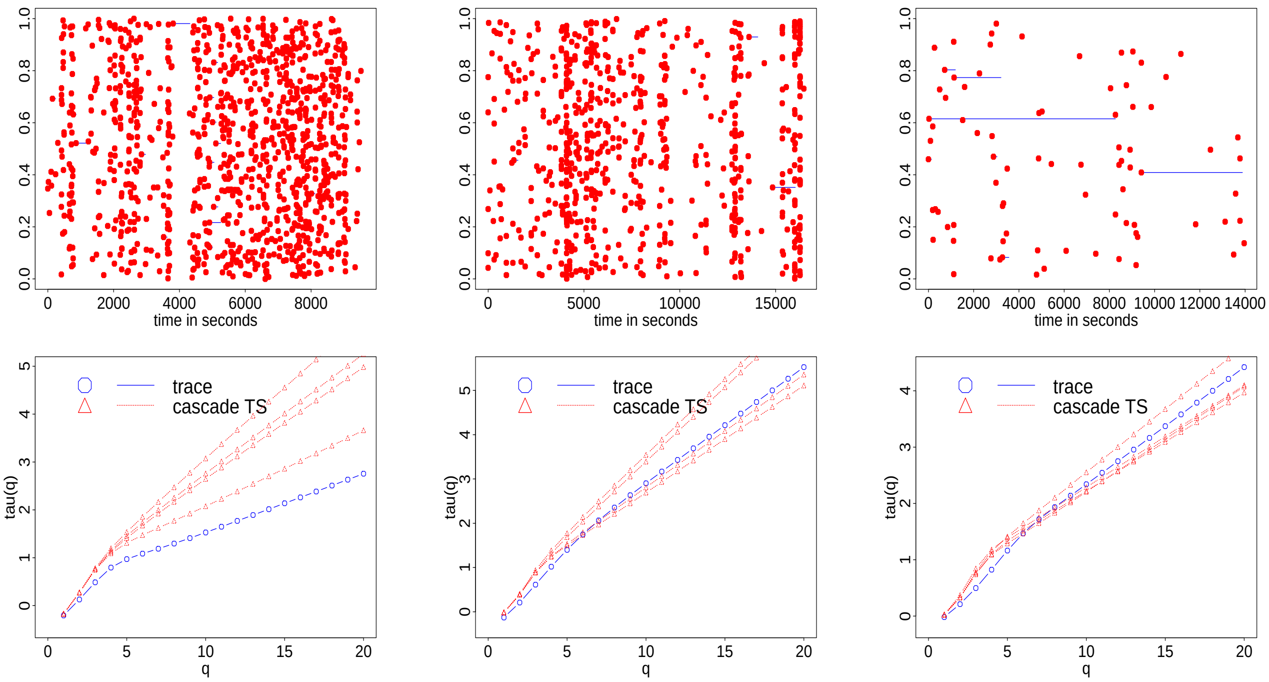

Figure 9: Top: Textured plots of arrivals of port-to-port flows, with port-to-port flow duration, for three different user sessions. Bottom: Multifractal analysis (structure function) of the three sessions with variable generators W_1 (left), W_2 (middle) and W_3 (right).

considering network traffic aggregated to the level of individual user sessions. The difficulties can be attributed to the observation that higher-layer aspects of network traffic can either be additive, or multiplicative, or a combination of both, depending on, for example, the mix of applications or the user behavior in a given session. Note however that by aggregating over all user sessions, the aggregate packet-level traffic fully reflects the combined additive and multiplicative aspects observed at the user session level (see Sections 2 and 3).

4.4 Impact and relevance on workload modeling

Our analysis shows that the clearest distinction between the additive aspect of measured network traffic (over large time scales) and its multiplicative property (over small time scales) can be seen at the level of port-to-port flows (or at slightly higher aggregation levels). Given the findings reported in Section 4.3, making such a distinction at the level of individual user sessions seems to be less obvious. However, to keep things simple, we will ignore in the following this difficulty related to user sessions when discussing an Internet workload modeling approach that assumes that user sessions (i) arrive in accordance to a Poisson process, (ii) bring with them a workload (e.g., number of packets or bytes, number of port-to-port flows, length of session) that is heavy-tailed with infinite variance, and (iii) distribute the workload over the lifetime of the session according to a multiplicatively generated multifractal with a semi-random cascade generator. Thus, for ease of presentation, we assume here that user sessions have within-session structure that agrees with what we observed for port-to-port flows.[7] This workload model is a generalization of Kurtz's model [16, 30] by allowing the within-session traffic rate process to be generated by a semi-random cascade model.

The attractive feature of this workload model is that it accounts in a parsimonious manner for both the global as well as local scaling characteristics observed in measured WAN traffic. While the global scaling behavior is already part of Kurtz's original model (via the relationship between heavy-tailed sizes or durations of the individual sessions and the asymptotic self-similarity of the aggregate packet stream) and is captured by the Hurst parameter H, the original model does not incorporate local scaling behavior. However, we have seen in Section 4.2 that choosing a variable generator W^* for a "typical" semi-random cascade model for the within-flow traffic rate process is relatively obvious (linearly decreasing logarithm of the standard deviation from small to coarse scales, with a slope of around -0.2) and gives rise to multifractal scaling as captured by the corresponding structure function $\tau(q)$. Note that the particular form of W^* retains the parsimonious nature of this workload model and preliminary results suggest that the aggregate traffic generated by these sessions is at the same time asymptotically self-similar (with Hurst parameter H) and multifractal (as expressed in terms of $\tau(q)$).

The practical relevance for such a workload model is that it allows for a more complete description of network traffic than exists to date in cases where higher-order statistics or multiplicative aspects of the traffic play an important role but cannot be adequately accounted for by traditional strictly second-order descriptions of network traffic. By aiming for a complete description of traffic, a comprehensive analysis of network performance-related problems becomes feasible and desirable. In the past, thorough analytical studies of which aspects of network traffic are important for which aspects of network performance have of-

[7]Future work will consider the more appropriate case, suggested by our findings earlier in this section, where port-to-port flow counts are self-similar and where the individual flows are modeled as in (iii).

ten been prevented due to a lack of models that provide provably complete descriptions of the traffic processes under study. This situation can lead to misconceptions and misunderstandings of the relevance of certain aspects of traffic for certain aspects of performance (e.g., see [27] and [12]). Finally, in terms of practical relevance, we also argue that by incorporating—via multifractals—local scaling characteristics of the traffic into a workload model, it may become in fact feasible to adequately describe traffic in a closed system (like the Internet) with an open model.

5 Conclusions and outlook

By analyzing a number of different packet-level WAN traces from different WAN environments and at different layers within the TCP/IP protocol architecture, we attempt in this paper to provide an answer to the question of why measured WAN traffic appears to be multifractal. In effect, we propose and empirically validate that measured network traffic conforms to an underlying cascade construction and identify aspects of network traffic where its multiplicative properties can be examined in detail. In this sense, we make rigorous and validate empirically the intuitive notion that networks act like certain cascades called semi-random cascades and illustrate that they give rise to intricate features in the temporal dynamics of network traffic that agree with the local and global scaling phenomena observed in measured WAN traffic. One of our main findings is that the cascade paradigm or multiplicative nature of network traffic over small time scales (i.e., where the influence of higher-layer activities is negligible) appears to be robust across different WANs and under changes in the underlying WAN environment and traffic conditions, and hence constitutes a new traffic invariant for WAN traffic that can co-exist with the concept of self-similarity. At the same time, through the implied complex local scaling structure, multiplicatively generated multifractals promise great flexibility in accounting for and, in turn, detecting and identifying network/application/user-specific features. While the paper puts in place a structure that provides for extensive and novel explorations of these areas of interest to the networking community, we have barely begun exploring this yet uncharted territory.

To study the local scaling phenomena of measured network traffic, we introduce and illustrate in this paper appropriate methods and techniques for analyzing and inferring multifractal scaling behavior. Our methods are based on wavelets and their natural ability for scale- and time-localization, and the techniques (as well as the practical implementation of the techniques) rely on the theoretical properties of the discrete wavelet transform in a multiresolution analysis. In particular, moving beyond the traditional applications of wavelets to study scale-dependent global features of traffic, we emphasize here their ability for time-localization which can be interpreted as providing a mathematical microscope for detecting and identifying local irregularities in a trace; e.g., location-dependent scaling features. Future work will focus on relating such features to specific networking conditions. By developing practical tools for distinguishing between monofractal and multifractal scaling, we make it easier for networking researchers and engineers to gain access to a new area of traffic analysis (i.e., investigating local structure) that has been off limits in the past.

A natural next step is to explore how this new and improved understanding of modern data networks and data traffic can be exploited for network engineering and traffic management. In addition to the already mentioned implication of multifractals for workload modeling, we conclude with a brief discussion of some aspects of network engineering where knowing either the global or local (or both) scaling behaviors is essential for tackling specific networking problems.

(1) Generating realistic data network traffic: The physical explanation advocated in this paper for the observed multifractal nature of measured WAN traffic gives rise to a simple recipe for synthetically generating realistic data network traffic. Indeed, we demonstrate in Section 3 how semi-random cascades are able to accurately match a given trace not only with respect to its first- and second-order statistical properties but also in terms of the higher-order statistics (e.g., multifractal scaling).

(2) Inferring fine-time scaling behavior from coarse-time measurements: Typical network operations systems collect link-level traffic statistics every 5–15 minutes. However, to predict, for example, that the traffic levels on a trunk stay within safe operating regions, we must to be able to infer the burstiness behavior of the traffic over small-time scales from the large-time scale operational measurements. For multifractal traffic, we can use the underlying semi-random cascade paradigm that determines (via the cascade generator) how a certain workload, measured over large-time scales, is distributed over smaller time scales. In effect one can use the cascade to extend the coarse-scale time series to finer time scales.

(3) Global and local scaling behavior and round-trip time of packets: The global scaling analysis plots in Figure 1 raise a natural question about the pronounced change in the global scaling behavior from small-time to large-time scales, around time scales on the order of a few hundred milliseconds or seconds. It may be that the location of the "knee" is related to properties of the round-trip time in the network or to some other aspects of the particular network under study. A natural approach to study in detail the precise relationship between round-trip time and local and global scaling behavior is to experiment with the generator of the semi-random cascade construction in a controlled network environment. The ns simulator [20] is an ideal tool for this, and initial ns-based experiments show promising results in support of these partly heuristic, partly empirical-based arguments.

(4) Exploiting a new dimension in network traffic analysis: To date, network traffic analysis has focused almost exclusively on first-order (e.g., mean, variance, marginal distributions) and second-order (e.g., autocorrelations, spectral density) statistical properties of measured data, and existing traffic models fully reflect that attitude. The empirical finding of multifractal scaling properties in measured WAN traffic opens up new opportunities for improving our current understanding of modern networks and the traffic that they carry by providing a new perspective (i.e., local scaling) and a new mathematical tool (i.e., multifractal analysis) to investigate aspects of measured traffic that have so far been off limits to network researchers and engineers. These aspects concern the detailed nature of the local irregularities in traffic caused by networking-related mechanisms operating on small-time scales and we can expect them to be of significant importance when we try to infer network-specific properties and/or user-perceived network performance from active network measurements. Intuitively, the relevant information contained in measurements obtained by sending certain test traffic into the network and recording specific responses (e.g., see Paxson [24]) is often contained in the measurements' local irregularities rather than in their global

statistical properties. As such, multifractal analysis is likely to impact how results from active network measurement experiments will be analyzed in the future and how active network measurements will be used to help manage tomorrow's data networks.

Acknowledgments

We acknowledge the help of many of our colleagues at AT&T Labs, especially of J. Friedmann and A. Greenberg, with the data collection effort. Special thanks go to S. Alexander and S. Gao from AT&T Labs-Research for their help with collecting the ATTLAB traces. The Bellcore trace was collected by D.V. Wilson and the LBL trace by V. Paxson; both traces are available from `http://www.acm.org/sigcomm/ITA/`.

We are very grateful to the anonymous Sigcomm reviewers and to V. Paxson for constructive criticism and invaluable suggestions for improving the presentation of the material. The paper also benefited from helpful discussions with R. Calderbank, A. Erramilli, O. Narayan, T. Kurtz, R. Riedi and T. Sweet. The idea of using cascades to explain multifractality in the networking context was originally suggested by F. Clerot during a discussion at Sigcomm'96 with one of the authors (W. W.).

Finally, we thank P. Abry and D. Veitch for making available their programs to perform the wavelet-based global scaling analysis in Section 2.1. Some of the local scaling analysis techniques presented in this paper use functions that are part of the wavelet package in *S-Plus*, described in detail in [3].

References

[1] P. Abry and D. Veitch. Wavelet analysis of long-range dependent traffic. *IEEE Transactions on Information Theory* **44**, pp. 2–15, 1998.

[2] A. Arneodo. Wavelet analysis of fractals: From the mathematical concept to experimental reality. In M. Y. Hussaini, editor, *Wavelets: Theory and Application*, pp. 349–502. Oxford University Press, New York, 1996.

[3] A. Bruce and H.-Y. Gao. *Applied Wavelet Analysis with S-Plus*. Springer-Verlag, New York, 1996.

[4] K. C. Claffy, H.-W. Braun, and G. C. Polyzos. A parameterizable methodology for internet traffic flow profiling *IEEE Journal on Selected Areas in Communications*, vol. 13, pp. 1481–1494, 1995.

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

[6] I. Daubechies. *Ten lectures on wavelets.* SIAM, Philadelphia, 1992.

[7] A. Erramilli, O. Narayan and W. Willinger. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking* **4**, pp. 209–223, 1996.

[8] 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.

[9] A. Feldmann, A. C. Gilbert. W. Willinger and T. G. Kurtz. The changing nature of network traffic: Scaling phenomena. *Computer Communication Review* **28**, No. 2, April 1998.

[10] A. Feldmann, J. Rexford and R. Caceres. Reducing Overhead in Flow-Switched Networks: An Empirical Study of Web Traffic. *Proc. of IEEE INFOCOM'98*, 1998 (to appear).

[11] A. C. Gilbert, W. Willinger and A. Feldmann. Scaling analysis of random cascades in network traffic. Preprint. 1998

[12] M. Grossglauser and J.-C. Bolot. On the relevance of long-range dependence in network traffic. *Proc. of the ACM/SIGCOMM'96*, Stanford, CA, pp. 15-24, 1996.

[13] 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.

[14] R. Holley and E. C. Waymire. Multifractal dimensions and scaling exponents for strongly bounded random cascades. *Annals of Applied Probability* **2**, pp. 819–845, 1992.

[15] G. Kaiser. *A friendly guide to wavelets*, Birkhauser, Boston, 1994.

[16] 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.

[17] 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.

[18] J. Levy-Vehel and R. Riedi. Fractional Brownian motion and data traffic modeling: The other end of the spectrum. In *Fractals in Engineering*, Springer-Verlag, Berlin, pp. 185–202, 1997.

[19] P. Mannersalo and I. Norros. Multifractal analysis of real ATM traffic: A first look. *COST257TD*, VTT Information Technology, 1997.

[20] S. McCanne and S. Floyd. NS: Network Simulator. `http://www-mash.cs.berkeley.edu/ns/`.

[21] Y. Meyer and R. Coifman. *Wavelets: Calderón-Zygmund and multilinear operators*, Cambridge University Press, New York, pp. 52–54, 1997.

[22] E. W. Montroll and M. F. Shlesinger. On 1/f noise and other distributions with long tails. *Proc. Natl. Acad. Sci. USA* **79**, pp. 3380–3383, 1982.

[23] V. Paxson. End-to-end routing behavior in the Internet. *Proc. of the ACM/SIGCOMM'96*, Stanford, CA, pp. 25–38, 1996.

[24] V. Paxson. End-to-end Internet packet dynamics. *Proc. of the ACM/SIGCOMM'97*, Cannes,France, pp. 139–152, 1997.

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

[26] R. H. Riedi and J. Levy-Vehel. TCP traffic is multifractal: A numerical study. Preprint, 1997.

[27] B. K. Ryu and A. Elwalid. The importance of long-range dependence of VBR video traffic in ATM traffic engineering: Myths and Realities. *Proc. of the ACM/SIGCOMM'96*, Stanford, CA, pp. 3–14, 1996.

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

[29] A. H. Tewfik and M. Kim. Correlation structure of the discrete wavelet coefficients of fractional Brownian motion. *IEEE Trans. on Info. Theory*, vol. 38, **2**, pp. 904–909, 1992.

[30] W. Willinger, V. Paxson, and M. S. Taqqu. Self-similarity and heavy tails: Structural modeling of network traffic. 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, 1998 (to appear).

[31] 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.