1 Toward an Improved Understanding of Network Traffic Dynamics

R.H. RIEDI[1] and W. WILLINGER[2]

[1] Department of Electrical and Computer Engineering
Rice University
Houston, TX 77251-1892

[2] Information Sciences Research Center
AT&T Labs-Research
Florham Park, NJ 07932-09771

Draft, April 9, 1999, to appear in
—"Self-similar Network Traffic and Performance Evaluation", Wiley, June 1999

1.1 INTRODUCTION

Since the statistical analysis of Ethernet LAN traces in [20], there has been significant progress in developing appropriate mathematical and statistical techniques that provide a physical-based, networking-related understanding of the observed fractal-like or self-similar scaling behavior of measured data traffic over time scales ranging from hundreds of milliseconds to seconds and beyond. These techniques explain, describe, and validate the reported large-time scaling phenomenon in aggregate network traffic at the packet level in terms of more elementary properties of the traffic patterns generated by the individual users and/or applications. They have impacted our understanding of actual network traffic, to the point where we now know why aggregate data traffic exhibits fractal scaling behavior over time scales from a few hundreds of milliseconds onwards. In fact, a measure of the success of this new understanding is that the corresponding mathematical arguments are at the same time rigorous and simple, are in full agreement with the networking researchers' intuition and with measured data, and can be explained readily to a non-networking expert. These developments have helped immensely in demystifying fractal-based traffic modeling and have given rise to new insights and physical understanding of the effects of large-time scaling properties in measured network traffic on the design, management and performance of high-speed networks.

However, to provide a complete description of data network traffic, the same kind of understanding is necessary with respect to the dynamic nature of traffic over small time scales, from a few hundreds of milliseconds downwards. Because of the predominant protocols and end-to-end congestion control mechanisms that play a central role in modern-day data networks and determine the flow of packets over those fine time scales and at the different layers in the TCP/IP protocol hierarchy, studying the fine-time scale behavior or local characteristics of data traffic is intimately related to understanding the complex interactions that exist in data networks such as the Internet between the different connections, across the different layers in the protocol hierarchy, over time as well as in space. In this chapter, we first summarize the results that provide a unifying and consistent picture of the large-time scaling behavior of data traffic and discuss the appropriateness of self-similar processes such as fractional Gaussian noise for modeling the fluctuations of the traffic rate process around its mean and for providing a complete description of the traffic on individual links within the network. Then we report on recent progress in studying the small-time scaling behavior in data network traffic and outline a number of challenging open problems that stand in the way of providing an understanding of the local traffic characteristics that is as plausible, intuitive, appealing and relevant as the one that has been found for the global or large-time scaling properties of data traffic.

1.2 THE LARGE-TIME SCALING BEHAVIOR OF NETWORK TRAFFIC

In this section, we demonstrate why the empirically observed large-time scaling behavior or (asymptotic) self-similarity of aggregate network traffic is an additive property, with the additional requirement that the individual component processes that generate the total traffic exhibit certain high-variability or heavy-tailed characteristics.

1.2.1 Additive structure and Gaussianity

When viewed over large enough time scales, the number of packets or bytes per time unit collected off a link in a network originate from all those connections that were active during the measurement period, utilized this link, and actively generated traffic during this time. In other words, if for "time scales" or "levels of resolution" $m >> 1$, $X^{(m)} = (X^{(m)}(k) : k \geq 0)$ denotes the overall traffic rate process, i.e., the total number of packets or bytes per time unit (measured at time scale m) generated by all connections, then we can write

$$X^{(m)}(k) = \sum X_i^{(m)}(k), k \geq 0 \tag{1.1}$$

where the sum is over all connections i that are active at time k and where $X_i^{(m)} = (X_i^{(m)}(k) : k \geq 0)$ represents the total number of packets or bytes per time unit (again measured at time scale m) generated by connection i.[1] Thus, Equation (1.1) captures the *additive* nature of aggregate network traffic by expressing the overall traffic rate process $X^{(m)}$ as a superposition of the traffic rate processes $X_i^{(m)}$ of the individual connections.

Assuming for simplicity that the individual traffic rate processes $X_i^{(m)}$ are independent from one another and identically distributed, then under weak regularity conditions on the marginal distribution of the $X_i^{(m)}$'s (including for example the existence of second moments), Equation (1.1) guarantees that the overall traffic rate process (or its deviations from its mean) exhibits Gaussian marginals, as soon as the traffic is generated by a sufficiently large number of individual connections.

1.2.2 Self-similarity through heavy-tailed connections

Focusing on the temporal dynamics of the individual traffic rate processes $X_i^{(m)}$, suppose for simplicity that connection i sends packets or bytes at a constant rate (say, rate 1) for some time (the "active" or "on" period) and does not send any packets or bytes during the "idle" or "off" period; we will return

[1] Note that the processes $X^{(m)}$ and $X_i^{(m)}$ are defined by averaging X and X_i over non-overlapping blocks of size m.

to the challenging problem of allowing for more realistic "within-connection" packet dynamics in Section 1.3 below. For example, in a LAN environment, a connection corresponds to an individual host-to-host or source-destination pair and the corresponding traffic patterns have been shown in [38] to conform to an alternating renewal process where the successive pairs of on- and off-periods define the inter-renewal intervals. On the other hand, in the context of wide-area networks or WANs such as the Internet, we associate individual connections with "sessions," where a session starts at some random point in time, generates packets or bytes at a constant rate (say, rate 1) during the life time of the connection and then stops transmitting packets or bytes. Here a session can be an FTP application, a TELNET connection, a Web session, sending email, reading Network News, etc., or any imaginable combination thereof. In fact, over 1/2- to 1-hour periods, session arrivals on Internet links have been shown to be consistent with a homogeneous Poisson process; e.g., see [25] for FTP and TELNET sessions, and [12] for Web sessions. Note that in the present setting, only global connection characteristics (e.g., session arrivals, life times of sessions, durations of the on/off periods) play a role, while the details of how the packets arrive within a connection or within an on-period have been conveniently modeled away by assuming that the packets within a connection are generated at a constant rate.

To describe the stochastic nature of the overall traffic rate process $X^{(m)}$, the only stochastic elements that have not yet been specified are the distributions of the lengths of the on/off-periods (in case of the LAN example) or the distribution of the session durations (for the WAN case) associated with the individual traffic rate processes $X_i^{(m)}$. Based on measured on/off-periods of individual host-to-host pairs in a LAN environment (e.g., see [38]) and measured session durations from different WAN sites (e.g., [25, 37, 12]), we choose these distributions to be heavy-tailed with infinite variance. Here, a positive random variable U (or the corresponding distribution function F) is called *heavy-tailed with tail index* $\alpha > 0$ if it satisfies

$$P[U > y] = 1 - F(y) \approx cy^{-\alpha}, \quad \text{as } y \to \infty, \tag{1.2}$$

where $c > 0$ is a finite constant that does not depend on y. Such distributions are also called *hyperbolic* or *power-law distributions*, and include, among others, the well-known class of *Pareto distributions*. The case $1 < \alpha < 2$ is of special interest and concerns heavy-tailed distributions with finite mean but *infinite variance*. Intuitively, infinite variance distributions allow random variables to take values that vary over a wide range scales and can be exceptionally large with non-negligible probabilities. Hence, heavy-tailed distributions with infinite variance allow for compact descriptions of the empirically observed high-variability phenomena that dominate traffic-related measurements at all layers in the networking hierarchy; e.g., see [12].

Mathematically, the heavy-tailed property of, for example, the durations during which individual connections actively generate packets implies that the temporal correlations of the stationary versions of an individual traffic rate processes $X_i^{(m)}$ and, because of the additivity property (1.1), of the overall traffic rate process $X^{(m)}$, decay hyperbolically slowly; that is, they exhibit long-range dependence. More precisely, if $r^{(m)} = (r^{(m)}(k) : k \geq 0)$ denotes the autocorrelation function of the stationary version of the overall traffic rate process $X^{(m)}$, then property (1.2) can be shown to imply *long-range dependence* (e.g., see [4] and [38]; for similar results obtained in the context of a fluid queueing system under heavy traffic, see Brichet's chapter in this book; that is, for all $m \geq 1$, $r^{(m)}$ satisfies

$$r^{(m)}(k) \approx ck^{2H-2}, \quad \text{as } k \to \infty, \quad 0.5 < H < 1, \tag{1.3}$$

where the parameter H is called the *Hurst parameter* and measures the degree of long-range dependence in $X^{(m)}$; in terms of the the tail index $1 < \alpha < 2$ that measures the degree of "heavy-tailedness" in (1.2), H is given by $H = (3 - \alpha)/2$. Intuitively, long-range dependence results in periods of sustained greater-than-average or lower-than-average traffic rates, irrespective of the time scale over which the rate is measured. In fact, for a zero-mean covariance-stationary process, Equation (1.3) implies (and is implied by) *asymptotic (second-order) self-similarity*; that is, after appropriate rescaling, the overall traffic rate processes $X^{(m)}$ have identical second-order statistical characteristics and "look similar" for all sufficiently large time scales m. In other words, Equation (1.3) holds if and only if for all sufficiently large time scales m_1 and m_2, we have

$$m_1^{1-H} X^{(m_1)} \approx m_2^{1-H} X^{(m_2)}, \tag{1.4}$$

where the quality is in the sense of second-order statistical properties and where $1/2 < H < 1$ denotes the self-similarity parameter and agrees with the Hurst parameter in Equation (1.3).

The ability to explain the empirically observed self-similar nature of aggregate data traffic in terms of the statistical properties of the individual connections that make up the overall traffic rate process shows that (asymptotically) self-similar behavior (i) is an intrinsically additive property (i.e., aggregate over many connections), (ii) is mainly caused by user/session/connection characteristics (i.e., Poisson arrivals of sessions, heavy-tailed distributions with infinite variance for the session sizes/durations), and (iii) has little to do with the network (i.e., the predominant protocols and end-to-end congestion control mechanisms that determine the actual flow of packets in modern data networks). In fact, for the self-similarity property of data traffic over large time scales to hold, all that is needed is that the number of packets or bytes per connection is heavy-tailed with infinite variance, and the precise nature of how the individual packets within a session or connection are sent over the network is largely irrelevant.

Note that this understanding of data traffic started with an extensive analysis of measured aggregate traffic traces, followed by the statistically well-grounded conclusion of their self-similar or fractal characteristics, and triggered the curiosity of networking researchers who wanted to know "Why self-similar or fractal?" In turn, this question for a physical explanation of the large-time scaling behavior of measured data traffic resulted in findings about data traffic at the connection level that are, at the same time, mathematically rigorous, agree with the networking researchers' experience, are consistent with data, and are intuitive and simple to explain in the networking context. In this sense, the progression of results proceeded the opposite of how traffic modeling has been traditionally done in this area; that is, by first analyzing in great detail the dynamics of packet flows within individual connections and then appealing to some mathematical limiting result that allowed for a simple approximation of the complex and generally over-parameterized aggregate traffic stream. In contrast, the self-similarity work has demonstrated that novel insights into and new and unprecedented understanding of the nature of actual data traffic can be gained by a careful statistical analysis of measured traffic at the aggregate level and by explaining aggregate traffic characteristics in terms of more elementary properties that are exhibited by measured data traffic at the connection-level.

1.2.3 Self-similar Gaussian processes as workload models

Notice that in the Gaussian setting discussed in Section 1.2.1, the self-similarity property (1.4) implies that for $1/2 < H < 1$ and for all sufficiently large time scales m, the traffic rate process $X^{(m)}$ (or more precisely, the deviations from its mean) satisfy

$$m^{1-H} X^{(m)} \approx X, \tag{1.5}$$

where in this case, the equality is understood in the sense of finite-dimensional distributions, and where $X = (X_k : k \geq 1)$ denotes *fractional Gaussian noise* (FGN), the only stationary (zero-mean) Gaussian process that is *(exactly) self-similar* in the sense that Equation (1.5) holds for all $m \geq 1$. Equivalently, FGN is uniquely characterized as the stationary (zero-mean) Gaussian process with autocorrelation function $r(k) = 1/2\ [(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}],\ k \geq 1, 1/2 < H < 1$.

For the purpose of modeling the dynamics of actual data traffic over a link within a network, FGN has the big advantage of providing a complete description of the resulting traffic rate process; that is, specifying its mean, variance, and Hurst parameter H suffices to completely characterize the traffic. Given this advantage over other – typically incomplete – descriptions of network traffic dynamics, it is important to know under what conditions FGN is an adequate and accurate process for modeling the deviations around the mean of actual data traffic. To this end, Erramilli et al. [8] note that the FGN model can be expected to be an appropriate model for data traffic provided (i) the traffic is aggregated over a large number of independent and not too wildly fluctuating connections (i.e., ensuring Gaussianity of expression (1.1)), (ii) the effects of flow control on any one connection is negligible (i.e., requiring, in fact, that we consider the traffic only over sufficiently large time scales where (1.4 holds), and (iii) the time scales of interest for the performance problem at hand coincide with the scaling region (i.e., where (1.5) holds). In practice, these conditions are often satisfied in the backbone (i.e., high levels of aggregation) and for time scales that are larger than the typical round-trip time of a packet in the network.

1.2.4 Toward self-similar non-Gaussian workload models?

One of the conditions mentioned above that justify the use of FGN as an adequate and accurate description of actual data traffic traversing individual links in a network states that the traffic over a specific link is made up from a large number of (more or less) independent connections, where each connection's own traffic rate cannot fluctuate too wildly; that is, $X_i^{(m)}$ is chosen from a distribution with finite variance. While this condition is generally applicable in many legacy LAN and WAN environments and can often be validated against measured traffic, due to changes in networking technologies, applications and user behavior, it can no longer be taken for granted in today's networks. For example, advanced networking technologies such as 100 Mbps Ethernets or Gigabit Ethernets can be expected – despite the presence of TCP, for example – to allow the traffic rates of individual connections to vary over many orders of magnitude, from Kbps to Mbps and beyond, depending on the networking conditions. Thus, for understanding modern-day network traffic, processes that combine heavy tails both in time and space (i.e., the distributions of the durations as well as of the rates at which individual connections emit packets are heavy-tailed with infinite variance) may become relevant in practice and may see genuine applications in the networking area in the near future.

To illustrate, let $X_i^{(m)}$ denote an on/off-type connection described earlier, where in addition to the duration of the on/off-periods, the rate at which the connection emits packets during the on-period is also heavy-tailed with infinite variance (with tail index β, say). Focusing on this modification of the renewal reward model investigated by Mandelbrot [22] and Taqqu and Levy [34], Levy and Taqqu [21] recently showed that when studying the overall traffic rate process $X^{(m)}$ defined in (1.1); i.e., aggregating many such independent connections, one can obtain a dependent, stationary process that has a stable marginal distribution with infinite variance and which is self-similar as in (1.5) with self-similarity parameter H given by

$$H = \frac{\beta - \alpha + 1}{\beta}. \tag{1.6}$$

Here β denotes the index characterizing the heaviness of the tail of the traffic rate of the individual connections, and α denotes the tail index associated with the distributions of the durations of on- and off-periods which we assume for simplicity to be identical. Observe that in the finite variance case ($\beta = 2$), Relation (1.6) reduces to the familiar $H = (3-\alpha)/2 \in (1/2, 1)$ which appears in connection with fractional Gaussian noise considered earlier. However, in contrast to FGN, the superposition process obtained under the assumption of heavy tails with infinite variance on the durations *and* rates is not Gaussian but has heavy-tailed marginals instead, implying that there is a much higher probability than in the Gaussian case that the overall traffic rate can differ greatly from the average value and that it can take extreme values (a phenomenon also known as *intermittency*). Being non-Gaussian, one of the obstacles at this stage for using these kinds of stable superposition processes in the context of modeling data traffic is that their statistical parameters α (which specifies the marginals) and H (1.5) do not define them completely; there exist a number of different dependent, stationary increment processes with stable marginals with the same α and same self-similarity parameter H – see for example [33]. This is in stark contrast to FGN where knowing the second-order statistical characteristics (i.e., variance and Hurst parameter H) uniquely defines the process, due to Gaussianity.

1.3 THE SMALL-TIME SCALING BEHAVIOR OF NETWORK TRAFFIC

The analysis of measured network traffic and resulting understanding of some of its underlying structure outlined in Section 1.2 above have led to the realization that while wide-area traffic is consistent with asymptotic self-similarity or large-time scaling behavior, its small-time scaling features are very different from those observed over large time scales. Thus, to provide an adequate and more complete description of actual network traffic, it is necessary to deal with these small-time scaling features and to ultimately understand their cause and effects. To this end, we summarize in this section our current understanding of this very recent development in network traffic analysis and modeling by introducing concepts that are novel to the networking area, e.g., multifractals, conservative cascades, and multiplicative structure, and illustrate their relevance to networking.

1.3.1 Multifractals

From a networking perspective, it comes as no surprise that protocol-specific mechanisms and end-to-end congestion control algorithms operating on small time scales and at the different layers in the hierarchical structure of modern data networks give rise to structural properties that are drastically different from the large-time scaling behavior which has been shown earlier to be mainly due to global user and/or session characteristics. Since these networking mechanisms determine largely the actual flow of packets across the network, they are likely to cause the traffic to exhibit pronounced local variations and irregularities which, per se, cannot be expected to have any obvious connection to the self-similar behavior of the traffic over large time scales.

To quantify these local variations in measured traffic at a particular point in time t_0, let $Y = (Y(t) : \ 0 \leq t \leq 1)$ denote the process representing the total number of packets or bytes sent over a link up to time t, and for some $n > 0$, consider the traffic rate process $Y((k_n + 1)2^{-n}) - Y(k_n 2^{-n}); k_n = 0, 1, \ldots, 2^n - 1$; that is, the total number of packets or bytes seen on the link during non-overlapping intervals of the form $[k_n 2^{-n}, (k_n + 1)2^{-n})$. We say that the traffic has a local scaling exponent $\alpha(t_0)$ at time t_0 if the traffic rate process behaves like $(2^{-n})^{\alpha(t_0)}$, as $k_n 2^{-n} \to t_0$ $(n \to \infty)$. Note that $\alpha(t_0) > 1$ corresponds to instants with low intensity levels or small local variations (Y has derivative zero at t_0), while $\alpha(t_0) < 1$ is found in regions with high levels of burstiness or local irregularities. Informally, we call traffic with the same scaling exponent at all instants t_0 *monofractal* (this includes exactly self-similar traffic, for which $\alpha(t_0) = H$, for all t_0), while traffic with non-constant scaling exponent $\alpha(t_0)$ is called *multifractal*.

More formally, the degree of local irregularity of a signal Y or its singularity structure at a given point in time t_0 can be characterized to a first approximation by comparison with an algebraic function, i.e. $\alpha(t_0)$ is the best (i.e., largest) α such that $|Y(t') - Y(t_0)| \leq C|t' - t_0|^{\alpha}$, for all t' sufficiently close to t_0. Since our process Y has positive increments, this *singularity exponent* can be approximated through the somewhat simpler quantity

$$\alpha(t) = \lim_{n \to \infty} \alpha_n(t), \tag{1.7}$$

where – assuming the limit exists – for $t \in [k_n 2^{-n}, (k_n + 1)2^{-n})$,

$$\alpha_n(t) := \alpha^n_{k_n} := -\frac{1}{n} \log_2 \left| Y((k_n + 1)2^{-n}) - Y(k_n 2^{-n}) \right|. \tag{1.8}$$

The aim of *multifractal analysis* (MFA) is to provide information about these singularity exponents in a given signal and to come up with a compact description of the overall singularity structure of signals in geometrical or in statistical terms. Before describing in more detail some of the commonly used MFA methods, we note that since wavelet decompositions contain information about the degree of local irregularity of a signal, it should come as no surprise that the singularity exponent $\alpha(t)$ is related to the decay of wavelet coefficients $w_{j,k} = \int Y(s)\psi_{j,k}(s)\,\mathrm{d}s$ around the point t, where ψ is a band-pass wavelet function and where $\psi_{j,k}(s) := 2^{-j/2}\ \psi\big(2^{-j}s - k\big)$ (e.g., in the case of the well-known *Haar wavelet*, *psi*(s) equals 1 for $0 \leq s \leq 1$, -1 for $1 \leq s \leq 2$, and 0 for all other s; for a general overview of wavelets, we refer to [5]). Indeed, assuming only that $\int \psi(s)\,\mathrm{d}s = 0$ one can show as in [18] that

$$2^{n/2} w_{-n,k_n} \leq C \cdot 2^{-n\alpha(t)}, \qquad \text{as } k_n 2^{-n} \to t. \tag{1.9}$$

Moreover, it is known that under some regularity conditions (for a precise statement see [18] or [5, Theorem 9.2]), relation (1.9) characterizes the degree of local irregularity of the signal at the point t. This suggests to define $\tilde{\alpha}(t)$ as in (1.8) but with $\alpha_n(t)$ replaced by $\tilde{\alpha}_n(t)$, where

$$\tilde{\alpha}_n(t) := \tilde{\alpha}^n_{k_n} := \frac{1}{-n\log 2}\log\left(2^{n/2}|w_{-n,k_n}|\right). \tag{1.10}$$

In general, this may give a different but nevertheless useful description of the singularity structure of Y, particularly for non-monotonous processes (for an example, see [13]). Using wavelets may also have numerical advantages. The remainder of this section remains true if $\alpha(t)$ is replaced by $\tilde{\alpha}(t)$ and (1.8) by (1.10), i.e. increments by normalized wavelet coefficients.

Conceptionally, the geometrical formulation of MFA in the time-domain is the most obvious one. Its objective is to quantify what values of the limiting scaling exponent $\alpha(t)$ appear in a signal and how often one will encounter the different values. In other words, the focus here is on the "size" of the sets of the form

$$K_\alpha = \{t : \alpha(t) = \alpha\}. \tag{1.11}$$

To illustrate, since for FGN there exists only one scaling exponent (i.e., $\alpha(t) = H$), the set K_α is either the whole line (if $\alpha = H$) or empty, and FGN is therefore said to be "mono-fractal." Similarly, for the concatenation of several FGNs with Hurst parameters H^i in the interval $I^i = [i, i+1]$, we have $K_{H^i} = I^i$. In general, however, the sets K_α are highly interwoven and each of them lies dense on the line. Consequently, the right notion of "size" is that of the *fractal Hausdorff dimension* $\dim(K_\alpha)$ which is, unfortunately, impossible to estimate in practice and severely limits the usefulness of this geometrical approach to MFA. Therefore, we will focus below on different statistical descriptions of the multifractal structure of a given signal.

One such description involves the notion of the *coarse Hölder exponents* (1.8). To illustrate, *fix a path of* Y and consider a histogram of the α^n_k ($k = 0,\dots 2^n-1$) taken at some finite level n. It will show a non-trivial distribution of values, but is bound to concentrate more and more around the expected value as a result of the LLN: values other than the expected value must occur less and less often. To quantify the frequency with which values other than the mean value occur, we make extensive use of the theory of large deviations. Generalizing the Chernoff-Cramer bound, the large deviation principle (LDP) states that probabilities of rare events (e.g., the occurrence of values that deviate from the mean) decay exponentially fast. To make this more precise consider a sequence of i.i.d. random variables W, W_1, $W_2,\dots$ and set $V_n := W_1 + \dots + W_n$. Using Chebyshev's inequality and the independence, we find for any $q > 0$,

$$P\left[(1/n)V_n \geq a\right] = P\left[2^{qV_n} \geq 2^{nqa}\right] \leq \frac{\mathbb{E}2^{qV_n}}{2^{nqa}} = \left(\mathbb{E}[2^{qW}]2^{-qa}\right)^n. \tag{1.12}$$

Since $q > 0$ is arbitrary, we can replace the right hand side in (1.12) by its infimum over $q > 0$. A symmetry argument shows that $P[b \geq (1/n)V_n] \leq (\mathbb{E}[2^{qW}]2^{-qb})^n$, for all $q < 0$. Combining all this yields the following two upper bounds

$$\frac{1}{n}\log_2 P\left[b \geq (1/n)V_n \geq a\right] \leq \left\{ \begin{array}{l} \inf_{q>0}\{\log_2 \mathbb{E}[2^{qW}] - qa\} \\ \inf_{q<0}\{\log_2 \mathbb{E}[2^{qW}] - qb\} \end{array} \right. . \tag{1.13}$$

For a discussion of this simple result, let $L(q) = \mathbb{E}[2^{q(W-a)}]$. Since $\log(\cdot)$ is a monotone function, finding the infimum of L is the same as finding the infimum of $\log(L)$. We note first that $L''(q) > 0$, for all $q \in \mathbb{R}$, whence L is a strictly convex function and must have a unique infimum for $q \in \mathbb{R}$. From $L(0) = 1$ we conclude that this infimum must be less than or equal to 1. Focusing now on $q > 0$, we infer from $L'(0) = \log(2)(\mathbb{E}[W] - a)$ that $\inf_{q>0} L(q)$ is assumed in $q = 0$ and equals 1 if and only if $\mathbb{E}[W] \geq a$. On the other hand, $\inf_{q>0} L(q) < 1$ if $\mathbb{E}[W] < a$. An analogous result holds for the second bound. In summary, if $b > \mathbb{E}W > a$ then the bounds on the right hand side in (1.13) are both zero and reflect, thus, the LLN which says that $(1/n)V_n \to \mathbb{E}[W]$ almost surely. On the other hand, if $\mathbb{E}[W]$ is not contained in $[a, b]$ and when $P\left[b \geq (1/n)V_n \geq a\right]$ is the probability of $(1/n)V_n$ deviating far from its expected value, then exactly one of the bounds will be negative proving (at least) exponential decay of this probability. LDP-theorems extend this result to a more general class of random sequences V_n and establish conditions under which the bound in (1.13) is attained in the limit $n \to \infty$ [7, 6].

To apply the LDP approach to our situation, we fix a realization of Y and consider the location t, encoded by k_n via $t \in [k_n 2^{-n}, (k_n + 1)2^{-n})$, as the only randomness relevant for the LDP. Since k_n can take only 2^n different values which we will assume to be all equally likely, the relevant probability measure for t is the counting measure P_t. The sequence of interest for our purpose is

$$V_n := -\log_2 |Y((k_n + 1)2^{-n}) - Y(k_n 2^{-n})| = n\alpha^n_{k_n}.$$

Trying to obtain more precise information about the singularity behavior and aiming at simplifying (1.13), we let not only n tend to ∞ but also let $[a, b]$ shrink down to a single point $\alpha = (a + b)/2$ which unifies the two bounds in the limit. All this suggests that the following limiting "rate function" f will exist under mild conditions [27, Theorem 7]:

$$f(\alpha) := \lim_{\varepsilon \to 0} \lim_{n \to \infty} \frac{1}{n} \log_2 f_n(\alpha, \varepsilon), \tag{1.14}$$

with

$$f_n(\alpha, \varepsilon) := 2^n P_t\Big[\alpha + \varepsilon \geq \alpha_n(t) \geq \alpha - \varepsilon\Big] = \#\{\alpha_n(t) \in (\alpha - \varepsilon, \alpha + \varepsilon)\} \tag{1.15}$$

The counting in (1.15) relates to the notion of dimension: if $f(\alpha) = 1$ then all or at least a considerable part of the α_k^n's are approximatively equal to α, i.e., $f_n(\alpha, \varepsilon) \simeq 2^n$. Such is the case for FGN with $\alpha = H$; but we also have $f(\alpha) = 1$ if only a certain constant fraction of α_n's equals α, as is the case with the concatenation of FGNs described earlier [36]. Only if certain values of α_n are considerably more spurious than others will we observe $f(\alpha) < 1$. In fact, it can be shown [29, 28] that the rate function $f(\alpha)$ relates to $\dim(K_\alpha)$ and that we have

$$\dim(K_\alpha) \leq f(\alpha). \tag{1.16}$$

It is in this sense that f provides information on the occurrence of the various "fractal" exponents α and has been termed *multifractal spectrum.* Also, note that the rate function f is a random element because it is defined for every path of Y.

Although f can, in principle, be computed in practice, it is a very delicate and highly sensitive object, mainly because of its definition in terms of a double limit (see (1.14)). Fortunately, the LDP-result suggests to use the RHS of (1.13), with $\mathbb{E}[2^{qW}]$ replaced by $(\mathbb{E}[2^{qV_n}])^{1/n}$ as in (1.12), as an alternative method for estimating f that avoids double-limit operations and is generally more robust because it involves averages. In fact, consider the *partition function* $\tau(q)$ defined by

$$\tau(q) := \lim_{n\to\infty} \frac{-1}{n} \log_2 \left(2^n \mathbb{E}_t[2^{qV_n}]\right) = \lim_{n\to\infty} \frac{-1}{n} \log_2 S_n(q), \tag{1.17}$$

where the so-called *structure function* $S_n(q)$ is given by

$$S_n(q) := \sum_{k=0}^{2^n-1} |Y((k+1)2^{-n}) - Y(k2^{-n})|^q = \sum_{k=0}^{2^n-1} 2^{-qn\alpha_k^n}. \tag{1.18}$$

According to the theory of LDP we will have equality in (1.13) under mild conditions, at least in the limit as $n \to \infty$ and $b \to a$. Appealing to such results it is possible to establish conditions under which $f(\alpha) = \inf(q\alpha - \tau(q))$. In fact, collecting the terms k in $S_n(q)$ with $\alpha_k^n(t)$ approximately equal to some given value, say α, for varying α and noting that we have about $2^{nf(\alpha)}$ such terms yields

$$S_n(q) = \sum_\alpha \sum_{\alpha_n \simeq \alpha} 2^{-nq\alpha} \simeq \sum_\alpha 2^{-n(q\alpha - f(\alpha))} \simeq 2^{-n \inf_\alpha (q\alpha - f(\alpha))};$$

that is,

$$\tau(q) = f^*(\alpha) := \inf_\alpha (q\alpha - f(\alpha)), \tag{1.19}$$

where * denotes the Legendre transform of a function (for a mathematically rigorous argument, see [27, 28]).

While the partition function $\tau(q)$ is clearly easier to estimate than f it has to be noted, though, that f may contain more information than τ. In fact, the Legendre back-transform yields only

$$f(\alpha) \leq f^{**}(\alpha) = \tau^*(\alpha) = \inf_q (q\alpha - \tau(q)) \tag{1.20}$$

where f^{**} is the concave hull of f (compare (1.13)[2]). The questions are when and for which α the equality $f^{**}(\alpha) = f(\alpha)$ holds. A simple application of the LDP theorem of Gärtner-Ellis [7] provides an answer to these questions – under the assumption that $\tau(q)$ is differentiable everywhere (see [27]). In this particular case, we obtain the appealing formula

$$f(\alpha) = \tau^*(\alpha) = q\alpha - \tau(q) \qquad \text{at } \alpha = \tau'(q). \tag{1.21}$$

Since $\tau(q)$ is the Legendre transform of f, it must always be *concave.* This follows also from the fact that $S_n(q)$ is a log-convex function of q. Consequently, $\tau(q)$ is differentiable in almost all q a priori. For FGN, however, we obtain the degenerate case of a concave function: with probability one, we have

$$\tau(q) = qH - 1 \qquad q > -1. \tag{1.22}$$

This is consistent with the fact that $\alpha(t) = H$ for all t, i.e. the set K_H has dimension 1. The formula (1.22) can be guessed directly from ergodicity and self-similarity:

$$S_n(q) \simeq \sum_{k=0}^{2^n-1} \mathbb{E}|Y((k+1)2^{-n}) - Y(k2^{-n})|^q \approx 2^{n-nqH}\mathbb{E}|Y(1)|^q.$$

For the example considered earlier where we concatenated a number of FGNs, we find $\tau(q) = \min_k(qH_k - 1)$ which is again consistent with $\alpha(t)$ taking the values H_k on sets of dimension 1 (compare (1.33), see also [36] for more details). This example shows also how non-concavity in $\tau(q)$ can result in loss of information: $\tau(q)$ and its Legendre transform reflect only the minimal and the maximal of the H_k. In contrast, truly concave behavior of $\tau(q)$ indicates that there is a whole interval of α-values present in the signal and not just a few (hence the term *multifractal*).

1.3.2 Multiplicatively generated multifractals or cascades

A construction that fragments a given set into smaller and smaller pieces according to some geometric rule and, at the same time, divides the measure of these pieces according to some other (deterministic or random) rule is called a *multiplicative process* or *cascade* (e.g., see [9]). The limiting object generated

[2]The factors 2^n appearing in f and $\tau(q)$ are for convenience. The sign of $\tau(q)$ is chosen as to render (1.20) and (1.19) symmetrical.

by such a multiplicative process defines, in general, a singular measure or multifractal and describes the highly irregular way the mass of the initial set gets redistributed during this simple fragmentation procedure. The *generator* of the cascade specifies the mass fragmentation rule, and we consider in the following the class of *conservative cascades*, introduced by Mandelbrot [23] characterized by a generator that preserves the total mass of the initial set at every stage of the construction (i.e., mass conservation). To illustrate, we will construct a binomial conservative cascade or measure μ on the interval $I := [0, 1]$. More precisely, we will construct its distribution function $Y(t) = \mu([0, t])$ and since the underlying generator will be random, Y will define a stochastic process. By construction it will have positive increments and $Y(0) = 0$ almost surely.

This iterative construction starts with a uniform distribution on the unit interval of total mass M^0 and then 'redistributes' this mass by splitting it among the two subintervals of half size in the ratio M_0^1 to M_1^1 where $M_0^1 + M_1^1 = 1$. Proceeding iteratively one obtains after n steps a distribution which is uniform on intervals $I_{k_n}^n := [k_n 2^{-n}, (k_n + 1)2^{-n}]$. The mass lying in $I_{k_n}^n$ is redistributed among its two dyadic subintervals $I_{2k_n}^{n+1}$ and $I_{2k_n+1}^{n+1}$ in the proportions $M_{2k_n}^{n+1}$ and $M_{2k_n+1}^{n+1}$ where $M_{2k_n}^{n+1} + M_{2k_n+1}^{n+1} = 1$ almost surely.

To summarize, for any n let us choose a sequence $k_1, k_2, \ldots, k_n$ such that the interval $I_{k_l}^l$ lies in $I_{k_i}^i$ whenever $i < l$. In other words, the k_i are the n first binary digits of any point $t \in I_{k_n}^n$. We call this a *nested sequence*, and it is uniquely defined by the value of k_n. Then we have

$$Y((k_n + 1)2^{-n}) - Y((k_n)2^{-n}) = \mu(I_{k_n}^n) = M_{k_n}^n \cdot M_{k_{n-1}}^{n-1} \cdots M_{k_1}^1 \cdot M_0^0. \quad (1.23)$$

The various M_l^i, which collectively define the generator of the conservative cascade, may have distributions which depend on i and l and which are arbitrary, as long as they are positive and provided that for all i and all m,

$$M_{2m}^i + M_{2m+1}^i = 1, \quad (1.24)$$

almost surely. Note that this mass conservation condition introduces a strong dependence between the two "children" of any parent node. Furthermore, we will require that for all n and k_n $(n = 1, 2, \ldots)$, all the multipliers appearing in (1.23) are mutually independent. We will call this property *nested independence*. As long as these two requirements on dependency are satisfied one is completely free in how to introduce further correlation structure.

It is obvious from this iterative construction and from relation (1.23) that a multiplicatively generated "multifractal process" has approximately *lognormal* marginals. Indeed, as a sum of independent random variables, the logarithms of the increments of Y are approximately Gaussian, provided that the random variables $\log M_l^i$ have finite second moments.

Note that as we move from stage n to $n + 1$ in our construction of a conservative cascade, the conservation property (1.24) insures that the values of Y at dyadic points of order less than n are not changed. As we let n tend to infinity, we see from (1.23) that the increments of Y between dyadic points tend to zero, whence Y is continuous (μ has no atoms) and well-defined. Moreover, Y has increments of all lags but no (meaningful) derivative in the usual sense. As we will see, $\alpha(t)$ equals the expected value $\overline{\alpha}$ almost everywhere with $\overline{\alpha} > 1$, whence in these points, the product in (1.23) behaves like $2^{-n\overline{\alpha}}$ and the conventional derivative Y' is zero. Thus, the essential growth of Y happens "in" the points where Y' does not exist. In other words, the true derivative of Y is a distribution or singular measure, i.e. μ.

To study the singularity structure of Y using $\alpha(t)$, we calculate the partition function $\tau(q)$ of the binomial conservative measure "in expectation". To this end, we assume that the M_k^n ($k = 0, \ldots, 2^n - 1$) are identically distributed with $M^{(n)}$. Note that $M^{(n)}$ is necessarily symmetrically distributed around $1/2$ due to (1.24). Then, (1.23) is equally distributed as $M^{(n)} \cdot \ldots \cdot M^{(1)} \cdot M^0$ for each of the 2^n nested sequences $k_1, \ldots k_n$ of length n. Using the "nested" independence we find

$$\mathbb{E}[S_n(q)] = 2^n \cdot \mathbb{E}(M^{(n)})^q \cdot \mathbb{E}(M^{(n-1)})^q \cdot \ldots \cdot \mathbb{E}(M^{(1)})^q \cdot \mathbb{E}(M^0)^q \qquad (1.25)$$

Assuming now further that the $M^{(n)}$ converge in distribution, say to M, we have

$$T(q) := \lim_{n\to\infty} \frac{-1}{n} \log_2 \mathbb{E} S_n(q) = -1 - \log_2 \mathbb{E}[M^q]. \qquad (1.26)$$

Using the relations (1.16), (1.20), and $\tau^* \leq T^*$ (see [28]), and combining them with results in [10, 1, 3, 31], we get that for every α,

$$\dim(K_\alpha) = f(\alpha) = \tau^*(\alpha) = T^*(\alpha) \qquad \text{almost surely.} \qquad (1.27)$$

To demonstrate how MFA applies to conservative cascades and what sort of numerical results it can yield in this case, we use the wavelet-based approach mentioned earlier. For convenience, we will also deal with the wavelet coefficients of the distribution μ rather than the ones of Y. The former are given by

$$w_{j,k} := \int \psi_{j,k}(t)\, \mathrm{d}\mu(t). \qquad (1.28)$$

Using the Haar wavelet, we get with (1.23) the explicit expression

$$2^{-n/2} w_{-n,k_n} = \mu(I_{2k_n}^{n+1}) - \mu(I_{2k_n+1}^{n+1}) = (M_{2k_n}^{n+1} - M_{2k_n+1}^{n+1}) \prod_{i=0}^{n} M_{k_i}^i. \qquad (1.29)$$

Thus, we compare the increment-based MFA (in terms of α, S and τ) of Y to the wavelet-based MFA (in terms of $\tilde{\alpha}$, $\tilde{S}$ and $\tilde{\tau}$) of μ . Due to the fact that

$M_{2k_n}^{n+1} - M_{2k_n+1}^{n+1} = 2M_{2k_n}^{n+1} - 1$ we have

$$\mathbb{E}\tilde{S}_n(q) = \mathbb{E}\sum_{k_n=0}^{2^n-1} |2^{n/2} w_{-n,k_n}|^q = 2^{nq} S_n(q) \cdot \mathbb{E}|2M^{(n+1)} - 1|^q.$$

This gives immediately

$$\tilde{T}(q) = -q + T(q). \tag{1.30}$$

More generally, this relation holds for any choice of mother wavelet which is supported on $[0,1]$, provided the multipliers M_k^n are all identically distributed. This holds because the scaling properties (1.23) of μ allow us to write the wavelet coefficients in this case as $2^{n/2} \cdot M_{k_n}^n \cdot \ldots \cdot M_{k_1}^1$ times a random factor which is independent of $M_{k_i}^i$ and which is distributed as $w_{0,0}$ (compare also [2]).

In order to be able to say more about $\tilde{\tau}(q)$ for the Haar wavelet, we make an assumption which guarantees that the Haar wavelet coefficients don't decay too fast (compare (1.9)), i.e. the prefactor RHS in (1.29) doesn't become too small. Therefore, let us assume in addition that there is some $\varepsilon > 0$ such that for all n, $|2M^{(n+1)} - 1| \geq \varepsilon$ almost surely. Then for all t, $(1/n)\log(2M_{2k_n}^{n+1} - 1) \to 0$, and

$$\tilde{\alpha}(t) = -1 + \lim_{n\to\infty} \frac{1}{-n\log 2} \log\left(2^{-n/2}|\mu(I_{k_n}^n|\right) = -1 + \alpha(t). \tag{1.31}$$

Observe that this is precisely the relation we expect between the scaling exponents of a process and its (distributional) derivative – at least in nice cases. Moreover, differentiating (1.30) and recalling (1.21), we get $\tilde{T}'(q) = -1 + T'(q)$ which is in agreement with (1.31). Thus, both the increment-based and wavelet-based MFA yield the same results for conservative binomial cascades with multipliers bounded away from $1/2$. For a more detailed wavelet-based analysis of conservative cascades, we refer to [14, 28].

1.3.3 On the multifractal nature of network traffic

While multifractals are new to the networking area, they have been applied in the past – mainly for descriptive purposes – to such diverse fields as the statistical theory of turbulence, the study of strange attractors of certain dynamical systems, and more recently, to physical based rain and cloud modeling; see for example [9, 17] and references therein. In the networking context, multifractals and their ability to account for time-dependent scaling laws offer great promise for describing irregular phenomena that are localized in time. The latter are typically associated with network-specific mechanisms that operate on small time scales and – depending on the state of the network – can be expected to have a more or less severe impact on how the packets within individual connections are sent across the network. Empirical evidence in support of complex within-connection or local traffic characteristics in measured wide-area traffic that can be traced to the dominant TCP/IP protocol hierarchy of IP networks has been reported in the original comprehensive analysis of WAN traces by Paxson and Floyd [25], and more recently, in work by Feldmann et al. [12]. The original findings of multifractal scaling behavior of measured aggregate WAN traffic are due to Riedi and Levy-Vehel [30] (see also [36]), followed by a similar study by Mannersalo and Norros [24] involving measured ATM WAN traces (for an earlier discussion on multifractal scaling and measured LAN traffic, see also [35]).

Motivated by the empirically observed multifractal scaling behavior in measured WAN traffic by Riedi and Levy-Vehel [30], Feldmann et al. [11] (see also [14]) present a more detailed investigation into the multifractal nature of network traffic and bring multifractals into the realm of networking by providing empirical evidence that WAN traffic is consistent with multifractal scaling because IP networks appear to act as conservative cascades. In particular, they demonstrate that (i) conservative cascades are inherent to wide-area network traffic, (ii) multiplicative structure becomes apparent when studying data traffic at the TCP layer, and (iii) the cascade paradigm appears to be a traffic invariant for WAN traffic that can co-exist with self-similarity. By systematically investigating the causes for the observed multifractal nature of measured network traffic, they observe that the packet arrival patterns within individual TCP connections (where one or more TCP connections make up a session) appear to be consistent with a multiplicative structure. The latter, they argue, seems to be mainly caused by networking mechanisms operating on small time scales, and results in aggregate network traffic that exhibits multifractal scaling behavior over a wide range of small time scales. Although it is tempting to invoke the TCP/IP protocol hierarchy of modern data networks for motivating the presence of an underlying conservative cascade construction (e.g., a web session generates requests, each request gives rise to connections, each connection is made up of flows, flows consist of individual packets), Feldmann et al. demonstrate that the multiplicative structure associated with a conservative cascade construction is most apparent when studying network traffic at the TCP layer, where the network behavior (i.e., the way the packets within a TCP connection are sent across the network) is largely decoupled from the user behavior. Moreover, Feldmann et al. suggest that the transition from multifractal to self-similar scaling occurs around time scales on the order of the typical round-trip time of a packet within the network under consideration.

While this work leaves open the "big" question "Why are packets within individual TCP connections distributed in accordance with a conservative cascade construction?" it clearly identifies the TCP layer as the most promising place in the networking hierarchy to search for the physical reasons behind the observed multifractal scaling behavior of measured network traffic and/or behind the conjecture that modern data networks act in a manner consistent with conservative cascades. Clearly, progress on these problems will require a close collaboration with networking experts. Realizing that it is difficult to think of any other area in the sciences where the available data provide such detailed information about so many different facets of behavior, there exists great potential for coming up with intuitively appealing, conceptually simple and mathematically rigorous statements as to the causes and effects of multifractals in data networking. Put differently, for multifractals to have a genuine impact on networking, their application has to move beyond the traditional descriptive stage and has to be able to answer question as to why network traffic is multifractal (i.e., physical explanation in the network context) and how it may or may not impact network performance (i.e., engineering).

1.3.4 Multiplicative structure and log-normality

The observed multifractal nature of measured WAN traffic over small time scales and the empirical evidence discussed above in support of an underlying conservative cascade mechanism responsible for the multifractal scaling phenomenon imply that over those fine time scales, network traffic is *multiplicatively* generated. In other words, at the microscopic level where the network (via the underlying protocols and end-to-end congestion control mechanisms) determines how the individual packets of a connection are sent across a given link in the network, the traffic rate process (i.e., total number of packets or bytes per small time unit) is the product of a large number of more or less independent "multipliers." In contrast, we have seen that at the macroscopic level or over large time scales, user and/or session characteristics are mainly responsible for the observed self-similar scaling behavior of network traffic and that over those time scales, the traffic rate process is *additive* in nature; that is, the sum of a large number of more or less independent "summands" where the individual summands or connections exhibit heavy-tailed distributions with infinite variance for their sizes or durations.

Intuitively, this distinction between the additive and multiplicative structure of measured network traffic over large and small time scales, respectively, can be best explained when considering an individual TCP connection. When viewed over large enough time scales, all we observe is the total workload M^0 (in bytes or packets) that is sent over the network during the connection's lifetime and for simplicity, we assume in general that the connections traffic rate $X_c^{(m)}$ is constant and that the connection's duration is unity. However, when zooming in onto finer time scales, we observe that a certain fraction of the total workload was sent during the first half of the connection's lifetime and the rest in the second half. Continuing inductively, we find that the workload emitted by the connection during a time interval of length 2^{-n} (which corresponds to a certain level of aggregation m) is of the form

$$X_c^{(m)} := M^n \cdot M^{n-1} \cdots M^1 \cdot M^0, \tag{1.32}$$

where the multipliers M^k reflect the "state of the network" and determine the amount of workload that the connection can send across the link at any given point in time. Small multipliers suggest heavy competition for the link, while large multipliers indicate that the connection can temporarily transmit at close to full speed.

As we have seen earlier, the idea of successively fragmenting the total workload into parts leads naturally to a *multiplicative* process or cascade. While the networking application justifies our choice of considering conservative cascades, our focus on an underlying binomial structure for the cascades is for simplicity. On mild independence assumptions on the multipliers (they should form a certain martingale) we are assured that we can talk about the limit of infinitely fine scales ($n \to \infty$) and that this limit has interesting statistical properties. In fact, by experimenting with turning a constant bit-rate connection into a highly bursty one via an appropriately chosen conservative binomial cascade construction [14], we find that the latter can closely match the way networking mechanisms operating on small time scales determine the actual flow of packets/bytes over the duration of a TCP connection. Moreover, when the traffic rate over a small time interval is described in terms of a conservative binomial cascade, it is explicitly multiplicative in nature (1.32); and as a result, the marginals of the traffic rate process over small time scales will automatically be approximately *lognormal* (e.g., apply CLT to the random variables $\log M^k$).

1.4 TOWARD COMPLETE DESCRIPTIONS OF NETWORK TRAFFIC

The empirical finding that measured WAN traffic contains an additive component as well as a multiplicative component provides new motivation for and insights into developing a more complete description of the dynamic nature of actual network traffic. In the following, we discuss a simple workload model that exhibits self-similar as well as multifractal scaling but is not consistent with measured network traffic. Then we illustrate the changes that are required to turn this simple model into one that is consistent with actual traffic, not only with respect to the large-time and small-time scaling behavior of measured aggregate traffic rate processes, but also at the different layers in the IP protocol hierarchy.

1.4.1 A simple multifractal workload model

To start, we consider the workload model discussed in Section 1.2.3 where (i) user-initiated sessions arrive in accordance to a Poisson process, (ii) bring with them a workload (e.g., number of bytes, packets, flows, or TCP connections, session duration) that is heavy-tailed with infinite variance, and (iii) distribute the workload over the lifetime of the session at a constant rate. A result by Kurtz [19] states that over large enough time scales, the fluctuations of the aggregate traffic rate around its mean value are well described by FGN, for a very general class of within-session traffic rate processes that includes the special case of constant bit-rate sessions. Recall that the self-similar scaling property over large time scales (or equivalently, long-range dependence) is essentially due to the fact that the session sizes exhibit infinite variance, and that approximate Gaussianity follows from an application of the CLT, i.e., from aggregating over a large number of independent sessions whose individual traffic rates are sufficiently "tame."

To incorporate multiplicative structure into this simple traffic description, we simply modify property (iii) above and require that the constant within-session traffic rate processes are replaced by multiplicative processes, or more precisely, by independent and identically distributed multifractals generated by appropriately chosen conservative binomial cascades with associated partition function $\tau(q)$ (or the more informative multifractal spectrum f). This modified workload process is a generalization of Kurtz's model by allowing within-session traffic rates to be multifractals. Since Kurtz's model is known to be insensitive to the particular within-session traffic dynamics, the self-similar scaling property over large time scales remains intact, even for multifractal within-session structure, and represents the additive component of network traffic, which is mainly due to the global characteristics of user-initiates sessions. However, when viewed over small time scales, this modified work load process will also exhibit multifractal scaling, not only at the session level, where it does so by definition, but also at the aggregate level. In fact, it can be shown that the superposition of i.i.d. conservative binomial cascades also exhibits multifractal structure, with a multifractal spectrum that is identical to the one of a "typical" session-related conservative binomial cascade. To illustrate, let μ and ν be two multifractals generated by two (possibly different) conservative binomial cascades. It is easy to see that independent of their supports, for the multifractal $\mu + \nu$ obtained by superposing μ and ν, we have

$$\tau^{\mu+\nu}(q) = \min(\tau^{\mu}(q), \tau^{\nu}(q)), \tag{1.33}$$

for all $q \geq 0$. For a proof, simply use that for all positive a, b and q, we have $(a^q + b^q)/2 \leq (\max\{a,b\})^q \leq (a+b)^q \leq (2\max\{a,b\})^q \leq 2^q(a^q + b^q)$, and hence

$$S_n^{\mu}(q) + S_n^{\nu}(q) \leq 2 \cdot S_n^{\mu+\nu}(q) \leq 2^{q+1}\left(S_n^{\mu}(q) + S_n^{\nu}(q)\right).$$

If the supports of μ and ν are disjoint, we have $S_n^{\mu+\nu}(q) = S_n^{\mu}(q) + S_n^{\nu}(q)$ and (1.33) holds for all q. However, for more general cascades with overlapping support, we will typically see $\tau(q) > \min(\tau^{\mu}(q), \tau^{\nu}(q))$ for negative q.

Assuming now that μ and ν have the same $\tau(q)$ and taking the Legendre transform we see that the superposition $\mu + \nu$ has the same spectrum $f(\alpha)$ in the increasing part, that is for small α ($\alpha < \overline{\alpha}$) which correspond to the bursty part of the multifractal. For α larger than the expected value $\overline{\alpha}$, corresponding to the smoother parts, we may observe a smaller $f(\alpha)$. In other words, the superposition has a tendency towards more bursts and fewer smooth parts.

This is natural since bursts of one multifractal may overwhelm some smooth parts of the other.

Thus, the built-in multifractal within-session structure causes the overall traffic rate process to be multiplicative over small time scales, thereby accounting in a parsimonious manner for the effect that the network has on the small-time scale dynamics of traffic rates on individual links within the network.

1.4.2 The additive and multiplicative nature of network traffic

Note that the above generalization of Kurtz's workload model that allows for multifractal within-session traffic rates is not consistent with measured data. In fact, Feldmann et al. [11] present empirical evidence that the observed within-session structure is itself a complicated mixture of additive and multiplicative components, and only by investigating network traffic at the TCP level (e.g., in terms of port-to-port flows) is it possible to clearly isolate the multiplicative structure in measured network traffic. Using the findings from yet another empirical traffic study (see Feldmann et al. [12]), we also know that the overall number of TCP connections per time unit exhibits self-similar scaling behavior for time scales on the order of seconds and beyond. Thus, to get a workload model for wide-area traffic that combines additive and multiplicative structure and is consistent with measured data, we simply modify the multifractal version of Kurtz's process and require that (i) TCP connections arrive in accordance to a self-similar process, i.e., the fluctuations around the mean of the total number of TCP connection arrivals per time unit follows a FGN; (ii) the TCP connections' workload is heavy-tailed with infinite variance; and (iii) the workload of a TCP connection is distributed over the connection's lifetime in a multifractal fashion, i.e., according to a conservative binomial cascade.

To see that the latter model has the desired large-time and small-time scaling properties and hence is in agreement with the observed additive and multiplicative properties of actual network traffic, we keep (ii) and (iii) as is, but note that the self-similar scaling property for the aggregate TCP connection traffic rate can be accomplished by relying on the underlying session structure of the original Kurtz model. That is, user-initiated session continue to arrive in a Poisson fashion, but the session workload is now expressed in terms of the number of TCP connection that make up a particular session and remains to be heavy-tailed with infinite variance; for consistency, we assume that the TCP connections within a session arrive in such a way that they don't overlap with one another. It is then easy to see that this two-tier approach to describing aggregate WAN traffic yields the additive traffic component via the TCP-connection-within-session structure and the multiplicative component via the dynamics prescribed for the packets within individual TCP-connections. Moreover, this two-tier approach is also fully consistent with measured Internet traffic at the different layers in the TCP/IP protocol hierarchy (e.g., see [12]).

1.4.3 Toward a comprehensive study of network performance

The attractive feature of the above structural model for wide-area traffic is that it is consistent with measured traffic at all levels of interest and that it accounts in a parsimonious manner for both the global or large-time scale as well as local or small-time scale 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 earlier that by choosing an appropriate generator for the generic underlying conservative binomial cascade for the within-connection traffic rate process, we are able to obtain the same overall multifractal scaling as captured by the multifractal spectrum associated with the generic cascade model for the individual TCP connections.

The practical relevance for such a structural 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 more 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 often 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 [15], [16], and [32]).

In a first attempt to allow for a more complete description of network traffic, Riedi et al. [31] (see also [26]) emphasize performance aspects of descriptive traffic models with additive and multiplicative structures. Working in the wavelet domain, they discuss in [31] a multiplicative model based on binomial cascades which exhibits the multifractal properties observed in measured network traffic at small scales and, in addition, matches the self-similar behavior of traffic over large time scales. Their model becomes approximately additive at large scales, as the variance of the cascade generator decreases with increasing scale, explaining why a purely multiplicative model can be consistent with an additive property in the limit of large scales. Riedi et al. also provide initial evidence that models which allow for a more complete description of network traffic, in particular its multifractal behavior, typically outperform additive Gaussian models in the context of specific performance problems [26].

1.5 CONCLUSION

One of the implications of the discovery of self-similar or multifractal scaling behavior in measured network traffic has been the realization that network traffic modeling and performance analysis can and should no longer be viewed as exercises in data fitting and queueing theory or simulations. Instead, relevant traffic modeling has become a natural by-product of a renewed effort that aims at gaining a physical (i.e., network-related) understanding of the empirically observed scaling phenomena. Moreover, the novel insights gained from such a physical-based understanding of actual network traffic dynamics often allows for a qualitative assessment of their potential impact on network performance, when more quantitative methods appear to be mathematically intractable or are not yet available. While traditional performance modeling has mainly lived in the confines of mathematically tractable queueing models, the observed scaling properties of measured network traffic and the constantly changing nature of today's networks strongly suggest a shift away from focusing exclusively on quantitative methods for assessing the wide range of network performance-related problems towards achieving instead a more qualitative understanding of the implications of the dominant features of measured network traffic on relevant networking issues. While supporting such a qualitative knowledge –where possible– through quantitative analysis is clearly desirable, we believe that the development of an ubiquitous, stable, robust and high-performance networking infrastructure of the future will depend crucially on a qualitative rather than quantitative understanding of networks and network traffic dynamics.

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. The vast majority of currently used models for network traffic completely ignore the fact that the dynamic nature of packet traffic over a given link is the result of a combination of source/user behavior and highly nonlinear interactions between the individual users and the network. The search for a physical explanation of the observed multifractal nature of measured traffic at the packet-level is intimately related to trying to sort out these complicated interactions and to abstract them to a level that is intuitively appealing, conforms to networking reality, and captures and explains in a mathematically rigorous manner empirically observed phenomena. Clearly, a pre-requisite for succeeding in this endeavor is a close collaboration with networking experts who are familiar with the details of the various protocols and control mechanisms that operate at the different layers within the hierarchical structure of modern-day data networks and who are aware of the problems that are associated with the highly dynamic, constantly changing, and extremely heterogeneous nature of today's communication networks.

REFERENCES

1. M. Arbeiter and N. Patzschke, "Self-similar random multifractals," *Math. Nachr.* **181**, pp. 5–42, 1996.

2. E. Bacry, J. Muzy, and A. Arneodo, "Singularity spectrum of fractal signals from wavelet analysis: Exact results," *J. Stat. Phys.* **70**, pp. 635–674, 1993.

3. J. Barral, "Continuity, moments of negative order, and multifractal analysis of Mandelbrot's multiplicative cascades," *Université Paris Sud, thèse no. 4704*, 1997.

4. D. R. Cox, "Long-range dependence: A review," In H. A. David and H. T. David, editors, *Statistics: An Appraisal*, pp. 55–74, Iowa State University Press, Ames, Iowa, 1984.

5. I. Daubechies, *Ten Lectures on Wavelets.* SIAM, Philadelphia, 1992.

6. J.-D. Deuschel and D. W. Stroock, *Large Deviations.* Academic Press, New York, 1994.

7. R. Ellis, "Large deviations for a general class of random vectors," *Ann. Prob.* **12**, pp. 1–12, 1984.

8. 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.

9. 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.

10. K. J. Falconer, "The multifractal spectrum of statistically self-similar measures," *J. Theor. Prob.* **7**, pp. 681–702, 1994.

11. A. Feldmann, A. C. Gilbert, and W. Willinger, "Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic," *Computer Communication Review* **28**, No. 4, Proc. ACM/Sigcomm'98, Vancouver, Canada, September 1998, pp. 42–55, 1998.

12. 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, pp. 5–29, April 1998.

13. A. C. Gilbert, W. Willinger and A. Feldmann, "Visualizing multifractal scaling behavior: A simple coloring heuristic," *Proc. of the 32nd Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, November 1998.

14. A. C. Gilbert, W. Willinger and A. Feldmann, "Scaling analysis of conservative cascades, with applications to network traffic," *IEEE Trans. Info. Theory, Special Issue on "Multiscale statistical signal analysis and its applications"*, April 1999 (to appear).

15. M. Grossglauser and J.-C. Bolot. On the relevance of long-range dependence in network traffic. *Computer Communication Review* **26**, No. 4, Proc. ACM/Sigcomm'96, Stanford, CA, pp. 15–24, 1996.

16. D. P. Heyman and T. V. Lakshman. "What are the implications of long-range dependence for VBR-video traffic engineering?" *IEEE/ACM Transactions on Networking* **4**, pp. 301–317, 1996.

17. 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.

18. S. Jaffard, "Local behavior of Riemann's function," *Contemporary Mathematics* **189**, pp. 287–307, 1995.

19. 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.

20. 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.

21. J. B. Levy and M. S. Taqqu, "Renewal reward processes with heavy-tailed interarrival times and heavy-tailed rewards," preprint, 1997.

22. B. B. Mandelbrot, "Long-run linearity, locally Gaussian processes, H-spectra and infinite variances," *International Economic Review* **10**, pp. 82–113, 1969.

23. B. B. Mandelbrot, "Intermittent turbulence in self similar cascades: Divergence of high moments and dimension of the carrier," *J. Fluid. Mech.* **62**, pp. 331–358, 1974.

24. P. Mannersalo and I. Norros, "Multifractal analysis of real ATM traffic: a first look," *COST257TD*, 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. V. Ribeiro, M. S. Crouse, R. Riedi, and R. G. Baraniuk, "Simulation of non-Gaussian long-range-dependent traffic using wavelets," *Proc. Sigmetrics'99*, May 1999 (to appear).

27. R. H. Riedi, "An improved multifractal formalism and self-similar measures," *J. Math. Anal. Appl.* **189**, pp. 462–490, 1995.

28. R. H. Riedi, "Multifractal processes," preprint, 1999.

29. R. H. Riedi and B. B. Mandelbrot, "Exceptions to the multifractal formalism for discontinuous measures," *Math. Proc. Cambr. Phil. Soc.* **123**, pp. 133–157, 1998.

30. R. Riedi and J. L. Véhel, "Multifractal properties of TCP traffic: a numerical study," *Technical Report No 3129, INRIA Rocquencourt, France*, Feb. 1997. Also available at: `www.dsp.rice.edu`.

31. R. Riedi, M. S. Crouse, V. Ribeiro, and R. G. Baraniuk, "A multifractal wavelet model with application to TCP network traffic," *IEEE Trans. Info. Theory, Special Issue on "Multiscale statistical signal analysis and its applications"*, April 1999 (to appear).

32. B. K. Ryu and A. Elwalid, "The importance of long-range dependence of VBR video traffic in ATM traffic engineering: Myths and realities," *Computer Communication Review* **26**, No. 4, Proc. ACM/Sigcomm'96, Stanford, CA, pp. 3–14, 1996.

33. G. Samorodnitsky and M. S. Taqqu, *Stable non-Gaussian processes: Stochastic models with infinite variance.* Chapman and Hall, New York, London, 1994.

34. M. S. Taqqu and J. Levy, "Using renewal processes to generate long-range dependence and high variability," In E. Eberlein and M. S. Taqqu, editors, *Dependence in Probability and Statistics*, pp. 73–89, Boston, Birkhäuser, 1986.

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

36. J. L. Véhel and R. Riedi, "Fractional Brownian motion and data traffic modeling: The other end of the spectrum," *Fractals in Engineering*, pp. 185–202, Springer 1997.

37. 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.

38. 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.