Long-Range Dependence and Data Network Traffic

Walter Willinger

Vern Paxson

Rolf H. Riedi

Murad S. Taqqu

ABSTRACT This is an overview of a relatively recent application of *long-range dependence* (LRD) to the area of communication networks, in particular to problems concerned with the dynamic nature of packet flows in high-speed data networks such as the Internet. We demonstrate that this new application area offers unique opportunities for significantly advancing our understanding of LRD and related phenomena. These advances are made possible by moving beyond the conventional approaches associated with the wide-spread "black-box" perspective of traditional time series analysis and exploiting instead the physical mechanisms that exist in the networking context and that are intimately tied to the observed characteristics of measured network traffic. In order to describe this complexity we provide a basic understanding of the design, architecture and operations of data networks, including a description of the TCP/IP protocols used in today's Internet. LRD is observed in the large scale behavior of the data traffic and we provide a physical explanation for its presence. LRD tends to be caused by user and application characteristics and has little to do with the network itself. The network affects mostly small time scales, and this is why a rudimentary understanding of the main protocols is important. We illustrate why multifractals may be relevant for describing some aspects of the highly irregular traffic behavior over small time scales. We distinguish between a time-domain and wavelet-domain approach to analyzing the small time scale dynamics and discuss why the wavelet-domain approach appears to be better suited than the time-domain approach for identifying features in measured traffic (e.g., relatively regular traffic patterns over certain time scales) that have a direct networking interpretation (e.g., "round trip" time behavior).

AMS Subject classification: Primary 94-02; secondary 60G18.

Keywords and phrases: self-similarty, multifractals, computer networks

1 Introduction

While the popular "black-box" approaches to time series have their merits, especially in areas of applications where the available measurements are limited in scope, we argue here that for highly engineered complex systems such as the Internet, they are of little value. Ignoring the rich semantic context present in traffic measurements means missing out on new discoveries.

As is cogently discussed in [34], it seems to be the rule rather than the exception that many naturally occurring empirical records or time series violate the assumption of independence or correlations that decay exponentially-fast. Instead, the data often suggest the presence of strong temporal correlations over large lags, manifesting themselves in a phenomenon called *persistence* that captures the intuition behind the empirical observation that measurements of many physical or engineered systems naturally tend to "cluster." That is, they show fluctuations above or below a given level over extended periods of times that are much more common than in the case of independent or Markovian-type observations. Mathematically, persistence or slowly decaying correlations can be parsimoniously described via the notion of *long-range dependence* (LRD), which was brought to the attention of statisticians and probabilists by Mandelbrot and his co-workers [51, 53, 54], mainly through applications in such areas as hydrology (e.g., annual river flow data) [56, 55, 57], geophysics [58, 59], and finance (e.g., stock prices) [49].

As far as the networking application area is concerned, the last decade has seen an enormous increase in empirical studies of high-quality and high-volume data sets of traffic measurements from a variety of different data networks, but especially from different links within the global Internet. These studies describe pertinent statistical characteristics of the temporal dynamics of the "packet" or bit rate processes (i.e., the time series representing the number of packets or bits per time unit, over a certain time interval) as seen on a link within the network. They provide ample evidence that measured packet traffic exhibits LRD, and hence when viewed within the right range of scales, the traffic appears to be *fractal-like* or *self-similar*, in the sense that a segment of the traffic measured at some time scale looks or behaves just like an appropriately scaled version of the traffic measured over a different time scale. In effect, this empirically-based effort toward describing actual data network traffic has demonstrated that self-similarity provides an elegant and compact mathematical framework for capturing the essence behind the observed "burstiness" or scale-invariance in measured traffic traces.

This behavior is in sharp contrast to what one observes in the usual models for data traffic, models that in general lack validation against measured data traffic time series or *traces*. The conventional models are based on presumed conditions of the dynamics of actual data network traffic, and are typically assumed to be relevant for data traffic because of their apparent similarities or close associations with a long line of highly successful models of voice traffic (e.g., see [79] and references therein). A hallmark of these traditional voice-based data traffic models is an exponentially fast decaying correlation function, implying that time-aggregation quickly results in white noise traffic characterized by the absence of any significant temporal correlations, and capable only of reproducing the observed bursty behavior of measured traffic over a narrow range of time scales.

During the last few years, the discovery of ubiquitous LRD in measured traffic from data networks has led to two very different and essentially disjoint research efforts. The first, purely descriptive in nature, has followed the line of traditional time series analysis, where the context in which the data are measured and collected is of little or no significance. The main emphasis is on statistical inference (e.g., model fitting and parameter estimation), and the end product is generally a model that describes the given data well. This approach ignores fundamental issues concerning the role of traditional model fitting when there are large numbers of voluminous data sets and when the actual data networks are known to undergo constant and often significant changes. The approach also revisits arguments concerning the validity of inferring asymptotic statistical properties such as LRD from a finite set of data, first put forward more than 20 years ago when LRD was a hotly debated issue in areas of applications such as hydrology and finance (see for example [43], and the discussions in [53]).

In contrast to this wide-spread descriptive approach to dealing with LRD traffic traces, the networking application has motivated a very different research effort; an effort that essentially abandons traditional time series analysis thinking all together, insists on fully exploiting the context in which the data are collected, and focuses on providing physical explanations for the observed LRD phenomenon that directly relate to underlying networking mechanisms (and can be validated against measured data). The results are constructive models of data network traffic that are mathematically solid, empirically consistent with the data, and can be easily explained to networking experts because the constructions capture the essence of how traffic is generated in the networking context in the first place. Clearly, this effort requires at times intimate familiarity with networking-specific details. However, learning about the application area so as to be able to converse with the experts can yield unexpected findings that compel revisiting with conventional wisdom. In the case at hand, we show how this application-centered approach has lead to new scientific discoveries related to the origins of LRD in measured aggregate traffic; we discuss how understanding the physics behind the LRD phenomenon in the networking context (i.e., user/application characteristics) points up the additional need to understand hitherto unexplored structural properties of measured data network traffic over small time scales; we explain why the mathematical framework of multifractals may be relevant for capturing and describing this newly observed, highly irregular fine-scale traffic behavior; and we put forward arguments for why there is hope for finding a physical explanation of this phenomenon that is as intuitive, rigorous, and appealing as the one that is now available for the LRD phenomenon.

In the spirit of convincing the interested reader of the value of the "getting to know your network" approach over the "black-box" approach, we follow closely the presentations in [66, 88]. Each section of this article starts by providing successively more detailed information about the networking application, starting from a discussion of the basic concepts of data networks (as compared to the traditional voice networks) in Section 2.1, to a brief sketch in Section 3.1 of the basic design principles behind many of today's data networks, including the notion of the TCP/IP protocol architecture. This is followed up in Section 4.1 with a primer on TCP, the predominant transfer protocol in today's Internet. The corresponding Sections 2.2, 3.2 and 4.2 illustrate how knowing about these different aspects of data networks helps in formulating new and relevant research problems that remain invisible to the "black-box" approach but contribute in a fundamental way to the ultimate goal of gaining a solid understanding of the highly complex nature of traffic flows in large-scale data networks such as the global Internet. We conclude in Section 5 with thoughts on how to achieve this ultimate goal. To assist the reader we include at the end a list of the main abbreviations used in this article.

2 Data networks, data network traffic, and LRD

To appreciate the fundamental differences between voice and data networks, we first give a brief description of the corresponding network architectures and technologies, and argue why a rudimentary understanding of these differences is crucial when analyzing traffic traces from these networks.

2.1 Networking 101: Voice versus data networks

The public switched telephone network (PSTN) is an example of a *connection-oriented* network, where a *connection* or *call* between two hosts (e.g., phones, fax machines, computers) represents the basic building block. The defining feature of connection-oriented networks is that resources are reserved at the beginning of a connection or call, freed up when the call terminates, and are not shared with any other connections for the entire duration of the call.

One major implication of this design principle for the PSTNs has been to engineer them in a *circuit-switching* fashion. There are *routers* or "waystations" internal to the network, which are responsible for forwarding traffic from one link to the next so that it ultimately reaches its destination and keep track of each currently active connection. When new traffic arrives, the routers look up its corresponding connection to determine where to forward the traffic. This abstraction is termed providing "virtual circuits," because the network behaves as though it provides a direct physical circuit from the traffic source all the way to its destination. Many concurrent connections can be easily "multiplexed," that is, share a common (expensive) wire or *link*, by allocating a fixed amount of the link's capacity to each connection. When a new call request arrives, it is easy to look at a link's current load and determine whether the link has sufficient capacity to carry the additional load. Circuit-switching and connection-oriented service are ideally suited for voice traffic that is relatively homogeneous and predictable, and, from a signaling perspective, spans long time scales. In this setting, *Quality-of-Service*, or QoS, in short, reduces to the simple notion of *call blocking* (denial of service). Billing is also easy, at least conceptually. Two important disadvantages are wasted resources (resources are dedicated to a connection irrespective of how much traffic is sent over the line) and the need for "smart" routers or switches (each switch has to know the "state" of every active connection it sees).

These disadvantages became apparent as circuit-switching networks began to be used more frequently for *data communications* between *hosts*, that is, endpoints which are themselves full-blown computers. First, a typical data connection is much "burstier" than a voice call because much of the time the line tends to be idle. This makes circuit-switching inefficient. Second, computer-generated data traffic is also much more variable than voice traffic because data connections range from extremely short duration to extremely long ones, from extremely low-rate to extremely high-rate. These fundamental differences between data and voice traffic have led to a design for data networks where the fundamental building block is a *packet* of data or "datagram." Each packet is self-contained in the sense that its *header* contains complete "addressing" information, and the routers need only inspect the header of the packet to determine its destination and forward it through the network. It is transmitted independently from the other packets. Consequently, the routers do *not* keep track of each currently active connection, and can forget a packet as soon as it has been forwarded. The shift from circuit-switching to a *packet-switching* technology for data traffic was advocated around 1970, when data networks were very small. Interestingly enough, the basic concept of packet-switching has remained essentially the same during the past 30 years and continues to represent one of the few effective technologies for data communication in networks such as the global Internet—an *internetwork* made up of 86,000+ separate networks, with a total of about 90 million hosts as of September 2000, and counting.

The implications of adopting a packet-switching technology have been profound and far-reaching. In contrast to circuit-switching, the concept of QoS can no longer be simply identified with call blocking; it has become more complex and multi-faceted. At the same time, packet-switching outperforms circuit-switching as far as efficiency, robustness, and flexibility are concerned. First, the physical links in packet-switched environments can be utilized much more efficiently, because there is no notion of reserving part of their capacity for each active connection. Newly arriving packets will grab any capacity available in the network and benefit from it. Each packet in the network competes with all the others—if there happens to be little competing traffic along a particular path, then a connection using it can enjoy the entire capacity (*bandwidth*) of the path, and transfer its data very quickly. If many connections compete along the same path, then each will receive a (perhaps unfair) portion of the available bandwidth. Second, the routers in a packet-switched network function without any notion of currently active connections. If a router or link *fails*, it is a simple matter to route around the failure—the traffic is simply sent to a different set of routers and links. The new routers have no problem accepting the rerouted traffic because, as far as they can tell, it is not in any way "new" traffic—they have no notion of "current" traffic and hence no problem accepting traffic they did not until that very moment know existed. This situation is very different from that in a circuit-switched network, in which the routers cannot easily accept rerouted traffic because they have no knowledge of the corresponding virtual circuit.

The ability of a packet-switched network to transparently route around failures without perturbing active connections buys enormous *robustness*: the network can continue to operate and successfully deliver data even in the face of major equipment failure. It also allows for great flexibility when connecting new hosts or disconnecting existing endpoints. Finally, links can become overloaded because packets arrive for transmission along them at a rate exceeding the capacity of the link. Such packets will be *buffered* awaiting transmission along the link, but if the excess rate is sustained—a condition termed *congestion*—then ultimately the buffers will be exhausted and some packets must be discarded, or *dropped*. To effectively communicate with one another in such a lossy packet-switched environment, the network's end hosts and routers rely on a common "language"—a standard set of *protocols*. Protocols will be discussed in more detail in the following sections.

2.2 Data network traffic and LRD

Voice traffic in PSTNs appears relatively easy to model. All that is needed to describe the temporal dynamic of voice traffic in a connection-oriented circuit-switched networking environment are the distributions of the interarrival times and service demands (i.e., durations) of the calls. Moreover, traditional PSTNs have (or at least used to have) highly predictable growth rates allowing for fine-tuned short- and long-term capacity planning, their offered services are strictly regulated and monitored, and their network controls and operations are fully centralized so as to be able to take full advantage of information about the network's global state. All this contributed to the popular belief of the existence of "universal laws" governing voice traffic. The most significant such law concerns the presumed and repeatedly validated *Poisson* nature of call arrivals at links in the network where traffic is highly aggregated. A close second is the well-documented observation that call durations or service demands are generally well described by *exponential distributions.** These two invariants of voice traffic provide in general a complete description of the dynamic nature of network traffic in PSTNs.

For data networks, the story is much more complicated. We begin with a description of one type of measurement, that made of traffic as it transits a given link. The measurements are made by copying either the initial part of each packet (i.e., the packet header) or the entire packet (i.e., header plus payload). In addition, packet time stamps (arrival time), sizes, and other information is saved. When viewed as "black boxes," i.e., when focusing on the mere existence of a measured packet (time stamp, packet size) and not on its "meaning" as revealed by its header and/or payload, packet traffic appears as not much more than a simple time series; univariate, if we simply focus on the packet arrival times; multivariate, if we also include packet size and possibly other status information. Following the example of voice traffic modeling where this time series analysis view was appropriate and highly successful, traditional traffic modeling and analysis for packet-switched networks has embraced an identical approach, resulting in a description of packet traffic solely in terms of arrival and service processes.

Given the arrival times and sizes of all packets that traversed a link within the network during a some time period, a natural object to derive from this data and to study is the time series $X = (X_k : k = 1, 2, \dots)$ representing the number of arriving packets (or bytes) in successive, non-overlapping time intervals of unit length (e.g., millisecond, 100 milliseconds, second, etc.). Assuming that X is second-order stationary and has zero mean (i.e., we think of X as describing the fluctuations of the traffic rate process around its mean rate), for each integer $m > 1$, we can define the *aggregated process*

*See however [19] for drastic changes in the "static" world of PSTNs, caused by—among other things—new pricing structures, the increased use of faxes and, more importantly, the use of the voice network to connect to the Internet.

$X^{(m)}$ of X at aggregation level m by setting

$$X^{(m)}(i) = \frac{1}{m}(X_{(i-1)m+1} + \cdots + X_{im}), \quad i \geq 1.$$

That is, the aggregated processes $X^{(m)}$ are obtained from X by partitioning the observation interval into non-overlapping blocks of size m and by viewing the block averages as new observations $X^{(m)}(i)$, where i denotes the block index. For each $m > 1$, $X^{(m)}$ defines a new second-order stationary and zero mean process, and the family $(X^{(m)} : m \geq 1)$ of aggregated processes is ideally suited for studying the temporal dynamic of network traffic at different time scales and for treating mathematical concepts such as self-similarity that relate statistical properties of X to those of $X^{(m)}$ through judicious scaling of time and space.

Following Cox [12], we call X *asymptotically second-order self-similar* (with *self-similarity parameter* or *Hurst parameter* $0 < H < 1$) if the second-order statistics of $m^{1-H}X^{(m)}$ converge as follows:

$$\lim_{m\to\infty} \mathrm{Var}(m^{1-H}X^{(m)}) = \sigma^2, \ 0 < \sigma < \infty, \tag{2.1}$$

and

$$\lim_{m\to\infty} r^{(m)}(k) = \frac{1}{2}((k+1)^{2H} - 2k^{2H} + (k-1)^{2H}), \tag{2.2}$$

where $r^{(m)} = (r^{(m)}(k), k \geq 0)$ denotes the autocorrelation function of the aggregated process $X^{(m)}$. X is called *exactly second-order self-similar* if for all $m > 1$, the processes $m^{1-H}X^{(m)}$ have the same second-order statistics as X; that is, $\mathrm{Var}(m^{1-H}X^{(m)}) = \mathrm{Var}(X) = \sigma^2$ and $r^{(m)}(k) = r(k) = \frac{1}{2}((k+1)^{2H} - 2k^{2H} + (k-1)^{2H})$, where $(r(k), k \geq 0)$ is the autocorrelation function of X. The form of the autocorrelation function appearing in the limit (2.2) points to the presence of long-range dependence.

Here, a second-order stationary stochastic process $X = (X_k : k = 1, 2, \ldots)$ with autocorrelation function $r(k)$ is said to exhibit *long-range dependence (LRD)* if for some $0 < \beta < 1$,

$$r(k) \sim c_1 k^{-\beta}, \quad \text{as } k \to \infty, \tag{2.3}$$

where c_1 is a positive finite constant.† In Mandelbrot's terminology [51], long-range dependence is also referred to as the *Joseph Effect* and captures the persistence phenomenon observed in many naturally occurring empirical time series; that is, the occurrence of pronounced clusters or runs of

†More general definitions of long-range and short-range dependence are possible, but for convenience, we will use here the working definition involving conditions (2.3) and (2.4) respectively. The symbol $\sim$ means "asymptotic to".

consecutive large or consecutive small values. Note that the hyperbolically slow decay of the autocorrelations of a long-range dependent process implies $\sum_k |r(k)| = \infty$. It is this non-summability of the autocorrelations that captures the essence of long-range dependence: even though the high-lag autocorrelations are individually small, their cumulative effect is of importance and gives rise to a behavior of the underlying stochastic process that is markedly different from that of the conventionally considered short-range dependent processes. Here, a second-order stationary stochastic process $X = (X_k : k = 1, 2, \dots)$ is called *short-range dependent (SRD)* if for some $0 < \rho < 1$,

$$r(k) \sim c_2 \rho^k, \quad \text{as } k \to \infty, \tag{2.4}$$

where c_2 is a positive finite constant. Thus, in contrast to LRD, SRD is characterized by an autocorrelation function that decays geometrically fast and is summable (i.e., $\sum_k |r(k)| < \infty$).

It is easy to see that when restricting the Hurst parameter H to values in the interval $(0.5, 1)$, exact or asymptotic second-order self-similarity implies LRD with $\beta = 2 - 2H$. In fact, the most striking feature of second-order self-similar processes, namely that in the limit, as $m \to \infty$, their aggregated processes $X^{(m)}$ possess the non-degenerate autocorrelation function (2.2), is intimately related to the hyperbolically slow decay (or, equivalently, to the non-summability of the autocorrelations of an LRD process). This behavior is in contrast to the traditionally studied SRD models with their geometrically fast decaying (and hence summable) autocorrelations, whose aggregated processes quickly converge to second-order pure noise; i.e., $r^{(m)}(k) \to 0, \text{ as } m \to \infty, \ k \geq 1$.

Consequently, to test for self-similarity, one uses the essentially equivalent characterizations of LRD in the time-, frequency-, or wavelet-domain. They can be exploited to measure this phenomenon quantitatively by estimating the Hurst parameter H. Time-domain estimation techniques are described in [83, 86] and include heuristics such as analysis of the *rescaled adjusted range statistics* (R/S statistic, in short; e.g., see [38, 57, 83]) and variance-time analysis of the aggregated processes [12, 83]; examples of frequency-domain techniques are the periodogram analysis [31, 37, 83] and Whittle's method [87, 9, 16, 28]. For a wavelet-domain approach, see [1, 25, 2, 3].

Leland et al. [47] introduced the self-similarity and LRD concepts in the modeling of data network traffic. Starting with the extensive analyzes of traffic measurements from Ethernet local-area networks (LANs) over a four-year period reported in [46, 47], there have been a number of crucial follow-up studies providing further evidence of the prevalence of self-similar traffic patterns in measured traffic from data networks. Most prominent among these studies are the in-depth analysis and exploration of large amounts of pre-Web wide-area network (WAN) traffic measurements in

[64, 65] and—right after the emergence of the Web—detailed investigations of Web-related traffic on the Internet [13, 14]. One of the most surprising findings from these and many other empirical studies concerns the ease with which it is possible to distinguish statistically between measured network traffic and traces generated from the traditional, commonly-used traffic models: actual traffic exhibits correlations over a wide range of time scales and the resulting aggregated processes don't converge to white noise. Traditional traffic model give rise to traffic traces that quickly become indistinguishable from white noise after aggregation over larger time scales. (For a simple visualization of this empirical observation, see Figure 4 in [47].) In short, commonly-used models for data traffic tend to focus on a very limited range of time scales and are inherently short-range dependent. On the other hand, measured traffic from actual data networks is fully consistent with long-range dependence or, equivalently, asymptotic second-order self-similarity.

2.3 When there is more to LRD than model fitting ...

The empirical finding that measured data network traffic is consistent with LRD (or equivalently, self-similar scaling over sufficiently large time scales; that is, asymptotic second-order self-similarity) has led to a revival of time series analysis for modeling packet traffic in data networks. However, traditional time series analysis, with its main emphasis on finding the "best-fitting" model for a given single data set of small–medium size is unfortunately ill-suited for identifying, capturing, and describing pertinent statistical characteristics that may be common among a large number of voluminous data sets. As a result, traditional time series analysis has resulted in an endless stream of proposed "black box" models for data network traffic, where the focus has been mainly on statistical issues related to model fitting and where the networking context in which the data had been generated and collected in the first place has been of little or no significance. Examples of such "black-box" models include processes that are inherently long-range dependent (e.g., the well-known classes of fractional Gaussian noise processes, fractional ARIMA models, and their stable counterparts) and processes that imitate LRD-behavior by relying on sufficiently highly parameterized SRD models (e.g., Markov-modulated Poisson processes, ARIMA and related time series models).

In this sense, traffic modeling has become yet another application of traditional time series methodologies. Furthermore, the ever-present questions related to the validity and/or appropriateness of inferring asymptotic statistical features such as LRD from finite data sets have been dealt with using similar arguments as, for example, put forward in the context of modeling annual river flow data or financial time series, where—in the absence of genuine physical understanding—the discussions have been mainly philosophical in nature (e.g., see [43].

For segments of the Internet research community, the situation couldn't be more different. The discovery of LRD in measured data network traffic has been met from the very beginning with a degree of curiosity that goes beyond the obvious question of how to model a single measured traffic trace that has been found to be consistent with LRD. Indeed, shortly after the initial publication of proposed self-similar models [46], an editorial in *IEEE Network* began: "This month's Proceedings of ACM SIGCOMM '93 has a fascinating paper that anyone interested in congestion control for data networks should read" [63]. These researchers wanted to know if there exists a physical explanation for the observed LRD nature of data traffic—an explanation that makes sense in the networking context and can be phrased and, more importantly, validated in terms of more elementary traffic-related entities. Thus, instead of emphasizing the purely descriptive aspect of LRD (as is done in traditional time series analysis), the question of how to physically explain LRD constitutes a new research challenge that has met with little success in areas such as finance or hydrology. To move beyond the conventional time series analysis perspective of data traffic and succeed in tackling the problem of explaining the physics behind LRD in data traffic, it is essential to know more about the design and architecture principles behind today's data networks.

3 LRD and large-time scaling: An application-layer view

In order to explain the relationship between networking and LRD, we first need to discuss in more detail the concept of a *layered protocol architecture* for data networks such as the Internet. Then, recalling that LRD involves large-time scales and leaves the small-time scales essentially unspecified, we will show how measured data network traffic lends itself naturally to a decomposition into individual components which typically include, at the highest layer, the features that explain the observed LRD property.

3.1 Networking 101: The TCP/IP protocol architecture

Today's data networks are highly engineered entities that enable the exchange of data between any set of end hosts. Due to the richness and complexity of the procedures governing such exchanges, an early design principle has been to avoid implementing a complex task such as a file transfer between two end hosts as a single module, but to instead break the task up into subtasks, each of which is relatively simple and can be implemented separately. The different modules can be thought of being arranged in a vertical stack, where each layer in the stack follows a certain "protocol;"

that is, it is responsible for performing a well-defined set of functionalities. Each layer relies on the next lower layer to execute more primitive functions, and provides services to the next higher layer. Two hosts with the same protocol architecture communicate with one another by having the corresponding layers in the two systems talk to one another. The latter is achieved by means of formatted blocks of data that obey a set of rules or conventions known as a *protocol*.

To illustrate this layering design, we consider below the five-layer TCP/IP protocol suite, which is the main protocol stack in the Internet. It consists of the physical, link, internetwork, transport, and application layers. The *physical layer* concerns the physical aspects of data transmission on a particular link, such as characteristics of the transmission medium, the nature of the signal, and data rates. Above the physical layer is the *link layer*. Its mechanisms and protocols (e.g., signaling rules, frame formats, media-access control) control how packets are sent over the raw media of individual links. Above that is the *internetwork layer*, responsible for getting a packet through an *internet*, i.e., a series of different networks. Here, the main challenge is to adequately implement all the mechanisms necessary to knit together divergent networking technologies into a single virtual network, "an internet," so as to enable data communication between sending and receiving hosts that reside on different types of networks.

The *Internet Protocol (IP)* is the internetworking protocol for TCP/IP. (*TCP* is involved in the next higher layer and will be discussed below). To send data, a host that resides on, say, an Ethernet, simply sends its data encapsulated in the internetworking protocol (i.e., IP over Ethernet). If the IP datagram needs to be sent to a host on, for example, an ATM network, the datagram will go through a router which will remove the datagram from its Ethernet encapsulation and put it into the ATM encapsulation. Clearly, internetworking protocols are crucial for enabling computers on different types of networks to communicate with one another so that networking boundaries become transparent to the users.

The layer above IP is the *transport layer*, where most commonly the *Transmission Control Protocol (TCP)* deals with end-to-end congestion control and assures that arbitrarily large streams of data are reliably delivered and arrive at their destination in the order sent. The top layer is the *application layer*, which includes a wide range of applications such as TELNET (to establish a remote connection), FTP (the File Transfer Protocol), and HTTP (the Hypertext Transfer Protocol, used for access to the World Wide Web).

With high-quality traffic measurements at hand, accurate accounting of this multi-level hierarchy of measured network traffic is possible because all the relevant information can be obtained by looking inside the collected packets. For layers up to *transport*, checking the header of each recorded packet suffices; for *application* layer analysis, though, it is also necessary to take a look at the actual payload of the packets. At this stage, network traffic is no longer a simple uni- or multivariate time series of packets, but manifests itself at the different networking layers in a variety of different forms. Consider, for example, the Internet. Starting at the top layer, namely at the application level, we can describe the traffic in terms of session arrivals, session duration, and session size (examples of sessions are remote login, file transfer, or web surfing). At the transport layer, the overall traffic can be characterized in terms of TCP connection arrivals, lengths and sizes (other components of the traffic that employ transport protocols other than TCP have their own characterization); there may be many connections corresponding to a single session. At the internetwork layer, traffic descriptions focus on either individual IP packets, or on flows (i.e., their arrival patterns, sizes, origination and destination addresses), where flows are made up of successive IP packets that satisfy certain common features. Again, there may be many packets or even flows corresponding to a single connection. Alternatively, we can view network traffic as an aggregate of packets generated by many host-host pairs. Finally, at the link layer, traffic can be dealt with by treating the individual packets as black boxes, i.e., by focusing on the mere existence of a measured packet (time stamp, packet size) and not on its "meaning" as revealed by its header.

Clearly, as a result of this hierarchy of protocol architectures, actual network traffic, that is, the flow of packets across a link inside the network, is the result of intertwined mechanisms and modes that exist at the different networking layers. Focusing on the "black box" view of packet traffic by using uni/multivariate time series does not do justice to the highly hierarchical and interconnected structure of actual data network traffic. Black boxes ignore nearly all of the gathered information and are therefore incapable of contributing significantly to an improved understanding of data networks and data network traffic.

3.2 Mathematical framework I: The on-off source model

In a first attempt to move beyond the traditional "black-box" approach to modeling packet traffic, we follow closely [91, 85] and consider a highly simplified abstraction of network traffic at the application layer where users or end hosts cycle through periods of activity (when packets are sent) and inactivity (when no packets are transmitted). In particular, as part of this abstraction, we are not concerned about the precise nature of how packets are transmitted during an activity period but assume for simplicity that they are sent at a constant rate. Thus, we consider a network with a number of users/sources or end hosts communicating with each other, where an individual source is modeled according to an on-off alternating renewal process, as follows. The source alternates between an *active* or *on* state where it sends packets into the network and an *inactive* or *off* state where it is idle and does not send any packets. Let $\{W(t),\ t \geq 0\}$, be a stationary process, where

$$W(t) = \begin{cases} 1 & \text{if time } t \text{ is an } \textit{on} \text{ interval,} \\ \\ 0 & \text{if time } t \text{ is in } \textit{off} \text{ interval.} \end{cases}$$

(We will simply say that $W(t)$ is 1 if t is *on* and 0 if t is *off*.) Viewing $W(t)$ as the reward at time t, we have a reward of 1 throughout an *on* interval, then a reward of 0 throughout the next *off* interval, then 1 again, and so on. The length of the *on* intervals are i.i.d., those of the *off* intervals are i.i.d. and the lengths of *on* and *off* intervals are independent. An *off* interval always follows an *on* interval, and it is the pair of *on* and *off* intervals that defines an interrenewal period.

Let F_{on} and F_{off} denote the cumulative distribution function of the *on* and *off* interval respectively, and let $\overline{F} = 1 - F$ denote a complementary cumulative distribution function. Let also σ_{on} and σ_{off} denote the respective variances. Assume as $x \to \infty$,

$$\text{either } \ \overline{F}_{on}(x) \sim \ell_{on} x^{-\alpha_{on}},\ 1 < \alpha_{on} < 2 \ \text{ or } \ \sigma_{on} < \infty,$$

and

$$\text{either } \ \overline{F}_{off}(x) \sim \ell_{off} x^{-\alpha_{off}},\ 1 < \alpha_{off} < 2 \ \text{ or } \ \sigma_{off} < \infty.$$

Here ℓ_{on} and ℓ_{off} are constants. (For simplification, we do not use general slowly varying functions.)

When $1 < \alpha_{on} < 2$, the distribution of the *on* times is said to be "heavy-tailed" with exponent α_{on}. Since it has infinite variance (but finite mean), the on time can be very long with relatively high probability. We will refer to the case $\sigma^2_{on} < \infty$ as "$\alpha_{on} = 2$." In this case, the *on* times are not likely to last very long. They can be for example exponential, as in the classical model of Poisson arrivals. Similar remarks hold for the distribution of the *off* times.

We are interested in understanding the behavior of the cumulative load, $A(t) = \int_0^t W(u)du$, at large times t. The load $A(t)$ has variance

$$V(t) = \text{Var}A(t) = \text{Var}\left(\int_0^t W(u)du\right) = 2\int_0^t \left(\int_0^v \gamma(u)du\right)dv \qquad (3.1)$$

where

$$\gamma(u) = EW(u)W(0) - (EW(0))^2 \qquad (3.2)$$

denotes the covariance function of W. One can show [85] that this implies that

$$V(t) \sim \sigma^2 t^{2H} \quad \text{as} \quad t \to \infty \qquad (3.3)$$

for some constant $\sigma > 0$. The constant H is given by

$$H = \frac{3 - \min(\alpha_{on}, \alpha_{off})}{2}. \tag{3.4}$$

To grasp the importance of (3.3), suppose momentarily that this relation holds. Then, with M independent and identically distributed sources, the aggregate load at time t will be $\sum_{m=1}^{M} W^{(m)}(t)$. Now consider the renormalized aggregate load $M^{-1/2} \sum_{m=1}^{M} (W^{(m)}(t) - EW^{(m)}(t))$. As $M \to \infty$, it satisfies

$$\mathcal{L} \lim_{M\to\infty} M^{-1/2} \sum_{m=1}^{M} (W^{(m)}(t) - EW^{(m)}(t)) = G(t), \quad t \geq 0,$$

by the usual Central Limit Theorem ($\mathcal{L}$ denotes convergence of the finite dimensional distributions). The process $\{G(t),\ t \geq 0\}$ is Gaussian and stationary (since the $W_m(t)$'s are stationary) and has covariance function $\{\gamma(t),\ t \geq 0\}$. If we also aggregate and renormalize time, considering for each $T > 0$ the process $(T^{2H} L(T))^{-1/2} \int_0^{Tt} G(u)du,\ t \geq 0$, we get, after applying (3.3),

$$\mathcal{L} \lim_{T\to\infty} T^{-H} \int_0^{Tt} G(u)du = \sigma B_H(t),\ t \geq 0. \tag{3.5}$$

The process $\{B_H(t),\ t \geq 0\}$ which now appears in the limit is fractional Brownian motion, a Gaussian process with stationary increments, mean 0 and variance t^{2H}. Indeed, the limit in (3.5) must be Gaussian with mean zero and have stationary increments since the integral of G has these properties. Moreover, by (3.3), its variance must be $\sigma^2 t^{2H}$ for fixed t. Since these properties characterize fractional Brownian motion (Samorodnitsky and Taqqu [80], Corollary 7.2.3, p. 320), Relation (3.5) follows, and hence we have

$$\mathcal{L} \lim_{T\to\infty} \mathcal{L} \lim_{M\to\infty} \frac{1}{T^H} \frac{1}{\sqrt{M}} \int_0^{Tt} \sum_{m=1}^{M} (W^{(m)}(u) - EW^{(m)}(u))du = \sigma B_H(t),\ t \geq 0. \tag{3.6}$$

This result states that when properly normalized, the aggregated total load converges to fractional Brownian motion as we first let the number of sources M go to infinity and then consider the aggregate load over larger and larger time intervals (i.e., let the time scale factor T tend to infinity); for details see [85].

To obtain (3.6), it was crucial to first let $M \to \infty$ and then $T \to \infty$. In fact, if we reverse the order and consider the limit where first $T \to \infty$ and then $M \to \infty$, after changing the normalization accordingly, one can show [85] that in this case,

$$\mathcal{L} \lim_{M\to\infty} \mathcal{L} \lim_{T\to\infty} \frac{1}{M^{1/\alpha}} \frac{1}{T^{1/\alpha}} \int_0^{Tt} \sum_{m=1}^{M} (W^{(m)}(u) - EW^{(m)}(u))du = cS_\alpha(t), \ t \geq 0, \tag{3.7}$$

where

$$\alpha = \min(\alpha_{on}, \alpha_{off})$$

and $S_\alpha(t)$ is a, generally skewed, Lévy stable motion, a process with independent and stationary increments but with infinite variance. For example, if $\alpha_{on} < \alpha_{off}$, the skewness parameter of S_α equals 1, that is, the distribution of $S_\alpha(t)$ is totally skewed to the right. Expressed in terms of the aggregated total load $A^{(M)}(t) = \int_0^t \sum_{m=1}^M W^{(m)}(u)du$, relations (3.6) and (3.7) become

$$\mathcal{L} \lim_{T\to\infty} \mathcal{L} \lim_{M\to\infty} \frac{A^{(M)}(Tt) - EA^{(M)}(Tt)}{T^H \sqrt{M}} = \sigma B_H(t), \ t \geq 0,$$

and

$$\mathcal{L} \lim_{M\to\infty} \mathcal{L} \lim_{T\to\infty} \frac{A^{(M)}(Tt) - EA^{(M)}(Tt)}{T^{1/\alpha} M^{1/\alpha}} = cS_\alpha(t), \ t \geq 0.$$

What happens if T and M go together to infinity? Mikosch, Resnick, Rootzén and Stegeman [61] answer this question by considering an integer-valued function $M(T)$ which is non-decreasing in T and which tends to infinity as $T \to \infty$. They show that under the fast growth condition

$$\lim_{T\to\infty} \frac{M(T)}{T^{\alpha-1}} = \infty, \tag{3.8}$$

the limit is the fractional Brownian motion $B_H(t)$ and under the slow growth condition

$$\lim_{T\to\infty} \frac{M(T)}{T^{\alpha-1}} = 0, \tag{3.9}$$

the limit is the Lévy stable motion $S_\alpha(t)$ of (3.7).

Finally, instead of on-off or 0/1-valued alternating renewal processes, one can also consider renewal processes where the rewards are chosen from a distribution that has infinite variance. In this case, the limit depends also on the relative magnitude of $1 < \alpha < 2$ ("heavy tailed" exponent of the renewal interval) and $0 < \beta < 2$ ("heavy-tailed" exponent of the reward). The limit is the Lévy stable motion if $\beta < \alpha$ in the case $T \to \infty$ first and then $M \to \infty$ or in the case where the order of M and T is reversed, and also in the case $\beta > \alpha$ when $T \to \infty$ first, followed by $M \to \infty$ (Levy

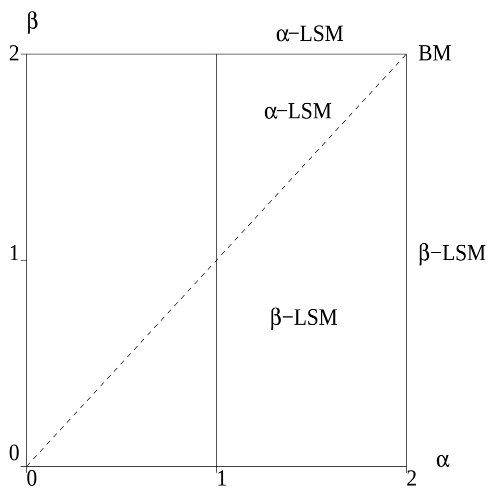

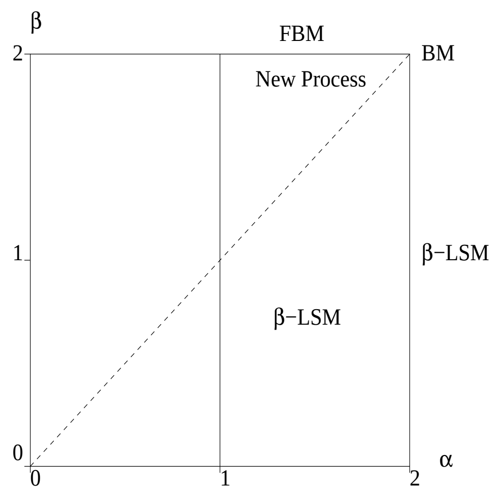

$\lim_{M\to\infty} \lim_{T\to\infty}$		$\lim_{T\to\infty} \lim_{M\to\infty}$	
$\alpha = 2$		$\alpha = 2$	
$\beta = \alpha = 2$:	BM	$\beta = \alpha = 2$:	BM
$0 < \beta < \alpha = 2$:	β-LSM	$0 < \beta < \alpha = 2$:	β-LSM
$1 < \alpha < 2$		$1 < \alpha < 2$	
$1 < \alpha < \beta = 2$:	α-LSM	$1 < \alpha < \beta = 2$:	FBM
		$1 < \alpha < \beta < 2$:	Z_β
$0 < \beta < \alpha < 2$:	β-LSM	$0 < \beta < \alpha < 2$:	β-LSM

TABLE 1.1. The table indicates the possible limits depending on the respective value of α, the interrenewal exponent and β, the reward exponent. The possible limits are BM = Brownian motion, FBM = fractional Brownian motion, LSM = Lévy stable motion, Z_β = the new self-similar stable process with index β and dependent increments.

and Taqqu [45]). However, the interesting case is $1 < \alpha < \beta < 2$, with $M \to \infty$ first followed by $T \to \infty$. This is the case which gives fractional Brownian motion when $\beta = 2$ (finite variance rewards). The limit turns out to be β-stable, self-similar with stationary increments (Levy and Taqqu [48]). But it is not fractional stable motion, which is the commonly used extension of fractional Brownian motion to the stable domain. The limit is a "mixed moving average" type process (for details, see Pipiras and Taqqu [69]). These results are summarized in Table 1.1.

Pipiras, Taqqu, and Levy [70] describe what happens in this situation when M and T go to infinity together. Interestingly, the fast growth and slow growth conditions are exactly the same as in (3.8) and (3.9) respectively, and therefore they do not depend on β. Under the fast growth condition (3.8), one obtains in the limit the β-stable, self-similar process with stationary increments mentioned above, and under the slow growth condition (3.9), one obtains a *symmetric* Lévy stable motion $S_\alpha(t)$. See Taqqu [81] for an overview.

The use of on-off models (or more generally, of renewal reward processes) to generate fractional Brownian motion and/or symmetric Lévy stable motion was originally proposed by Mandelbrot [49] in an economic context (see Taqqu and Levy [82] for a description). Resnick and Samorodnitsky [71] provide an overview of different approaches to the subject.

3.3 Mathematical framework II: The infinite source Poisson model

An alternative to the superposition of on-off sources is the *infinite source Poisson model*. In this model, sources arrive at the link at a rate of λ, each source transmits for a random time with distribution $\overline{F}_{on}(x) \sim lx^{-\alpha}$, $x \to \infty$, $1 < \alpha < 2$, and upon completing transmission, the sources remain inactive. The variance of the transmission times is infinite but their mean μ_{on} is finite. The transmission times are assumed i.i.d. This model provides another appealing abstraction of network traffic at the application layer in the sense that it views traffic as being generated by "sessions" which arrive in a Poisson fashion and whose sizes (in bytes or packets) or durations (in time) are heavy-tailed with infinite variance. In fact, the infinite source Poisson model is closely related to the on-off source models considered in Section 3.2 above— it is λ, here, which plays the role of M.

Let $W(u)$ denote the number of active sources in the infinite source Poisson model at time $u > 0$ and let $A(t) = \int_0^t W(u)du$ be the aggregated load from time 0 to time $t \geq 0$. The behavior of the system depends on the relative size of the arrival rate λ and the time horizon T. Mikosch, Resnick, Rootzén and Stegeman [61] show that under the fast growth condition

$$\lim_{T\to\infty} \frac{\lambda(T)}{T^{\alpha-1}} = \infty,$$

one has

$$\mathcal{L}\lim_{T\to\infty}\frac{A(Tt)-EA(Tt)}{\lambda^{1/2}T^{(3-\alpha)/2}}=c_1B_H(t),\ t\geq 0,$$

and under the slow growth condition

$$\lim_{T\to\infty}\frac{\lambda(T)}{T^{\alpha-1}}=0,$$

one has

$$\mathcal{L}\lim_{T\to\infty}\frac{A(Tt)-EA(Tt)}{\lambda^{1/\alpha}T^{1/\alpha}}=c_2S_\alpha(t),\ t\geq 0.$$

Observe that one has $EA(t)=\mu_{on}\lambda t$.

The infinite source Poisson model was originally considered by Cox [12] and is also known as $M/G/\infty$ *model* or *immigration-death process*. As a traffic workload model, the convergence result that yields fractional Brownian motion in the limit was originally established by Kurtz [44].

3.4 Empirical validation: Source-level traffic characteristics

The structural modeling approaches involving the infinite source Poisson model and the different kinds of aggregations of on-off sources provide a networking-related understanding of the self-similarity phenomenon that is observed in measured data traffic, by relating it to more elementary properties of the traffic patterns generated by individual users and/or applications. We can now explain why data traffic exhibits self-similar scaling properties over large enough time scales. Part of the appeal of this new understanding is that the corresponding mathematical arguments are rigorous, simple, and closely related to familiar concepts (even though the resulting mathematical objects defy tradition and fall outside the commonly dealt-with Markovian framework). They are in agreement with the networking researchers' intuition, and can readily be explained to non-networking experts and non-mathematicians. These developments have helped immensely in demystifying self-similar traffic modeling. They have given rise to a physical understanding of the effects of LRD on the design, management and performance of modern data networks.

More importantly, these structural models involve elementary quantities (e.g., arrivals and durations for sessions, TCP connections, IP flows, or on/off-periods) that *(i)* correspond to different networking mechanisms, and *(ii)* can be measured. One can check, in particular, whether the statistical characteristics of these elementary quantities are in agreement with the mathematical assumptions made in Sections 3.2 and 3.3. For example, Willinger et al. [90, 91] revisited the data from [47] and performed an extensive analysis of the data at the level of individual source-destination pairs. The results confirmed that the measured traffic at that level was consistent with the assumed on-off behavior. In particularly, they found that the durations of the on- and/or off-periods conformed to heavy-tailed distributions with finite mean and infinite variance; i.e., $1 < \alpha_{on}, \alpha_{off} < 2$. In addition, an intuitive argument was given that relates the heavy-tailedness of the on-periods to the empirically observed heavy-tailedness of file sizes on file servers. Similar results were reported in [14] and used to attribute the observed self-similar scaling behavior of aggregate Web traffic to the fact that the measured traffic is the superposition of heavy-tailed on-off sources, where the on-periods correspond to the transmission durations of individual Web documents (which have been observed to be consistent with heavy-tailed behavior) and the off-periods correspond to times when no data are transmitted.

In contrast to on-off source models and their variants, the infinite source Poisson models tend to be especially appropriate when one can reconstruct from packet-level measurements network activity at the application level in the form of session arrivals and session sizes (number of packets or bytes) or durations (e.g., in seconds). Because of how the different applications are structured, determining session entities such as arrival times and sizes is straightforward for FTP and TELNET, but is in general not possible in the case of the HTTP (i.e., Web application). For Internet traffic, Paxson and Floyd in [65] showed that the arrivals of both FTP and TELNET sessions are consistent with (nonhomogeneous) Poisson processes, with rates that are constant over about an hour, and that the distributions of measured FTP session sizes or durations are heavy-tailed, with upper tails that are consistent with finite mean but infinite variance. As far as measured HTTP sessions are concerned, indirect evidence in favor of their heavy-tailed size or durations is given in [89], where HTTP connections are shown to be consistent with heavy-tailedness with $1 < \alpha < 2$. (An HTTP session is typically made up of many HTTP connections.) For more direct empirical evidence for the appropriateness of the infinite source Poisson model for HTTP traffic, we refer to a study by Feldmann et al. [25]. These authors associate HTTP sessions within a commercial Internet Service Provider (ISP) with individual modem calls and, by matching the appropriate packets with appropriate data collected about each modem call, they are able to validate that session arrivals are consistent with an (inhomogeneous) Poisson process, and that session sizes or durations conform to the desired heavy-tailedness property.

3.5 When there is more to data network traffic than LRD ...

That we can explain the empirically observed LRD of aggregate data traffic in terms of the statistical properties of the individual connections that make up the overall traffic rate process suggests that the LRD nature of data traffic is mainly caused by user/application characteristics (i.e., Poisson arrivals of sessions, heavy-tailed distributions with infinite variance for the session sizes/durations) and has little to do with the network (i.e., with 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 LRD or asymptotic self-similarity property of aggregate data traffic over large time scales to hold, all that is needed is that the number of packets or bytes per individual connection be 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. The network is simply here to transmit data, and how the network accomplishes the data transfer has—to a first order—no significant impact on the apparent characteristics of data traffic viewed over large time scales.

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 LRD property. It triggered the curiosity of networking researchers who wanted to know "Why LRD or self-similar?" In turn, this searching for a physical explanation resulted in findings about data traffic at the connection level. In this sense, the progression of results proceeded in an opposite way of how traffic modeling has been traditionally done in this area; namely, by first analyzing in great detail the dynamics of packet flows within individual connections and then by 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 finding of LRD in data traffic traces has demonstrated that new insights about the nature of actual data traffic can be gained by first performing a careful statistical analysis of measured traffic at the aggregate level, and then by explaining the aggregate traffic characteristics in terms of more elementary properties that are exhibited by measured data traffic at the connection-level.

A natural next step is the realization that measured data traffic exhibits further structure when analyzing it over small time scales. From a networking perspective, this observation translates into "If LRD is mainly a property of user/application behaviors, where is the network?" Phrased even more succinctly, the questions are "What is large-scale? What is small-scale? Where is the transition and what is its relevance for networking?" In fact, since networking algorithms such as protocol-specific rules typically operate on small time scales and largely determine the actual flow of traffic across the network, they can be expected to create pronounced local variations in the small-scale features of data traffic, and render these small-scale features different from the observed large-time (i.e., LRD) behavior. Studies of the fine-time scale structure of data traffic require a more detailed knowledge of the dominant protocols, or at least of their most pronounced features. To this end, we focus in Section 4.1 below on the most commonly used transport layer protocol, TCP, and some of its critical functionalities.

4 LRD and small-time scaling: A network-layer view

Before discussing evidence that suggests the presence of additional structure when studying data traffic over small time scales, we first discuss (at a relatively high level) TCP, one of the most important transport protocols. We illustrate how its rules and mechanisms for flow, error and congestion control are capable of introducing traffic patterns, usually localized in time, that range from very regular to highly irregular. Recalling that LRD leaves the small-time scale behavior (or, equivalently, the small-lag correlation structure or high-frequency component of the spectrum) essentially unspecified, we introduce, in this section, the mathematical concept of *multifractals* that may be relevant in some situations for dealing with the observed irregular behavior of measured data traffic over small scales.

4.1 Networking 101: A TCP primer

As mentioned above, when more packets arrive at a router than it can forward along its links, these packets are queued in memory buffers. If the excess rate is sustained (*congestion*), the buffers will be ultimately exhausted and some packets will have to be dropped from the buffer. Consequently, "socially responsible" protocols used for exchanging data across the Internet include various types of *end-to-end congestion control* mechanisms that automatically decrease the rate at which data is transmitted when congestion is detected. These introduce significant, complicated correlations across time, among active connections, and between the different layers in the protocol hierarchy. Many models of aggregate data traffic on a network link ignore the fact that the link has finite capacity. Yet it is precisely this finite capacity that drives the dynamics of protocols such as TCP and couples the different simultaneous connections sharing the link in intricate and complex ways.

As a specific illustration, consider TCP, the predominant transport layer protocol in the Internet. The service provided by TCP is to deliver a stream of data to a receiver such that the entire stream arrives in the same order, with no duplicates, and reliably even in the presence of packet loss, reordering, duplication, and rerouting. TCP splits the data into *segments*, with one segment transmitted in each packet. The receiver acknowledges the receipt of segments if they are "in order" (it has already received all data earlier in the stream). Each acknowledgment (ACK) also implicitly acknowledges all of the earlier-in-the-stream segments, so the loss of a single ACK is rarely a problem; a later ACK will cover for it, as far as the sender is concerned.

The sender runs a timer so that if it has not received an ACK from the receiver for data previously sent when the timer expires, the sender will conclude that the data (or all of the subsequent ACKs) was lost and retransmit the segment. In addition, whenever a receiver receives a segment that is out of order (does not correspond to the next position in the data stream), it generates a "duplicate ACK," that is, another copy of the same ACK as it sent for the last in order packet it received. If the sender observes the arrival of three such duplicate ACKs, then it concludes that a segment must have been lost (leading to a number of out-of-order segments arriving at the receiver, hence the duplicate ACKs), and retransmits it *without* waiting first for the timer to expire.

Imagine, for example, that segments 1 and 2 are sent, that 1 is received and 2 is lost. The receiver sends the acknowledgment ACK(1) for segment 1. As soon as ACK(1) is received, the sender sends segments 3, 4, and 5. If these are successfully received, they are retained by the receiver, even though segment 2 is missing. But because they are out of order, the receiver sends back three ACK(1)'s (rather than ACK(3), ACK(4), ACK(5)). From the arrival of these duplicates, the sender infers that segment 2 was lost (since the ACKs are all for segment 1) and retransmits it.

A key requirement for attaining good performance over a network path is that the sender must in general maintain several segments "in flight" at the same time, rather than just sending one and waiting an entire round trip time (RTT) for the receiver to acknowledge it. However, if the sender has too many segments in flight, then it might overwhelm the receiver's ability to store them (if, say, the first is lost but the others arrive, so the receiver cannot immediately process them), or the network's available capacity.

The first of these considerations is referred to as *flow control*, and in TCP is managed by the receiver sending an *advertised window* informing the sender how many data segments it can have in flight beyond the latest one acknowledged by the receiver. This mechanism is termed a "sliding window," since each ACK of new data advances a window bracketing the range of data the sender is now allowed to transmit. From the perspective of network dynamics, a very important property of a sliding window protocol is that it leads to *self-clocking*. That is, no matter how fast the sender transmits, its data packets will upon arrival at the receiver be spaced out by the network to reflect the network's current carrying capacity; the ACKs returned by the receiver will preserve this spacing; and consequently the window at the sender will advance in a pattern that mirrors the spacing with which the previous flight of data packets arrived at the receiver, which in turn matches the network's current carry capacity. Thus, TCP has intrinsic to it a mechanism that generates structure on the time scale of a RTT.

In addition, TCP maintains a *congestion window*, or CWND, that controls how the sender attempts to consume the path's capacity. At any given time, the sender confines its data in flight to the lesser of the advertised window and CWND. Each received ACK, unless it is a duplicate ACK, is used as an indication that data has been transmitted successfully, and allows TCP to increase CWND. At startup, CWND is set to 1 and the *slow start* mechanism takes place, where CWND is increased by one segment for each arriving ACK. The more segments that are sent, the more ACKs are received, leading to exponential growth. (The *slow start* procedure is "slow" compared to the old mechanism which consisted of immediately sending as many packets as the advertised window allowed.) If TCP detects a packet loss, either via duplicate ACKs or via timeout, it sets a variable called the slow start threshold, or SSTHRESH to half of the present value of CWND. If the loss was detected via duplicate ACKs, then TCP does not need to cut back its rate drastically: CWND is set to SSTHRESH and TCP enters the *congestion avoidance* state, where CWND is increased linearly, by one segment per RTT. If the loss was detected via a timeout, then the self-clocking pattern has been lost, and TCP sets CWND to one, returning to the slow start regime in order to rapidly start the clock going again. When CWND reaches the value of SSTHRESH, *congestion avoidance* starts and the exponential increase of CWND shifts to a linear increase. (For more details, see e.g. Peterson and Davie [68]).

The key point concerning TCP's management of the congestion window is that it leads to TCP traffic *dynamically adapting* to changing network conditions, and it does so on time scales of a few RTTs. Accordingly, TCP traffic is *shaped* by current conditions; it cannot be played back in some other context (for example, trace-driven simulation) without great care in understanding how the traffic would have adapted differently in that different context.

4.2 Multifractals as a mathematical framework

We have seen in Section 3 that user behavior and/or application-layer characteristics are primarily responsible for the observed LRD property of network traffic and hence for the self-similar scaling behavior of aggregate network traffic over large time scales. Specifically, we have seen that the fluctuations of measured traffic rate processes around their mean rate (when measured over sufficiently large time scales) tend to be consistent with fractional Gaussian noise. The latter provides a complete statistical description of the traffic as seen on a single link and when aggregated over all users/connections and over sufficiently large time periods. However, by their very definition, asymptotic properties (in the limit as we look over larger and larger time scales) such as LRD or, equivalently, asymptotic second-order self-similarity, leave the small-time scaling behavior of the measured traffic essentially unspecified and allow for different possible descriptions of the fine structure of the observed traffic. Thus, in order to develop models of measured network traffic capable of accounting for small time scales, we need a mathematical framework that extends beyond fractional Gaussian noise, one that can handle highly irregular or regular behavior well-localized in time. This leads us to the study of the differentiability, fractality and regularity of stochastic processes; that is, properties that are inherently related to the high-frequency behavior or small-lag correlations and have little to do with low-frequency components and/or high-lag correlations, except that the latter dominate as soon as smoothing or aggregating in time has successfully wiped out any interesting small-time peculiarities, leaving us with LRD or its close relative, i.e., self-similarity, as predominant and universal characteristics.

Local irregularity/regularity: Using first the more traditional time domain-based approach to quantifying locally irregular behavior in a measured signal at a particular point in time t_0, we follow closely [78] and 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. 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. $\alpha(t_0)$ is called the *singularity* or *Hölder exponent* at time t_0 and can be shown—passing to Fourier space—to be a generalization of the degree of differentiability. If Y has positive increments, this *singularity exponent* can be approximated through the somewhat simpler quantity

$$\alpha(t_0) = \lim_{n \to \infty} \alpha_n(t_0), \tag{4.1}$$

where—assuming the limit exists—for $t_0 \in [k_n 2^{-n}, (k_n + 1)2^{-n})$,

$$\alpha_n(t_0) := \alpha^n_{k_n} := -\frac{1}{n} \log_2 \left| Y((k_n + 1)2^{-n}) - Y(k_n 2^{-n}) \right|. \tag{4.2}$$

A second approach for dealing with locally highly irregular/regular signals is wavelet domain- rather than time domain-based. The wavelet-based method exploits the fact that the wavelet decomposition of a given signal contains information about its locally irregular behavior. In fact, the singularity or Hölder exponent $\alpha(t_0)$ 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_0, where ψ is a bandpass wavelet function and where $\psi_{j,k}(s) := 2^{-j/2}\ \psi(2^{-j}s - k)$ (e.g., in the case of the well-known *Haar wavelet*, $\psi(s)$ equals 1 for $0 \le s \le 1$, -1 for $1 \le s \le 2$, and 0 for all other s; for a general overview of wavelets, we refer to [17]). For example, assuming only that $\int \psi(s)ds = 0$, one can show as in [40] that

$$2^{n/2} w_{-n,k_n} \le C \cdot 2^{-n\alpha(t_0)}, \qquad \text{as } k_n 2^{-n} \to t_0. \tag{4.3}$$

Moreover, it is known that under some regularity conditions (for a precise statement see [40] or [17, Theorem 9.2]), relation (4.3) characterizes the degree of local irregularity of the signal at the point t_0. This suggests to define a new exponent $\tilde{\alpha}(t_0) = \lim_{n\to\infty} \tilde{\alpha}_n(t_0)$, where

$$\tilde{\alpha}_n(t_0) := \tilde{\alpha}^n_{k_n} := \frac{1}{-n\log 2}\log\left(2^{n/2}|w_{-n,k_n}|\right). \tag{4.4}$$

While this wavelet-based approach may give rise to a different description of the singularity structure of Y than the earlier-described time domain method, particularly for non-monotone processes (for an example, see [30]), we will discuss later in this section some definitive advantages that the former seems to have over the latter when analyzing measured TCP/IP traces.

Multifractal analysis (MFA): What follows is an intuitive way of understanding multifractals. (For a mathematical account of the subject, see [74].) The aim of *multifractal analysis* (MFA) is to provide information about the singularity exponents in a given signal (as defined in equations (4.2) and (4.4), resp.) and to come up with compact geometrical or statistical descriptions of the signal's overall singularity.

Conceptually, the time domain is the most obvious one for a geometrical formulation of MFA. 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

$$K_\alpha = \{t\,:\,\alpha(t) = \alpha\}. \tag{4.5}$$

To illustrate, tale Y to be FBM. Then 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 FBM is therefore said to be "mono-fractal." Similarly, for the concatenation of several FBMs 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, because its definition involves taking a double limit (see e.g. Falconer [23]). This drawback severely limits the usefulness of the geometrical approach to MFA and suggests considering instead statistical descriptions of the multifractal structure of a given signal.

One such description involves the notion of the *coarse Hölder exponent* (4.2). To illustrate, we fix a path of Y and consider a histogram of the α_k^n's ($k = 0, \dots 2^n - 1$) 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 apply the LDP approach to our situation, consider the fixed path of Y at the location t, where the encoding of t by k_n via $t \in [k_n 2^{-n}, (k_n+1)2^{-n})$ represents the only randomness relevant for applying 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 random variables of interest for our purpose is

$$V_n := -\log_2 |Y((k_n+1)2^{-n}) - Y(k_n 2^{-n})| = n\alpha_{k_n}^n. \tag{4.6}$$

Trying to obtain more precise information about the singularity behavior, we define the following limiting "rate function" f which will exist under mild conditions [73, Theorem 7]:

$$f(\alpha) := \lim_{\epsilon \to 0} \lim_{n \to \infty} \frac{1}{n} \log_2 f_n(\alpha, \epsilon), \tag{4.7}$$

with

$$f_n(\alpha, \epsilon) := 2^n P_t \Big[\alpha + \epsilon \geq \alpha_n(t) \geq \alpha - \epsilon \Big] = \#\{\alpha_n(t) \in (\alpha - \epsilon, \alpha + \epsilon)\}. \tag{4.8}$$

The counting in (4.8) 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., $f_n(\alpha, \epsilon) \simeq 2^n$. Such is the case for FBM 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 FBMs described earlier [76]. 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 [77, 74] that the rate function $f(\alpha)$ relates to the Hausdorff dimension $\dim(K_\alpha)$ and that we have

$$\dim(K_\alpha) \leq f(\alpha). \tag{4.9}$$

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.

Multifractal formalism: 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 (4.7)). Fortunately, the LDP-result suggests an alternative method for estimating f that avoids double-limit operations and is generally more robust because it involves averages. In fact, in an attempt to investigate the scaling behavior of the higher-order moments of the traffic rate process, consider the *partition function* $\tau(q)$ defined by

$$\tau(q) := \lim_{n\to\infty} \frac{-1}{n} \log_2 \left(2^n E_t[2^{qV_n}]\right) = \lim_{n\to\infty} \frac{-1}{n} \log_2 S_n(q), \tag{4.10}$$

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

$$S_n(q) := \sum_{k=0}^{2^n-1} |Y((k+1)2^{-n}) - Y(k2^{-n})|^q = \sum_{k=0}^{2^n-1} 2^{-qn\alpha_k^n} \tag{4.11}$$

and note that $\tau(q)$ and $S_n(q)$ are related via the scaling relation

$$S_n(q) \sim 2^{-n\tau(q)}. \tag{4.12}$$

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

$$S_n(q) = \sum_{\alpha} \sum_{\alpha_n \simeq \alpha} 2^{-nq\alpha} \simeq \sum_{\alpha} 2^{-n(q\alpha - f(\alpha))} \simeq 2^{-n \inf_\alpha (q\alpha - f(\alpha))}; \tag{4.13}$$

that is,

$$\tau(q) = f^*(\alpha) := \inf_{\alpha}(q\alpha - f(\alpha)), \tag{4.14}$$

where * denotes the Legendre transform of a function (for a mathematically rigorous argument, see [73, 74]).

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

$$f(\alpha) \leq f^{**}(\alpha) = \tau^*(\alpha) = \inf_{q}(q\alpha - \tau(q)) \tag{4.15}$$

where f^{**} is the concave hull of f. Conditions under which equality holds in (4.15) are of particular interest and have been studied extensively. For example, if the partition function $\tau(q)$ is everywhere differentiable, then using an application of the Ellis-Gärtner theorem of LDP [20] (or similar methods [73, 74]) the following explicit formula for the Legendre transform can be obtained:

$$f(\alpha) = q\alpha - \tau(q) \qquad \text{at } q \text{ such that } \alpha = \tau'(q). \tag{4.16}$$

A different case of interest in which we have the equality $f(\alpha) = \tau^*(\alpha)$ concerns the class of processes called *multiplicative cascades* and their variations [52, 42, 11, 5, 22, 62, 36, 67, 75]. This mathematically well-studied class of signals provides a powerful link between the coarse approximation of local regularity provided by the rate function $f(\alpha)$, the global scaling captured in the partition function $\tau(q)$, and the pointwise Hölder exponent α_t: the Hausdorff-dimension [24] of the set of time instances t for which α_t equals a given value α is precisely equal to $f(\alpha) = \tau^*(\alpha)$ for almost every path of the process [42, 5, 22, 62, 36, 75] (see also [74] and [41, 21, 77, 73]). These results constitute what is called the *multifractal formalism.* The term alludes to the "thermodynamical formalism" and dates back to the discovery of multifractals in the context of turbulence [50, 29, 32, 35, 33, 10, 41, 21, 75, 77, 73]). In view of this formalism, numerically established equality in (4.15) is generally taken as an indication of the presence of strong multiscale structure in the signal and as justification for using f as an appropriate measure for describing and capturing local irregularity. In particular, genuinely concave behavior of $\tau(q)$ is viewed as empirical evidence that there is a whole interval of α-values present in the signal and not just a few, hence the term *multifractal.*

While in this section the time domain-based approach was used to discuss most of the concepts related to MFA, we note that the basic conclusions remain true if $\alpha(t)$ is replaced by $\tilde{\alpha}(t)$ and (4.2) by (4.4); that is, the time domain-based method involving the increments is replaced by the wavelet domain-based technique that focuses on the signal's wavelet coefficients; for a more detailed description of the wavelet domain-based approach and its properties, see for example [6, 30, 75, 2].

4.3 Multifractals and network traffic: Empirical observations

Taqqu, Teverovsky and Willinger [84] discuss the relevance of monofractals and multifractals for describing LAN and WAN traffic. The first full-blown multifractal analysis of measured aggregate TCP/IP WAN traffic traces, however, is due to Riedi and Lévy Véhel [76]. Performing a MFA using the time domain-based approach, they report on multifractal scaling in unprecedented high quality and over essentially the whole range of scales. In fact, their plots of $\log S(n, q)$ against n (the logarithmic time scale) show nearly perfect linear behavior, allowing for reliable estimates of the corresponding slopes; i.e., the partition function $\tau(q)$ defined by (4.10). In addition, the log-normalized histograms of logarithmically spaced $\alpha^n_{n_k}$-values (see (4.2) above) collapse nicely for a set of large n-values and approach the Legendre transform $\tau^*(\alpha)$ of $\tau(q)$, suggesting an apparent scale-invariant behavior of the measured traffic over small time scales. In view of the multifractal formalism discusses earlier in this section, these observations suggest an underlying $\tau(q)$ function that is differentiable, adding additional regularity to the scaling. Using the same methodology as employed by Riedi and Lévy Véhel, Mannersalo and Norros [60] analyzed various short traces of measured aggregate WAN ATM traffic and basically confirmed the results reported in [76]. (ATM is a networking technology often used in IP networks as a substrate for an IP-layer, i.e., internetworking layer, hop.)

From a networking perspective, the apparent insensitivity of Riedi and Lévy Véhel's techniques to aspects of the measured traffic that one encounters (e.g., the pronounced self-clocking behavior of TCP on the order of RTT-sized time scales discussed in Section 4.1) is problematic. This insensitivity casts doubts on the appropriateness of applying a time domain-based technique when the traces tend to exhibit—in addition to possibly multifractal scaling—relatively regular patterns when observed over RTT-sized time scales. Clearly, such regular behavior can negatively affect any sort of scaling investigations, either in the small or large scales. Alternative approaches to the time domain-based method introduced in [76] for inferring multifractal scaling are needed—approaches that are capable of relating network-specific features with dominant aspects of the resulting trace analysis.

One such approach is wavelet domain-based, and exploits the fact that irregular/regular patterns in a given signal can be well localized (in time and space) and studied by considering the wavelet coefficients associated with the signal's wavelet decomposition. This wavelet-based approach to investigating the possibility of multifractal scaling in network traffic was originally proposed by Feldmann et al. [26, 27] in an effort to validate the findings reported in [76] using different data sets as well as a different technique than the unproven (in the network setting) time domain-based approach pursued in [76]. Wavelets were first brought to the attention of the networking community by Abry, Veitch and their co-workers, who demonstrated the wavelets' natural abilities for investigating scaling-related phenomena in the context of analyzing, estimating, and synthesizing LRD processes [1, 2, 3]. Influenced by the work of Arneodo and his collaborators [6, 7], Feldmann et al. proposed in [26, 27] an application of wavelet-based techniques for inferring scaling phenomena such as LRD and multifractal scaling in a manner that is more qualitative in nature and focuses less on a quantitative assessment of the scaling behaviors for which these techniques were originally designed for. In [26] and, particularly, in [27], Feldmann et al. demonstrate that one can associate known network-specific features with qualitative characteristics (e.g., scaling regions, "breakpoints", "spikes" or "dips") in the output of the wavelet scaling analysis of the data. For example, recalling the definition of a signal's wavelet coefficients given earlier in Section 4.2 (using, say, the Haar wavelet), it is easy to see that a relatively regular pattern that is well localized in time and scale will give rise to corresponding small wavelet coefficients; in contrast, well-localized irregular patterns can be expected to yield large wavelet coefficients in the corresponding time-scale or wavelet-domain "neighborhood" of the underlying signal. To illustrate, consider a sinusoidal signal and its corresponding wavelet coefficients using Haar wavelets. Clearly, for suitably chosen scales, the signal and the wavelets will "cancel" each other, resulting in wavelet coefficients that are negligible.

The fact that it is possible to relate wavelet domain-based features to network characteristics is very useful and promising. For example, Feldmann et al. [27] comment on (i) how RTT (the round-trip time) behavior impacts both large-time and small-time scaling behavior, (ii) when a small-time scaling analysis does and does not make sense, depending on the conditions under which the measured data were collected, and (iii) under what situations measured network traffic can be expected to conform to multifractal-like scaling behavior over small-time scales. In short, the findings in [27] suggest that highly accurate TCP traces (e.g., time stamp accuracy in the microsecond range) collected, if possible, from networking environments that favor large RTTs (e.g., on the order of a few hundreds of milliseconds) are a prerequisite for performing any reasonable scaling-type analysis over a sufficiently wide range of small time scales. Under current networking conditions, the TCP layer appears to be the most promising place within the layered networking hierarchy for trying to study the irregularities or regularities of network traffic and for attempting to relate them to network-specific features.

5 Outlook–There is more to network traffic than multifractals!

Starting with the discovery of self-similar scaling in measured traffic traces from modern-day data networks, we have argued that to encounter truly interesting and challenging problems related to the mathematical modeling of empirically observed phenomena, it is necessary to move beyond the traditional descriptive or "model-fitting" stage and aim instead for physical explanations of the phenomena that can be validated against actual data. Despite progress in relating certain network-specific features to the output of the wavelet domain-based scaling analysis, nothing close-to a physical explanation of the observed small-time scale behavior (including multifractal scaling) exists to date. Yet, for any type of observed scaling behavior to have a genuine impact on networking, it is essential that their applications move beyond the traditional descriptive stage and yield answers as to why network traffic exhibits a certain scaling property. Since the available data provide detailed information about so many different facets of network behavior, there exists great potential for coming up with intuitively appealing, conceptually simple and mathematically rigorous statements as to the causes the various scaling phenomena that are observed in data networks, on both small and large scales.

List of Abbreviations

ACK = acknowledgment packet
ATM = asynchronous transfer mode (sometimes used in high-speed wide-area networks)
CWND = control window (used in TCP)
FBM = fractional Brownian motion
IP = Internet protocol
LAN = local area network (for example, in a building)
LDP = large deviation principle
LRD = long-range dependence
MFA = multifractal analysis
PSTN = public switched telephone network
QoS = quality of service
RTT = round trip time
SRD = short-range dependence
SSTHRESH = slow start threshold (used in TCP)
TCP = transmission control protocol
WAN = wide area network (for example, the Internet)

Acknowledgments

W. Willinger and M. Taqqu were supported in part by the NSF Grant ANI-9805623 at Boston University. R. Riedi was supported by the NSF under Grants ANI-9979465 and CCR-9973188.

References

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

[2] P. Abry, P. Flandrin, M. S. Taqqu and D. Veitch. Wavelets for the analysis, estimation and synthesis of scaling data. In K. Park and W. Will-

inger, editors, *Self-Similar Network Traffic Analysis and Performance Evaluation*, Wiley, New York, 2000, pp. 39–88.

[3] P. Abry, P. Flandrin, M. S. Taqqu, and D. Veitch. Self-similarity and long-range dependence through the wavelet lens. In P. Doukhan, G. Oppenheim, and M. S. Taqqu, editors, *Long-range Dependence: Theory and Applications*. Birkhäuser, 2000. Appears in this volume.

[4] R. Adler. *The Geometry of Random Fields*. J. Wiley & Sons, New York, 1981.

[5] M. Arbeiter and N. Patzschke. Self-similar random multifractals. *Mathematische Nachrichten*, 181:5–42, 1996.

[6] 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.

[7] A. Arneodo, E. Bacry and J. F. Muzy. Random cascades on wavelet dyadic trees. *Journal of Mathematical Physics*, 39:4142–4164, 1998.

[8] E. Bacry, J. F. Muzy and A. Arneodo. Singularity spectrum of fractal signals from wavelet analysis: Exact results. *Journal of Statistical Physics*, 70:635–674, 1993.

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

[10] G. Brown, G. Michon, and J. Peyriere. On the multifractal analysis of measures. *Journal of Statistical Physics*, 66:775–790, 1992.

[11] R. Cawley and R. D. Mauldin. Multifractal decompositions of Moran fractals. *Advances in Mathematics*, 92:196–236, 1992.

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

[13] M. E. Crovella and A. Bestavros. Self-similarity in world wide web traffic—evidence and possible causes. *Proceedings ACM/SIGMETRICS 96*, Philadelphia, PA, pp. 160–169, 1996.

[14] M. E. Crovella and A. Bestavros. Self-similarity in world wide web traffic—evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5:835–846, 1997.

[15] M. E. Crovella, M. S. Taqqu, and A. Bestavros. Heavy-tailed probability distributions in the World Wide Web. In R. Adler, R. Feldman, and M. S. Taqqu, editors, *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, pages 3–25, Boston, 1998. Birkhäuser.

[16] R. Dahlhaus. Efficient parameter estimation for self similar processes. *The Annals of Statistics*, 17(4):1749–1766, 1989.

[17] I. Daubechies. *Ten Lectures on Wavelets*. SIAM, New York, 1992.

[18] J.-D. Deuschel and D. W. Stroock. *Large Deviations*, volume 137 of *Pure and applied mathematics*. Academic Press, 1984.

[19] D. E. Duffy, A. A. McIntosh, M. Rosenstein, and W. Willinger. Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks. *IEEE Journal on Selected Areas in Communications*, 12:544–552, 1994.

[20] R. Ellis. Large deviations for a general class of random vectors. *The Annals of Probability*, 12:1–12, 1984.

[21] C. J. G. Evertsz and B. B. Mandelbrot. Multifractal measures. Appendix B in *Chaos and Fractals* by H.-O. Peitgen, H. Jürgens and D. Saupe, Springer Verlag, New York, pages 849–881, 1992.

[22] K. J. Falconer. The multifractal spectrum of statistically self-similar measures. *Journal of Theoretical Probability*, 7:681–702, 1994.

[23] K. Falconer. *Fractal Geometry: Mathematical Foundations and Applications*. J. Wiley & Sons, New York, 1993.

[24] K. J. Falconer. *Techniques in Fractal Geometry*. J. Wiley and Sons, Chichester, England, 1997.

[25] A. Feldmann, A. C. Gilbert. W. Willinger and T. G. Kurtz. The changing nature of network traffic: Scaling phenomena. *Computer Communication Review*, 28:5–29, April 1998.

[26] A. Feldmann, A. C. Gilbert, and W. Willinger. Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic. *Proceedings ACM/SIGCOMM 98*, Vancouver, Canada, 28:42–55, 1998.

[27] A. Feldmann, A. C. Gilbert, P. Huang and W. Willinger. Dynamics of IP traffic: A study of the role of variability and the impact of control. *Computer Communication Review* **29**, No. 4, Proceedings ACM/SIGCOMM'99, Cambridge, MA, September 1999, pp. 301–313, 1999.

[28] R. Fox and M. S. Taqqu. Large-sample properties of parameter estimates for strongly dependent stationary Gaussian time series. *The Annals of Statistics*, 14:517–532, 1986.

[29] U. Frisch and G. Parisi. Fully developed turbulence and intermittency. *Proceedings of the International Summer School on Turbulence and Predictability in Geophysical Fluid Dynamics and Climate Dynamics*, pages 84–88, 1985.

[30] A. C. Gilbert, W. Willinger and A. Feldmann. Scaling analysis of conservative cascades, with applications to network traffic. *IEEE Transactions on Information Theory*, 45: 971–991, 1999.

[31] C. W. J. Granger and R. Joyeux. An introduction to long-memory time series and fractional differencing. *Journal of Time Series Analysis*, 1:15–30, 1980.

[32] P. Grassberger and I. Procaccia. Characterization of strange attractors. *Physical Review Letters*, 50:346–349, 1983.

[33] T. Halsey, M. Jensen, L. Kadanoff, I. Procaccia, and B. Shraiman. Fractal measures and their singularities: The characterization of strange sets. *Physical Review A*, 33:1141–1151, 1986.

[34] F. R. Hampel. Data analysis and self-similar processes. In *Proceedings of the 46th Session of the International Statistical Institute*, Tokyo, Japan, September 1987. International Statistical Institute.

[35] H. Hentschel and I. Procaccia. The infinite number of generalized dimensions of fractals and strange attractors. *Physica D*, 8:435–444, 1983.

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

[37] J. R. M. Hosking. Fractional differencing. *Biometrika*, 68(1):165–176, 1981.

[38] H. E. Hurst. Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116:770–808, 1951.

[39] S. Jaffard. Pointwise smoothness, two-microlocalization and wavelet coefficients. *Publications Matematiques*, 35:155–168, 1991.

[40] S. Jaffard. Local behavior of Riemann's function. *Contemporary Mathematics*, 189:287–307, 1995.

[41] S. Jaffard. Multifractal formalism for functions, part 1: Results valid for all functions. *SIAM Journal on Mathematical Analysis*, 28:944–970, 1997.

[42] J.-P. Kahane and J. Peyrière. Sur certaines martingales de Benoit Mandelbrot. *Advances in Mathematics*, 22:131–145, 1976.

[43] V. Klemes. The Hurst phenomenon: a puzzle? *Water Resources Research*, 10:675–688, 1974.

[44] 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.

[45] J. Levy and M. S. Taqqu. On renewal processes having stable inter-renewal intervals and stable rewards. *Les Annales des Sciences Mathematiques du Quebec*, 11:95–110, 1987.

[46] W. E. Leland, M. S. Taqqu, W. Willinger and D. V. Wilson, On the self-similar nature of Ethernet traffic. *Proceedings ACM/SIGCOMM 93*, San Francisco, CA, 28:42–55, 1993.

[47] 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*, 2:1–15, 1994.

[48] J. B. Levy and M. S. Taqqu. Renewal reward processes with heavy-tailed interrenewal times and heavy-tailed rewards. *Bernoulli*, 6(1):23–44, 2000.

[49] B. B. Mandelbrot. Long-run linearity, locally Gaussian processes, H-spectra and infinite variances. *International Economic Review*, 10:82–113, 1969.

[50] B. B. Mandelbrot. Intermittent turbulence in self similar cascades: Divergence of high moments and dimension of the carrier. *Journal of Fluid Mechanics*, 62:331, 1974.

[51] B. B. Mandelbrot. *The Fractal Geometry Of Nature.* W. H. Freeman and Company, New York, 1983.

[52] B. B. Mandelbrot. Limit lognormal multifractal measures. *Physica A*, 163:306–315, 1990.

[53] B. B. Mandelbrot. *Fractals and Scaling in Finance.* Springer-Verlag, New York, 1997.

[54] B. B. Mandelbrot. *Multifractals and* $1/f$ *Noise.* Springer-Verlag, New York, 1999.

[55] B. B. Mandelbrot and J. W. Van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM Review*, 10:422–437, 1968.

[56] B. B. Mandelbrot and J. R. Wallis. Noah, Joseph and operational hydrology. *Water Resources Research*, 4:909–918, 1968.

[57] B. B. Mandelbrot and M. S. Taqqu. Robust R/S analysis of long run serial correlation. In *Proceedings of the 42nd Session of the ISI*, Bulletin of the ISI, Vol. 48, Book 2, pp. 69–104, Manila, 1979.

[58] B. B. Mandelbrot and J. R. Wallis. Computer experiments with fractional Gaussian noises. *Water Resources Research*, 5:228–267, 1969.

[59] B. B. Mandelbrot and J. R. Wallis. Some long-run properties of geophysical records. *Water Resources Research*, 5:321–340, 1969.

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

[61] T. Mikosch, S. Resnick, H. Rootzén, and A. Stegeman. Is network traffic approximated by stable Lévy motion or fractional Brownian motion? Preprint, 1999. To appear in *The Annals of Applied Probability.*

[62] L. Olsen. Random geometrically graph directed self-similar multifractals. *Pitman Research Notes in Mathematics Series*, 307, 1994.

[63] C. Partridge. The end of simple traffic models. *IEEE Network* p. 3, Sept. 1993.

[64] V. Paxson and S. Floyd. Wide area traffic: The failure of Poisson modeling. *Proceedings ACM/SIGCOMM 94*, London, UK, 24:257–268, 1994.

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

[66] V. Paxson and W. Willinger. Finding order in chaos: Statistical challenges in analyzing the Internet. *Computing Science and Statistics*, Vol. 30, S. Weinberg (Ed.), Proceedings of the 30th Symposium on the Interface, Minneapolis, MN, May 1998, pp. 16–25, 1998.

[67] Y. Pesin and H. Weiss. A multifractal analysis of equilibrium measures for conformal expanding maps and Moran-like geometric constructions. *Journal of Statisitical Physics*, 86:233–275, 1997.

[68] L. L. Peterson and B. S. Davie. *Computer Networks.* Morgan Kaufmann Publishers, San Francisco, CA, 2000.

[69] V. Pipiras and M. S. Taqqu. The limit of a renewal-reward process with heavy-tailed rewards is not a linear fractional stable motion. *Bernoulli*, 6:607–614, 2000.

[70] V. Pipiras, M. S. Taqqu, and J. B. Levy. Slow and fast growth conditions for renewal reward processes with heavy-tailed renewals and either finite-variance or heavy-tailed rewards. Preprint, 2000.

[71] S. I. Resnick and G. Samorodnitsky. Fluid queues, leaky buckets, on-off processes and teletraffic modeling with highly variable and correlated inputs. In K. Park and W. Willinger, editors, *Self-Similar Network Traffic and Performance Evaluation.* Wiley (Interscience Division), 2000, pp. 171–192.

[72] V. Ribeiro, R. Riedi, M. S. Crouse, and R. G. Baraniuk. Simulation and queuing analysis of non-Gaussian long-range-dependent traffic using wavelets. IEEE/ACM Transactions on Networking, submitted 2001.

[73] R. H. Riedi. An improved multifractal formalism and self-similar measures. *Journal of Mathematical Analysis and Applications*, 189:462–490, 1995.

[74] R. H. Riedi. Multifractal processes. In P. Doukhan, G. Oppenheim, and M. S. Taqqu, editors, *Long-range Dependence: Theory and Applications.* Birkhäuser, 2001. Appears in this volume.

[75] R. Riedi, M. S. Crouse, V. Ribeiro, and R. G. Baraniuk. A multifractal wavelet model with application to TCP network traffic. *IEEE Transactions on Information Theory*, 45:992–1018, April 1999.

[76] R. Riedi and J. Lévy Véhel. Multifractal properties of TCP traffic: A numerical study. *Technical Report No 3129, INRIA Rocquencourt, France*, Feb, 1997. see also: J. Lévy Véhel and R. Riedi, "Fractional Brownian Motion and Data Traffic Modeling," in: *Fractals in Engineering*, pp. 185–202, Springer 1997.

[77] R. H. Riedi and B. B Mandelbrot. Exceptions to the multifractal formalism for discontinuous measures. *Mathematical Proceedings of the Cambridge Philosophical Society*, 123:133–157, 1998.

[78] R. H. Riedi and W. Willinger. Toward an improved understanding of network traffic dynamics In K. Park and W. Willinger, editors, *Self-Similar Network Traffic Analysis and Performance Evaluation*, Wiley, New York, 2000, pp. 507–530.

[79] J. Roberts, U. Mocci, and J. Virtamo (Eds.). *Broadband Network Teletraffic: Final Report of Action COST 242.* Lecture Notes in Computer Science, Vol. 1155, Springer-Verlag, New York, 1996.

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

[81] M. S. Taqqu. The modeling of ethernet data and of signals that are heavy-tailed with infinite variance. To appear in *Scandinavian Journal of Statistics*, 2001.

[82] 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*, pages 73–89, Boston, 1986. Birkhäuser.

[83] M. S. Taqqu, V. Teverovsky, and W. Willinger. Estimators for long-range dependence: an empirical study. *Fractals*, 3:785–788, 1995.

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

[85] M. S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review*, 27(2):5–23, 1997.

[86] M. S. Taqqu and V. Teverovsky. On estimating the intensity of long-range dependence in finite and infinite variance series. In R. Adler, R. Feldman, and M. S. Taqqu, editors, *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, pages 177–217, Boston, 1998. Birkhäuser.

[87] P. Whittle. Estimation and information in stationary time series. *Arkiv för Matematik*, 2:423–434, 1953.

[88] W. Willinger and V. Paxson. Where mathematics meets the Internet. *Notices of the American Mathematical Society*, Vol. 45, pp. 961–970, 1998.

[89] 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, pp. 27–53, 1998.

[90] 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. *Proceedings ACM/SIGCOMM 95*, Cambridge, MA, 25:100–113, 1995.

[91] 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 in Networking*, 5(1):71–86, 1997. Extended Version of the paper with the same title that appeared in *Computer Communications Review* Vol. 25 (1995) 100-113.

Walter Willinger, AT&T Labs-Research, 180 Park Avenue, Room C284, Florham Park, NJ 07932-0971, U.S.A., Email: `walter@research.att.com`

Vern Paxson, AT&T Center for Internet Research at ICSI (ACIRI), Berkeley, CA 94704, U.S.A., email: `vern@aciri.org`

Rudolf H. Riedi, Department of Electrical and Computer Engineering, Rice University, Houston TX 77005 - 1892, U.S.A., email: `riedi@rice.edu`

Murad S. Taqqu, Department of Mathematics, 111 Cummington St., Boston University, Boston, MA U.S.A., e-mail: `murad@math.bu.edu`