Where Mathematics meets the Internet

Walter Willinger and Vern Paxson*

Introduction

The Internet has experienced a fascinating evolution in the recent past, especially since the early days of the Web, a fact well-documented not only in the trade journals, but also in the popular press. Unprecedented in its growth, unparalleled in its heterogeneity, and unpredictable or even chaotic in the behavior of its traffic, "the Internet is its own revolution", as Anthony-Michael Rutkowski, former Executive Director of the Internet Society, likes to put it. At the same time, folklore has it that mathematics lies at the heart of Internet operation. After all, the argument goes, mathematics is the language of computers, and the Internet is currently connecting tens of millions of them, and still doubling every year [Lo98]. Yet the Internet is a new world, one where engineering reality wins over tradition-conscious mathematics and requires "paradigm shifts" that favor a combination of mathematical "beauty" and high potential for contributing to pragmatic Internet engineering. In this article, we take a look at how the Internet differs in fundamental ways from the conventional voice networks, how the (r)evolution of the Internet is impacting the world of mathematics in the small as well as in the large—both on *how* mathematics is done, and, for understanding the network itself, on *what sort of* mathematics is done—and why this, in turn, makes Internet engineering a gold mine for new, exciting and challenging research opportunities in the mathematical sciences.[1]

Teletraffic Theory and Internet Engineering

The term "teletraffic theory" originally encompassed all mathematics applicable to the design, control and management of the public switched telephone networks (PSTN): statistical inference, mathematical modeling, optimization, queueing and performance analysis. Later, its practitioners would extend this to include data networks such as the Internet, too. Internet engineering, an activity that includes the design, management, control and operations of the global Internet, would thus become part of teletraffic theory, relying on the mathematical sciences for new insights into and a basic understanding of modern data communications. However, from its early days, the Internet emphasized engineering and experimentation and was less concerned with mathematics and theory. In fact, some in the Internet community are quick to point out that today's Internet "works" because it ignored mathematics—in particular, teletraffic theory—and herein lies an interesting tale.

Mathematics and POTS

For someone living in an industrialized country, what is the likelihood of not getting a dial-tone when trying to make a phone call?[2] Now, what about not being able to connect to a popular web server over the Internet ...?

The answers to these questions range from once in a month or year in the first case, to once in an hour or day in the second case. Indeed, traditional teletraffic theory—as applied to POTS (plain old telephone service)—has arguably been one of the most successful applications of mathematical techniques in industry. It has led to first-rate telephone networks, whose quality of service we fully rely on and take for granted. It has enabled enormous efficiencies in the deployment and day-to-day operations of telecommunications networks and has resulted in near-universal telephony throughout the industrialized world. Among the main reasons for this tremendous success of

*W. Willinger is with AT&T Labs-Research, Florham Park, NJ 07932-0971, email: walter@research.att.com. V. Paxson is with the Network Research Group, Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94720, email: vern@ee.lbl.gov.
V. Paxson's work was supported by the Director, Office of Energy Research, Office of Computational and Technology Research, Mathematical, Information, and Computational Sciences Division of the United States Department of Energy under Contract No. DE-AC03-76SF00098. This paper appears in the *Notices of the American Mathematical Society,* 45(8), pp. 961–970, Sept. 1998.

[1]Note that this article is not intended to provide a comprehensive bibliographical guide to the latest developments and advances in this area; for such a guide, the interested reader may want to consult, for example, the recent survey paper [WTE96].

[2]Here we mean a voice call. We address the interesting case of dialing up to an Internet Service Provider below.

teletraffic theory and practice in traditional telephony are the *highly static* nature of conventional PSTNs and a well-defined and ever-present notion of *limited variability*, a trademark of homogeneous systems where one can talk about "typical" users and "generic" behavior and where averages describe system performance adequately. Another important reason has been the special appeal of the most widely used models to the engineering community, mainly because of their simplicity, physical interpretation, and practical relevance: they required only a few inputs that could be readily estimated in practice.

The static nature of traditional PSTNs contributed to the popular belief in the existence of "universal laws" governing voice networks, the most significant of which is the presumed *Poisson* nature of call arrivals at links in the network where traffic is heavily aggregated, such as at inter-office trunk groups. This law states that call arrivals are mutually independent and that the call interarrival times are all exponentially distributed, with one and the same parameter λ.

Equivalently, if $X = (X_k : k \geq 1)$ denotes the number of call arrivals in successive, non-overlapping time intervals of length $\Delta t > 0$, then X is the increment process of a Poisson process with parameter λ if and only if the random variables X_k are independent and identically distributed with

$$P[X_k = n] = e^{-\lambda \Delta t} \frac{(\lambda \Delta t)^n}{n!}, \quad n \geq 0. \tag{1}$$

Traffic models such as the Poisson process, whose full dynamics can be described with one or just a few parameters, are termed *parsimonious*, a highly desirable property for reasons we develop later.

The Poisson law has remained valid for modeling purposes for at least fifty years. So has a related invariant of POTS traffic that specifies that call "holding times" (durations) follow more or less an exponential distribution. Three other important teletraffic laws are: growth rates are highly predictable, allowing for fine-tuned short- and long-term capacity planning; network controls and operations are fully centralized, so one can envision taking advantage of information about the network's global state; and offered services are strictly regulated and monitored.

On the other hand, the static PSTN environment has resulted in a steady decline of the perceived importance of continued measurements, and has emphasized instead the need for analytical techniques. While teletraffic theory was originally based on empirical studies and on traffic measurements that were collected laboriously from the public telephone networks[3], soon the belief in the Poisson process and the exponential distribution as "universal laws" for POTS overcame the curiosity associated with collecting and analyzing more data. Moreover, new mathematical results provided a sound physical basis for the observed Poisson nature of call arrivals on trunk groups. As an example, the Palm-Khintchine theorem states that the superposition of many independent and properly normalized renewal processes—each one describing the call arrivals on a single phone line—forms a Poisson process. The resulting traffic models were, in general, mathematically tractable and could be used to predict accurately many performance measures of interest. Queueing theory was born. A faith in "true" traffic models took over, the need for further traffic measurements was glossed over, and the main focus shifted to turning queueing theory into a full-fledged mathematical discipline.

Ironically, the complacency engendered by this mathematical elegance and (particularly) success has recently been rocked by changes in the "static" world of telephony. Fifty-year-old patterns of telephone use, the bedrock of the teletraffic modeling success story, now have been greatly undermined by two major new uses of the telephone network. These changes began with the advent of FAXes in the 1980s and have continued and become more drastic with the popularity of the Web.

The key change is that telephone calls used for FAX transmission and Internet access have statistical characteristics dramatically different from a typical voice call. They tend to be significantly longer and much more variable in their duration than a voice call, and their numbers have recently increased dramatically, especially in terms of Internet access calls. Both types of calls are now playing havoc with the existing PSTN engineering infrastructure designed to deal with voice calls only. In some places, call "blocking" has increased to unacceptable levels, especially during late evening hours (popular with Web surfers), and ad-hoc engineering methods have become necessary to prevent temporarily Internet access calls from saturating access to the public telephone network. Clearly, theory no longer meets reality and as a result, capacity planning becomes dicey and inexact, and concentrated, industry-wide efforts for off-loading Internet traffic from the PSTN are under way.

Goodbye Poisson

One might expect that the voice network modeling success story would enjoy another triumph when applied to data networks, and indeed this has been attempted. But in fact much of the voice traffic modeling has proven nothing

[3]Pioneering work by Erlang, Palm, Wilkinson and others, over half a century ago.

short of disastrous when applied to data networks, for the simple but profound reason that *the rules all change* when it is computers and not humans doing the talking.

Voice traffic has the property that it is relatively homogeneous and predictable, and, from a signaling perspective, spans long time scales. Consequently, many concurrent voice connections can be easily "multiplexed" to 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 determine whether a given link has sufficient capacity to carry the additional load. As a result, voice networks have been engineered in a *circuit-switching* fashion. That is, the "routers" internal to the network, which are responsible for forwarding traffic from one link to the next so that it ultimately reaches its destination, keep track of each currently active connection, and when new traffic arrives, look up its corresponding connection to determine where to forward the traffic. The principal abstraction is known as providing "virtual circuits," because the network behaves as if it provides a direct circuit from the traffic source all the way to its destination.

In contrast to voice traffic, data traffic is much more variable, with individual connections ranging from extremely short to extremely long and from extremely low-rate to extremely high-rate. These properties have led to a design for data networks in which each individual data "packet" or "datagram" transmitted over the network is forwarded through the network independently of previous packets that may have been transmitted by the same connection. Each packet is self-contained, and the routers need only inspect the "header" of the packet to determine its destination and forward it through the network. Consequently, the routers do *not* keep track of each currently active connection.

This shift away from circuit-switching toward *packet-switching* has profound implications. On the one hand, it results in highly *efficient* networks. Any time capacity is available in the network, newly arriving packets can 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 "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. In addition, packet switching buys enormous *robustness*: it enables networks to route transparently around router or link failures without perturbing active connections. 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—a situation very different from a circuit-switched network, in which the routers cannot easily accept rerouted traffic because they have no knowledge of the corresponding virtual circuit.

However, links can become overloaded because packets arrive for transmission along them at rates 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 in the routers will fill up and some packets must be discarded, or *dropped*. To ensure that sources behave properly in the presence of congestion in the network, the protocols used for transmitting data in the Internet include *end-to-end congestion control* mechanisms that decrease automatically the rate at which data are transmitted when congestion is detected. An important consequence of the use of congestion control is that traffic in the network is *shaped* by the conditions each connection has encountered in its past. Thus, Internet traffic includes a basic mechanism that introduces significant, complicated correlations across time, as well as complex interactions among the active connections.

A damaging legacy of the telephony influence on data network research was a virtually complete absence in the 1970s and 1980s of attempts to validate crucial modeling assumptions against actual data network traffic measurements. Yet, just a few measurement studies suffice to discover that data traffic is highly variable or very *bursty*. That is, it does not come at a steady rate, but instead in starts and fits, with lulls in between. The term "bursty" has a readily understood intuitive meaning, but it turns out that nailing down its precise, mathematical meaning has profound implications for developing mathematical models of network traffic. The natural approach for getting a handle on burstiness is to define it in terms of a *time scale* over which activities and lulls occur. For telephony, this time scale is related to the rate λ of the Poisson process (1) that describes the dynamics of call arrivals. For example, if $\lambda = 100/\text{sec}$, then the time scale of burstiness is around 10 msec, and periods of sustained, greater-than-average activity or sustained, lower-than-average lulls over much smaller or larger time scales occur with rapidly vanishing probability.

However, practitioners have long observed that traffic bursts in data networks do indeed occur on many *different* time scales and that such multi-scale burstiness simply does not fit the world of traditional Poisson-based traffic modeling. The Poisson framework does not even provide a vocabulary for discussing burstiness of this sort.[4] Figure 1 is a visual demonstration of the failure of Poisson modeling to capture the burstiness present in actual network traffic. The plots were generated based on an hour-long trace of Internet traffic collected off a network link connecting a

[4]Indeed, we find researchers falling back on metaphors to try to characterize their observations: "traffic 'spikes' (which cause actual losses) ride on longer-term 'ripples', that in turn ride on still longer-term 'swells'." [FL91]

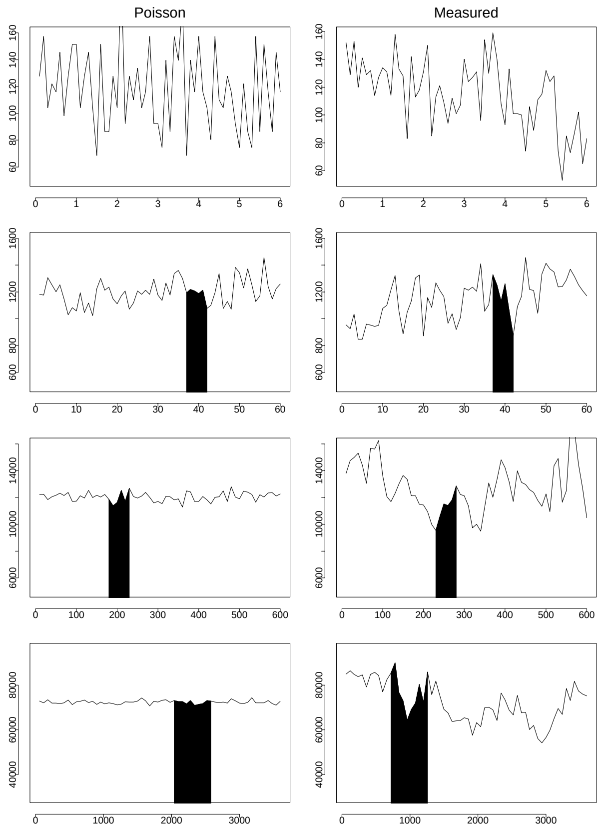

Figure 1: Synthesized traffic from a Poisson model vs. Internet traffic to which its mean and variance were fit, viewed over three orders of magnitude

large corporation to the Internet.[5] From this trace we synthesized another, hour-long series of packet arrivals created by fitting a simple Poisson-based model to the mean and variance of the measured sample. More elaborate modeling could be done, but the end effect (see below) would be the same.

We then observe visually the burstiness of the original trace and the synthetic trace as we vary the time scale of observation. The top row shows a randomly selected subset of each trace on a time scale of 100 msec; that is, each point in the plots reflects the number of packets observed during a 100 msec interval, for a total of 6 sec. The second row shows a time scale a factor of ten larger: now, each point reflects the number of packets per 1 sec, spanning 60 sec total. The black regions illustrate from where the plots in the row above were made. A key point is that we not only have increased the scale of the X-axis by a factor of 10, but we have done the same to the Y-axis. With the third row we have again increased the scale in both X and Y by a factor of ten, and in the final row by another factor of six, such that here the plots span the entire hour of the traces.

The difference between the Poisson model and the measured traffic is obvious and striking: as the time scale increases, the Poisson traffic "smooths out," becoming quite tame, while the measured traffic shows no such predilection. This difference is *absolutely crucial* from an engineering perspective: traffic that behaves as shown in the left column can be easily engineered for. Above a certain time scale, there are no surprises—everything boils down to knowing the long-term arrival rate; no need for big buffers in routers or switches, no reasons for being conservative in choosing safe operating points for engineering backbone trunks, and why even think of user-perceived quality-of-service as being a relevant issue? In stark contrast, measured traffic like that shown in the right column is *wild*, remains so even on quite coarse time scales, and plays havoc with conventional traffic engineering: routers require big buffers to accommodate the traffic fluctuations across many time scales; in the absence of any effective controls, safe operating points have to be set conservatively because the traffic can saturate the link at any time and over any time scale; and adequate overall network performance can no longer be taken as a guarantee of happy individual users.

[5] The measurements were gathered by J. Mogul in 1995, and are available from the Internet Traffic Archive, http://www.acm.org/sigcomm/ITA/.

The edifice of Poisson modeling repeatedly told Internet network engineers to expect the behavior shown on the left—but what they really observe is the roller-coaster ride on the right! The Internet engineering community has thus come to consider teletraffic theory as irrelevant (and, actually, detrimental) to the development of the Internet. More specifically, it has criticized the Poisson-based approach on grounds that the models: (i) have little in common with network engineers' practical experience observing their networks, (ii) are theoretical constructs based on assumptions lacking validation against measured data, especially when extended with additional parameters for describing burstiness, (iii) are too complex to aid in developing intuition or a physical understanding of actual network traffic dynamics ("black boxes"), and (iv) require inputs (parameter estimates) that, in practice, cannot be specified, collected or estimated.

Hello Fractals

Many networking experts argue that the only way to gain an in-depth understanding of data network traffic is—simply put—doing away with teletraffic tradition and starting from scratch. Interestingly, mathematics, which has been largely responsible for the success story of teletraffic theory for the voice network, has recently provided strong ammunition in support of the networking experts' arguments. However, as voice traffic turns out to differ drastically from data traffic, so too do the underlying mathematical ideas and concepts. The relevant mathematics for POTS is one of *limited variability* in both time—traffic processes are either independent or have temporal correlations that decay exponentially fast—and in space, i.e., the distributions of traffic-related quantities have exponentially decaying tails. But for data networks, the mathematics is one of *high or extreme variability.* Statistically, temporal high variability in traffic processes is captured by *long-range dependence*, i.e., autocorrelations that exhibit a power-law decay. On the other hand, extreme forms of spatial variability can be described parsimoniously using *heavy-tailed distributions* with infinite variance, i.e., probability distributions F with the property that for large x-values,

$$1 - F(x) \approx \kappa_1 x^{-\beta}, \tag{2}$$

where κ_1 is a positive finite constant that does not depend on x and where the tail index β is in the interval $(0, 2)$. This property is, for example, satisfied by the well-known family of "Pareto distributions," originally introduced for modeling the distribution of income within a population.

It turns out that power-law behavior in time or space of some of their statistical descriptors often cause the corresponding traffic processes to exhibit *fractal* characteristics. In the present context, we say that a traffic process has fractal characteristics if there exists a relationship between certain quantities Q of the underlying process and the resolution τ, of the general form

$$Q(\tau) \approx \kappa_2 \tau^{f(D)}, \tag{3}$$

where κ_2 is a positive finite constant that does not depend on τ. Here τ denotes a resolution in time or space at which Q is evaluated, and (3) specifies how Q must vary as a function of the resolution τ; $f(\cdot)$ is a simple, often linear, function of D; and D is a fractal dimension. To declare fractality, the above relationship is supposed to hold for a range of different τ-values, with a value of D that is less than the embedded dimension.

Fractal concepts have been non-existent in teletraffic theory. Yet, a look at Figure 1 (right side) shows fractal-like behavior, over a wide range of time scales, from hundreds of milliseconds to seconds to tens of seconds and beyond. In fact, Figure 2 shows the same sort of plot as in Figure 1, except now instead of using a Poisson-based model, we use a very simple mathematical model called *fractional Gaussian noise* that is strictly fractal in a sense to made precise shortly. For now, a covariance-stationary Gaussian process $X = (X_k : k \geq 1)$ is called a *fractional Gaussian noise* with *Hurst* parameter $H \in [0.5, 1)$ if the autocorrelation between X_n and X_{n+k}, $k \geq 0$, is given by $\mathbf{cor}(X_n, X_{n+k}) = 1/2(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H})$. Along with fitting the model to the measured traffic's mean and variance, it requires one extra parameter, the Hurst parameter H, which quantifies the strength of the fractal scaling. Visually, the synthesized traffic using the fractal model is in the right ballpark, and we can achieve this using only *one* additional parameter! This last comment is a *crucial* property of the fractal modeling approach: it preserves *parsimony*, meaning that the model is sufficiently simple that we have some hope of applying it across a wide range of conditions without requiring too many guesses as to how to set its parameters.

In view of the general skepticism that exists in the different circles in the mathematical community concerning the need, usefulness, and appropriateness[6] of fractals, what can we say about fractal-like scaling in measured data network traffic? To examine this question, we call a discrete-time, covariance-stationary, zero-mean stochastic process

[6] See for example the recent survey by Avnir et al. [ABLM98] that reports, for all the *Physical Review* journals from 1990 to 1996, a scaling range of experimentally-declared fractality that averages a mere 1.3 decades (orders of magnitude, base 10).

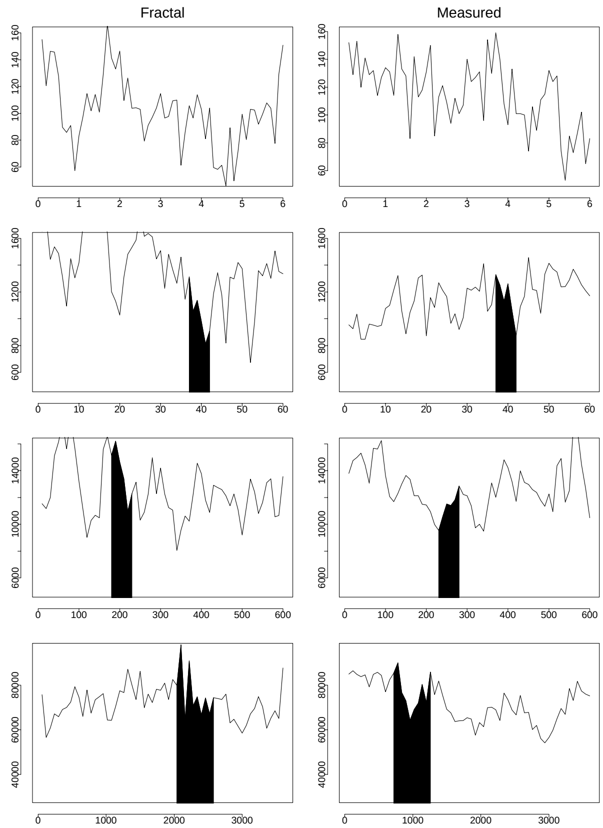

Figure 2: Synthesized traffic from a simple fractal model vs. Internet traffic to which its mean, variance and Hurst parameter (H) were fit, viewed over three orders of magnitude

$X = (X_k : k \geq 1)$ *exactly self-similar* or *fractal* with scaling parameter $H \in [0.5, 1)$ if, for all levels of aggregation or "resolution," $m \geq 1$,

$$X^{(m)} = m^{H-1} X,$$

where the equality is understood in the sense of finite-dimensional distributions, and where the aggregated processes $X^{(m)}$ are defined by:

$$X^{(m)}(k) = m^{-1}(X_{(m-1)k+1} + \cdots + X_{km}), \quad k \geq 1.$$

For example, the fractional Gaussian noise process introduced earlier is exactly self-similar with scaling parameter equal to the Hurst parameter. It is easy to check that for an exactly self-similar process with scaling parameter H, the functional relationship given by

$$\mathbf{Var} X^{(m)} = \kappa_1 m^{2H-2}$$

holds and fits the form of (3). The resulting linear log-log representation of $\mathbf{Var} X^{(m)}$ vs. m is called the *variance-time plot.* An illustration of it based on the Internet traffic trace used for Figures 1 and 2 is shown in Figure 3. Clearly, the observed scaling range spans 3 decades, indicating compelling evidence of fractal-like scaling. In other traces from different networks, scaling ranges spanning 3–5 decades are common.

When assessing the validity of describing a process using a self-similar model, one must be very careful not to mistake actual non-stationarities (e.g., connection arrival rate varying with time) for highly-variable but stationary fractal behavior. The two can appear very similar, both to the eye and to a number of statistical tests. However, this concern can be addressed by making good use of the very large size of network traffic traces. For example, we can extract numerous five-minute portions of a trace, analyze them for possible fractal behavior, and then compare the results to those for neighboring five-minute portions, and also to encompassing ten-minute portions, to see if the analyses yield consistent results, which then supports arguing that the data are well-modeled as stationary. Ordinarily, this process might encounter problems as we run out of data points and the subsamples become too small for compelling analysis. But for network traffic, we have "data to burn." In turn, this motivates the development of novel statistical techniques that exploit fully the impressive sample sizes and replace traditional concepts that

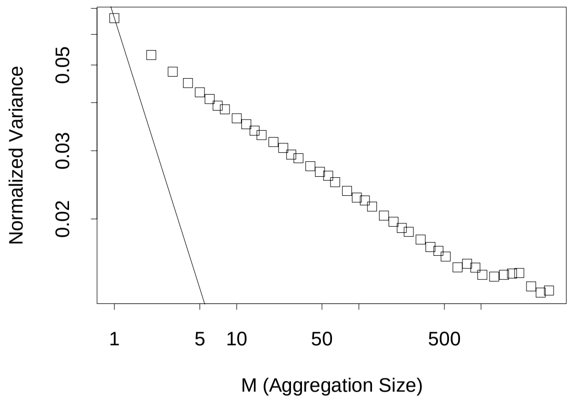

Figure 3: Variance-time plot for measured Internet traffic. The steep downward line to the left corresponds to the variance-time plot expected for Poisson traffic.

have been fine-tuned over the years to work to perfection when the sample size is small. As shown above, these new methodologies have little in common with traditional inference techniques such as assessing goodness-of-fit or hypothesis testing, but instead demonstrates how the concept of "borrowing strength from large data sets," coined by J. Tukey, applies in practice and can result in compelling evidence for, in this case, the fractal nature of Internet traffic.

Why getting to know the Internet is painfully hard

Although the finding of the fractal nature of Internet traffic can be viewed as a promising start toward solid characterizations of Internet traffic, the truth is that it has barely made a dent, because there are a number of often underappreciated features that make it immensely difficult to characterize and understand the Internet in any sound fashion. Each property reflects a form of *change*—the fact that virtually nothing about the network is "typical" in any sense. There is change over time, between sites, and in the basic assumptions about how the network is used [Pa94].

The changing Internet

The first, basic element of change concerning the Internet is that of *growth*. Simply put, the network grows exponentially, has done so for well over a decade, and shows no signs of slowing down. Figure 4 illustrates one growth statistic: the volume of traffic in bytes/day flowing through the USENET bulletin board system. The data start in 1984 and continue to 1994. The measurements fit beautifully a straight line, reflecting *sustained exponential growth* of about 80%/year for *over a decade* (note log-linear scale). Clearly, Internet growth is nothing new—it in no way began with the Web—and current statistics are consistent with the growth continuing completely unabated.

Another respect in which the Internet exhibits striking change concerns characteristics of its traffic as measured at different sites and moments in time. For example, a 1991 sample of Internet traffic at a research laboratory found that about 67% of the data bytes going into or out of the site were from use of the Internet's standard file transfer protocol, FTP. Yet, at the same point in time, a sample of a university's traffic found that only 18% of the data bytes were from FTP. For another university, the figure was 50%. Clearly, if researchers studied just one of these sites, regardless of how carefully, they would have come to a conclusion completely incorrect for some other Internet sites!

The same problem also occurs over modest amounts of time. An October 1992 sample of file transfer (FTP) traffic at a research laboratory found that the median data transfer size was 4,500 bytes. This median was computed over more than 60,000 transfers—presumably a highly robust statistic. Yet five months later, the same statistic computed over more than 80,000 transfers yielded 2,100 bytes, less than half the earlier value! In March 1998, the same median, now computed over 450,000 transfers, was 10,900 bytes. Thus, regardless of how diligently one studied any one of these three points in time, conclusions about this "highly robust" statistic were grievously inaccurate for the other points in time!

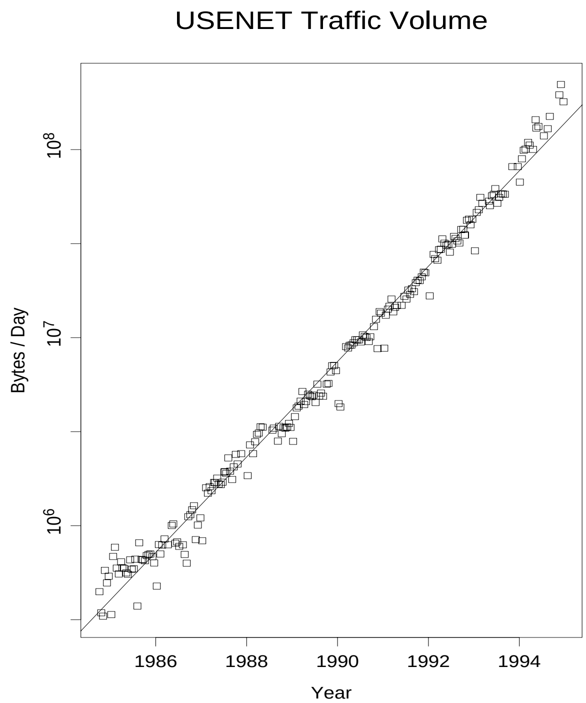

Figure 4: Bytes per day sent through the Internet's USENET bulletin board system. Data courtesy of R. Adams.

Another basic source of problematic Internet change concerns the extremely rapid advent of new applications. In October 1992, the aforementioned research lab participated in a total of 45 WWW connections, over the entire month. But in the months that followed, the Web *exploded*, and that same site's Web traffic began to *double every six weeks* and continued doing so for *two entire years*. Today the site participates in more than half a million WWW connections each day. From the perspective of Internet researchers, the Web came from *out of nowhere*. No one, not even its staunchest proponents, had predicted such rapid success in their wildest dreams. Worse, its traffic has properties that until then were not common in the Internet. Carefully considered extrapolations regarding future traffic became *obsolete* virtually overnight![7]

Toward scientific inference

Clearly, if one's goal is to understand and predict Internet behavior in any sound fashion, then the difficulties outlined in the previous section must be sobering; they demonstrate that there are enormous hurdles to overcome in terms of how rapidly the network changes, and the great diversity embodied within it. One approach for dealing with the pace at which the Internet changes, as well as with its extreme heterogeneity, is to base findings about the Internet on careful examinations of a *wide range* of Internet measurements, taken at different points in time, at different points in the network, and under a variety of different networking conditions. However, this approach has its own drawbacks: not only does it mean diligent and immensely time-consuming "digging around" in data, but statistical inference as it is currently taught and practiced has little to offer when faced with the task of drawing statistically sound conclusions from a *large number* of *large* data sets. Conventional statistical inference emphasizes the analysis of *single* data sets that are typically *small*; works to near perfection when it comes to the testing of "true models" using small samples; and has developed over the years an arsenal of techniques and tools that help an analyst "squeeze a data set dry" [Ch95].

What Internet traffic researchers instead require are inference methods that can fruitfully span a large collection of high-volume data sets. They need tools for searching for law-like relationships across different data sets that generalize to a wide range of different conditions. These approaches define what is generally called *scientific inference*. They have a long history in the physical sciences but have been all but ignored in the social sciences and in the traditional

[7] We note that this phenomenon is not unique to the Web. Other applications have also exploded. The Web—*so far*—is the only one that has continued to explode for more than a few years.

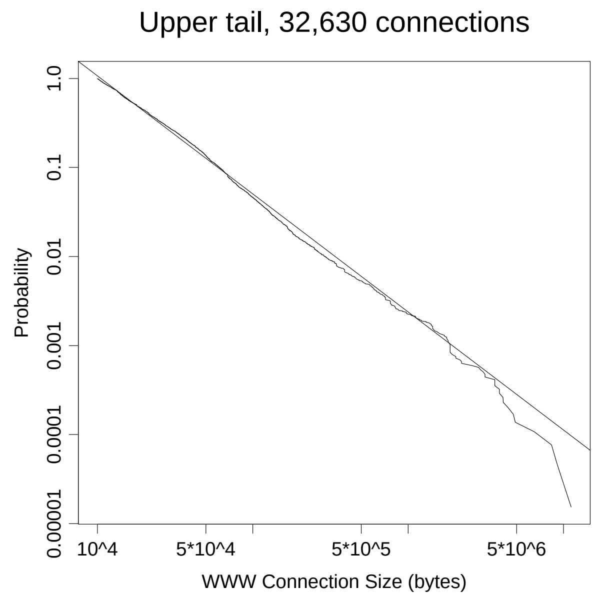

Figure 5: (Conditional) log-log complementary distribution plot of WWW connection sizes, given that the connection size is at least 10,000 bytes.

statistics literature. To study Internet traffic, we want scientific inference for what J. Tukey calls "broadening the basis," which means trying to uncover *traffic invariants*, i.e., features in traffic that are insensitive to the constantly changing conditions that networks experience. Such an approach emphasizes building intuition and physical understanding over traditional black-box descriptions or conventional data fitting. At the same time, scientific inference provides the proper framework for Internet traffic researchers who are desperate for *parsimonious* models of Internet traffic: any model with an immodest number of parameters is *doomed* to impracticality, because knowledgeable researchers know there is no hope of assigning meaningful values to all of the parameters. Put another way: parsimony in the context of Internet traffic is achievable only if the search for traffic invariants turns out to be successful.

Is there hope? Is there Math?

Unfortunately, there exist no scientific inference recipes for the identification of traffic invariants in an abundance of high-quality, high-volume data sets of Internet traffic measurements. One can, of course, still try to find invariants by diligent manual analysis and hard or unconventional thinking. Any successfully-identified invariant becomes worth its weight in gold, as it offers hope that some sort of coherent and parsimonious model of the network might actually prove attainable. In the following, we outline briefly what we consider to be some of the successfully-identified invariants to date.

First, while Poisson models have been decisively rejected as a basis for characterizing the arrivals of individual data packets in the Internet, there is solid evidence that these models *do* apply for the much more modest domain of characterizing the "arrivals" of humans to the Internet. That is, the times at which people begin using the Internet for a specific task do indeed conform to a memoryless process with an arrival rate that can be deemed constant over time intervals of many minutes to perhaps an hour. The basis for this invariant are data sets of Internet-related measurements that contain information about the start times of, for example, TELNET and FTP connections [PF95], or WWW sessions [FGWK98], collected over a number of years and at numerous locations.

Another, much more intriguing invariant, is that when considering the sizes (in number of bytes or packets) or durations (measured in seconds) of a set of network sessions or connections, one almost always finds that the empirical distribution exhibits the heavy-tailed property (2), with $\beta < 2$ and, sometimes, even $\beta \approx 1$. These cases indicate extreme variability: $\beta < 2$ means that the traffic process at hand is well-modeled as exhibiting *infinite variance*, and in the case $\beta \leq 1$, as having an *infinite mean*.

Figure 5 illustrates that these heavy tails are very well grounded in measured data. The data for this plot came from a day's WWW traffic at a large research laboratory. The day was chosen arbitrarily from (literally) hundreds of days' worth of recorded traffic. We now look at the size of each WWW connection. All in all, there were 226,000 connections. If we restrict ourselves to those connections transferring at least 10,000 bytes (the upper 14% tail), and plot their complementary distribution function (Eqn 2) against the corresponding size, both on a log-log scale, then we get the plot shown in Figure 5. A straight line on such a plot corresponds to tail behavior that agrees with that of a Pareto distribution, and its slope gives $-\beta$. It is strikingly clear that the more than 32,000 points plotted in Figure 5 do indeed fall on a line, and that, with $\beta \approx 1.3$, the data are indeed consistent with *infinite variance*.

Note that the heavy tail property is for the distribution of an *aggregate* property of a traffic source, such as how much total data it will send. It says nothing about *how* the source will in fact send the data when dividing them into a series of packets for transmission across the network. Consequently, one might well wonder to what use we can possibly put the finding of the infinite variance property at the session or connection level as a traffic invariant. The surprising answer is that there are new mathematical results that relate the presence of connection sizes or durations with infinite variance directly with the finding of fractal scaling in aggregate network traffic at the packet level! Thus, the compelling presence of the infinite variance property in data set after data set of connection-level Internet measurements has also become the bedrock of the shift away from Poisson-based modeling of data traffic over to fractal-based modeling. That it is an invariant therefore explains why fractal scaling is an invariant. In addition, it turned out to be the basis for a very simple physical explanation of the empirically observed fractal nature of aggregate network traffic (i.e., total number of packets or bytes per time unit). As such, heavy tails aided immensely in de-mystifying fractal traffic modeling.

Even more striking is that, in fact, the progression of results proceeded the opposite of what we outlined above. It was *not* the case that researchers observed the heavy-tailed or infinite variance property of individual connections and then went from there to postulate fractal traffic models. Instead, based on an extensive analysis of numerous traces collected from different local area networks during a 4 year period, and by applying the principles of "broadening the basis" and "borrowing strength from large data sets," some researchers *first* made the—at the time, nearly crazed—leap to postulate fractal traffic models [LTWW94]. Furthermore, while at that time the researchers could not directly answer the natural question of "Why fractal??" they *did* speculate as to possible mechanisms—speculation that basically told the network research community "go look for heavy tails." Once the researchers knew what to look for, they started finding them everywhere! For example, heavy tails can be found in: CPU time consumed by different processes; sizes of files in a file system; Web item sizes; inter-keystroke times when a person types; sizes of FTP bursts; and sizes and durations of bursts or idle periods of individual Ethernet connections.

These examples serve to illustrate that some of the hard-won progress to date toward getting to know the dynamics of Internet traffic has come from close collaborations between mathematicians and networking researchers. On the one hand, mathematicians feel generally overwhelmed by all the details related to the architecture of modern-day data networks, the underlying protocol hierarchies and the different network technologies. Nevertheless, many of these apparently minor details need to be understood to ensure that the mathematics research does not become detached from the networking application. On the other hand, networking researchers are generally less interested in the fine details of a mathematical proof or definition but want to be convinced at an intuitive level and/or through empirical arguments. When left on their own, Internet experiments and instrumentation, and the resulting measurements and prototypes, are impressive engineering achievements; and the theoretical results in fractal geometry represent intriguing and beautiful mathematics. However, the prospect is that in combination, they will contribute to a significantly improved understanding of the Internet.

Should mathematicians care?

The original finding of fractal scaling phenomena in Internet traffic was greeted with skepticism by many mathematicians. They considered it as yet another example of a "fad" that comes and goes, with ultimately nothing to show for it, similar to what had happened in other areas in the natural or social sciences such as hydrology, economics, or biophysics, where the fractal "craze" proved to be short-lived and had absolutely no impact beyond some philosophical discussions about the general purpose of modeling.

What these mathematicians missed was that the application of fractal analysis to networking was fundamentally different from these other applications. In addition to Internet engineering reality being the driving force, the available data sets are unique and outstanding, not only with regard to volume and quality, but, more importantly, with respect to the amount of information that is contained in every observation (i.e., data packet). This information provides detailed knowledge about the different layers in the hierarchical structure of modern-day networks, about how the different protocols that operate on those layers interact with one another, and, indirectly, about interactions between the different connections that share a given link. The richness of the data has a profound impact on how the data sets are analyzed, interpreted and modeled. It is difficult to think of any other area in the sciences where the available data provide such detailed information about so many different facets of behavior.

The switch from Poisson to fractal thinking in network traffic research has had a major impact on our understanding of actual network traffic, to the point where we now know why aggregate Internet traffic exhibits fractal scaling behavior over time scales from a few hundreds of milliseconds onwards. A measure of the success of this shift in thinking is that the corresponding mathematical arguments are at the same time rigorous and simple, are in full agreement with the networking researchers' intuition, and can be explained readily to a non-networking expert. An equally important part of this new understanding is the realization that we *do not* yet have a similarly clear picture of the dynamics of Internet traffic over fine time scales, from hundreds of milliseconds downwards, where the end-to-end congestion control mechanisms determine the flow of packets at the different layers in the networking hierarchy. However, recently reported empirical findings suggest that measured Internet traffic over those small time scales exhibits pronounced local irregularities that are consistent with *multifractal* scaling behavior and can be analyzed effectively using wavelet-based techniques. While wavelets can be expected to advance significantly the multifractal analysis of Internet traffic, the networking application is equally likely to influence the development of new wavelet-based techniques: ways to exploit fully the properties and rich structure of the available data sets.

Finally, we mention developments of a different nature that again beg for mathematical attention, and hold promise for interesting mathematical problems. Several networking research projects are now working on *systematic* Internet measurements: sets of potentially thousands of "probe platforms" deployed throughout the network that engage in both independent and orchestrated measurement of network paths in attempts to characterize the network's behavior and to locate trouble spots. Such a network-wide view of Internet traffic dynamics includes both temporal and spatial dimensions, as well as a dimension defined by the different layers in the networking hierarchy. Clearly, the interesting problems here are those of interactions, correlations and heterogeneities in time, space, and across the different networking layers. In addition, when the ultimate goal is to enable the tens of millions of Internet users to determine what performance they can obtain from the network, irrespective of where they are and when they want this information, and how to improve the engineering of the network to meet their myriad needs, then the analysis problems acquire a central element of *scale*, extending well beyond what has previously been attempted. While the sheer scale may appear daunting, we still have the significant advantage of superb data sets with which to work. The problems then acquire a character of tantalizing challenge, and solving them moves beyond mere interesting mathematics, into the regime of answering questions that will help determine just how effective this monstrous emerging global infrastructure actually proves.

Acknowledgments

We thank Rob Calderbank, Sally Floyd, and Andrew Odlyzko for their careful reading of various drafts of this article, and for their very helpful comments; we also acknowledge the useful suggestions of three anonymous reviewers. We are especially grateful to the Editor, Anthony Knapp, for his encouragements, support, guidance and enthusiasm, and to Hugo Rossi, the former Editor, for the invitation to write this article.

References

[ABLM98] D. Avnir, O. Biham, D. Lidar and O. Malcai, "Is the geometry of nature fractal?," *Science*, 279(5347), pp. 39-40, Jan. 2, 1998.

[Ch95] C. Chatfield, "Model uncertainty, data mining and statistical inference," *J. R. Statist. Soc.* **A**, 158:419–466, 1995.

[FGWK98] A. Feldmann, A. C. Gilbert, W. Willinger and T. G. Kurtz, "The changing nature of network traffic: Scaling phenomena," *Computer Communication Review* **28**, No. 2, April 1998.

[FL91] H. Fowler and W. E. Leland, "Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management," *IEEE Journal on Selected Areas in Communications*, 9(7), pp. 1139-1149, Sept. 1991.

[LTWW94] W. E. Leland, M. S. Taqqu, W. Willinger and D. V. Wilson, "On the self-similar nature of Ethernet traffic (Extended Version)," *IEEE/ACM Transactions on Networking* **2**, pp. 1–15, 1994.

[Lo98] M. Lottor, *http://www.nw.com/zone/WWW/top.html*; February 1998.

[Pa94] V. Paxson, "Growth Trends in Wide-Area TCP Connections", *IEEE Network*, 8(4), pp. 8-17, July/August 1994.

[PF95] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Transactions on Networking*, 3(3), pp. 226-244, June 1995.

[WTE96] W. Willinger, M. S. Taqqu and A. Erramilli, "A bibliographical guide to self-similar traffic and performance modeling for modern high-speed networks." *Stochastic Networks: Theory and Applications*, F.P. Kelly, S. Zachary and I. Ziedins (Eds.), Royal Statistical Lecture Note Series, Vol. 4, pp. 339–366, Clarendon Press, Oxford, UK, 1996.