Investigating the Scaling Behavior, Crossover and Anti-persistence of Internet Packet Delay Dynamics

Qiong Li* David L. Mills†
Department of Electrical and Computer Engineering
University of Delaware

September 23, 1999

Abstract

Relying on the wavelet-based time-scale analysis techniques and a so called height difference correlation function method, we studied the dynamics of the delay traces measured in the Internet with sampling interval ranging from 20 ms to 1 min. Scaling analyses showed that the delay traces have a complicated scaling behavior changing with both the time scales and the sampling intervals. A significant observation of our research was that the nature of the delay traces changes from a LRD noise-like time series (e.g., fractional Brownian noise) to a self-similar one with an outer cutoff scale when the sampling intervals decrease (we refer this as crossover). It turns out that the delay jitter time series derived from a sampled delay trace with small sampling interval is LRD noise with a Hurst parameter $H < 0.5$ (anti-persistence), indicating that the current Internet works in a controlled state. We also used a wavelet-based partition function method to investigate the possible multifractality of the delay traces. We analyzed the underlying physical mechanisms of the observed scaling behavior, crossover and anti-persistence, and discussed the implications of our findings on several aspects of networking engineering.

1 Introduction

Among the many aspects of the Internet dynamics [15], end-to-end packet delay dynamics plays an important role in characterizing the end-to-end behavior. There are several reasons for studying the end-to-end delay dynamics. First, the transport layer protocols such as TCP, rely on measuring end-to-end delay as a feedback mechanism to do flow control [8]. Better understanding about the delay dynamics, especially in a small time scale, will provide clues and guidelines to design and tune the protocols to achieve a better performance. Secondly, the Internet moving on to provide more real-time, interactive multimedia services. Many of them may have strict requirements on the quality of service (QoS) which is quantitatively measured by throughput, loss rate, delay and delay jitter. Understanding and modeling the end-to-end delay are essential to predicting the viability of the real-time services on the Internet.

Third, The possibility of building a realistic traffic model for simulation study depends on how well the characteristics of the traffic (include LANs and WANs) being understood. Recent researches [9, 16, 6] have empirically demonstrated the self-similarity characteristic of LAN traffic and the multifractal nature of Internet WAN traffic. A type of structural models taking into account specific features of the underlying network structure and hence providing a physical explanation for the observed scaling phenomena have been proposed [18]. These approaches reflect a source-centered view of traffic modeling, ignoring the interaction between the source and the network side, such as in the case of TCP connection, the sender would adjust the sending rate based on the current conditions of the network. The multifractal nature of the traffic found in the environment of WANs [6], especially in small-time scale, puzzles researchers with the question of the underlying physical mechanism of the multifractality. A cascade traffic model [6], viewing the aggregated traffic as multiplicative instead of additive, can mathematically reconstruct the multifractality in the WAN traffic traces, suggesting that the Internet protocol stack hierarchy causes the multifractality. There is a common agreement that the changing nature of the scaling phenomena in small-time scale must be due to "real life" TCP dynamics, specifically, the flow control mechanism, which implies that the end-to-end delay dynamics naturally describe the WAN traffic. So again, better understanding of the delay dynamics can provide deep insight into the physical explanation which would de-mystify the multifractality found in the WAN traffic.

Our research is closely related to the studies carried out in [3, 10]. The study in [3] first explored the long-range dependence (LRD) characteristics of the measured round-trip delay time series. The study in [10] also found the LRD characteristics of the round-trip delay time series measured by NTP. Due to the non-stationary characteristics of the packet delay pro-

*Department of Electrical and Computer Engineering, University of Delaware, Newark, DE19716, USA, email: qli@eecis.udel.edu.

†This work is supported by DARPA Information Technology Office Contract DABT 63-95-C-0046, and U. S. Navy Contract A303 276-93.

cess, the question arises about the conclusions drawn in the studies: whether the LRD-like effects in the delay time series is actually caused by the non-stationary process. Recently, a wavelet-based tool for the analysis of long-range dependence and a related semi-parametric estimator of the Hurst parameter H — the intensity measurement of the LRD effect, was introduced in [1]. The estimator was shown to be unbiased under very general conditions, and efficient under Gaussian assumptions. The most desired feature of the estimator is that it is highly robust against the presence of deterministic trends — possible types of the non-stationarity.

In our research, we use the wavelet-based techniques to study the scaling behavior of the measured delay traces. We also use the wavelet-based multifractal analysis techniques proposed in [6] to investigate the multifractal nature of the delay traces. In order to observe the crossover phenomena, we use another robust and direct technique called *height difference correlation function* method used mostly in surface science[12]. Using these tools we conduct an extensive study on scaling behavior, and provide empirical evidences showing the LRD nature and the possible multifractality of the delay traces measured by ICMP ping and NTP. A significant observation of our research was that we observed a crossover phenomena and an anti-persistence phenomena in the delay trace.

The rest of the paper is organized as follows. Section 2 gives a brief description on the measurements. In section 3, we introduce the mathematical foundations of the wavelet-based tools, and discuss issues related to non-stationarity. Section 4 presents the emperical results of the analyses. Section 5 introduces the height difference correlation function method, and illustrates the crossover phenomena, and the anti-persistence discovered in the delay traces. In section 6 we use a reservoir model to demonstrate the underlying physical mechanisms of the crossover, anti-persistence and LRD of the delay traces, and also discuss the implications of our findings on several issues related to networking engineering. Our conclusions are presented in section 7.

2 Measurements

We believe it is very important to study the packet delay dynamics in different time scales since a protocol or application may only use a sampled delay time series from an underlying random process, possibly not in a regular sampling basis. All the data used in this paper are either measured by NTP [13] or ICMP ping [17].There are some concerns that the ping measurements may not realistically reflect the statistics of the delays experienced by packets issued from a higher layer than IP. NTP is built on top of UDP, so the purpose of using the NTP measurements in our analysis is simply to generalize the findings of this paper.

The NTP delay (file name au_ntp_74sec) trace used in this paper were collected in September, 1996 by running NTP in a machine in our Lab. This machine is configured to make measurement with a primary NTP time server located in Australia. For more detail about the experiment we refer reader to [10]. We run a ping program in the same machine in our lab. All measurements were done in December, 1998. Table 1 gives a qualitative description of the set of data.

We use the smallest ping packet (36 bytes: 20 bytes for IP head, 8 bytes for ICMP head, and 8 bytes for the two timestamps) in our experiment in wish to minimumize the interference to other users. In the original raw traces, some pings were lost or came back out of order. If a ping was lost we do noting for it . We sorted the raw traces to put every received ping back to its sending order. Although there are some doubts about ping's fidelity to represent the packets finally going up or down to the applications, We believe it serves the purpose in this study.

3 Wavelet-based Analysis Techniques

In this section we give a brief introduction to the wavelet-based techniques (for more detail, see [1, 6]), and discuss the capability of the techniques to detrend the deterministic types of non-stationarity. We start this section by introducing some concepts we meet in this study.

3.1 Miscellaneous Concepts

If stochastic signal $X(t)$ has a power-law spectra as $f \to 0$, we call $X(t)$ a $1/f$ *process* or $1/f$ *noise*. The spectra property of a $1/f$ process indicates a slowly-decaying correlation function, So $X(t)$ is a *long-range dependence* process. A popular model for a long-range dependence process is the *fractional Gaussian (or Brownian) noise* derived from the so-called *fractional Brownian motion* (in short fBm) process [11]. The fBm is a non-stationary process with a property called *self-similarity*. Let $B_H(t)$ be a fBm process, and $B_H(0) = 0$, then the self-similarity can be expressed by:

$$B_H(at) \stackrel{\mathrm{d}}{=} a^H B_H(t) \tag{1}$$

for any $a > 0$, where $\stackrel{\mathrm{d}}{=}$ denotes equality of finite-dimensional distribution, H is called Hurst parameter, and $0 < H < 1$. Self-similarity (more precisely, self-affinity) indicates that the graph $(t, B_H(t))$ remains statistically unchanged when the time axis and the amplitude are simultaneously scaled by a factor a and a^{-H}, respectively. Or, we say fBm has the same *scaling behavior* on all time scales. We call the sample paths of fBm the *fractal* curves, because if we calculate the fractal dimension on these curves, we could get one, such as the Hausdorff dimension [5], with value $dim_H = 2 - H$.

The fractional Gaussian noise $\{G_{H,\delta}(t), (t,\delta) \in \mathcal{R} \times \mathcal{R}_+\}$ is the incremental process of fBm, defined by:

$$G_{H,\delta}(t) = \frac{1}{\delta}(B_H(t+\delta) - B_H(t)) \tag{2}$$

Table 1: Qualitative description of the data measured by ping

File Name	Sampling Interval	Measurements	Minimum Delay(ms)	Loss rate	Destination
ar_ping_20ms	20ms	60000	413.2	13%	Argentina
ar_ping_1s	1s	180000	407.3	4.7%	Argentina
ar_ping_1min	1min	10000	589.1	3.5%	Argentina
au_ping_20ms	20ms	60000	445.8	0.65%	Australia
au_ping_1s	1s	180000	487.881	0.49%	Australia
au_ping_1min	1min	10000	421.3	2.0%	Australia
uk_ping_20ms	20ms	60000	124.4	6.46%	England
uk_ping_1s	1s	180000	123.3	0.79%	England
uk_ping_1min	1min	10000	124.3	0.97%	England

where $G_{H,\delta}(t) \stackrel{d}{=} \mathcal{N}(0, \sigma\delta^{H-1})$[1]. With time lags large enough, Fractional Gaussian noise has a correlation function as the following,

$$\mathbf{E}[G_{H,\delta}(t+\tau)G_{H,\delta}(t)] \sim \sigma^2 H(2H-1)|\tau|^{2H-2} \tag{3}$$

and a $1/f$ power spectrum density at low frequencies, as:

$$\Gamma_{B_{H,\delta}}(\nu) \sim |\nu|^{1-2H}, 0 < |\nu| \ll \delta^{-1}. \tag{4}$$

so when $|\nu| \to 0$, $\Gamma_{B_{H,\delta}}(\nu) \to \infty$.

In reality, the stochastic description of physical phynomenon are rarely exactly self-similar, i.e., (1) holds for all time scales (or for any $a > 0$) with a constant H. Usually, with a constant H, (1) only holds for some range of time scales between inner and outer cutoff time scales. Very often, we could find that (1) only holds piecewise in small ranges of a with different constant Hs. In this case we need to check if this random process has a *multifractal nature*.

3.2 Discrete Wavelet Transform and Scaling Analysis

Discrete Wavelet Transform (DWT) is a signal analysis tool [4]. There exists a function ψ_0, called the *mother wavelet*, where a set of base functions can be derived from the scaled and shifted versions of ψ_0, such as $\{\psi_{j,k}(t) = 2^{-j/2}\psi_0(2^{-j}t - k), (j,k) \in \mathcal{Z}\}$. Using discrete wavelet analysis, $X(t)$ can be decomposed as the follow:

$$X(t) = \sum_j \sum_k d_X(j,k)\psi_{j,k}(t) \tag{5}$$

Where $d_X(j,k)$ can be computed by driving $X(t)$ through a fast recursive filter-bank-based pyramidal algorithm. The mapping from $X(t)$ to the set of coefficients $d_X(j,k)$ is the DWT.

Decomposition can be stopped at any stage $J > 0$ by using an approximation of $X(t)$ at scale 2^J to collectively represent the sum of the left wavelet decomposition terms; namely, (5) can be rewritten as:

$$X(t) = approx_J(t) + \sum_{j=1}^{J} detail_j(t)$$

$$= \sum_k a_X(J,k)\phi_{J,k}(t) + \sum_{j=-\infty}^{J} \sum_k d_X(j,k)\psi_{j,k}(t) \tag{6}$$

where $\phi_{J,k}(t)$ is called the *scaling function*, which is derived from a base function $\phi_0(t)$ by the templates $\{\phi_{j,k}(t) = 2^{-j/2}\phi_0(2^{-j}t - k), (j,k) \in \mathcal{Z}\}$. Equation 6 is called the multi-resolution analysis (MRA) of $X(t)$. Both $\phi_0(t)$ and $\psi_0(t)$ are not arbitrary; their constructions are based on certain rules, which require that $\psi_0(t)$ has at least one vanishing moment; i.e.,

$$\int_{-\infty}^{+\infty} t^m \psi_0(t)dt = 0, \quad m = 0, 1, ..., M-1, \tag{7}$$

where $M > 0$. Compactly supported wavelets with M vanishing moments and a support domain equal to the interval $[-(M-1), M]$ were provided in [4]. The higher the order of the vanishing moments of the wavelet, the better its detrending capability when used in the long-range dependence analysis.

It is shown in [1] that, if $X(t)$ is the incremental process of a self-similar process[2] with Hurst parameter $H \in (0.5, 1)$, then the expectation of the energy E_j that lies within a given 2^{-j} scale around frequency $2^{-j}\lambda_0$ (λ_0 is a reference frequency which depends on the wavelet $\psi_0(t)$) is given by

$$\mathbf{E}[E_j] = \mathbf{E}[\frac{1}{N_j}\sum_k |d_{j,k}|^2] = c|2^{-j}\lambda_0|^{1-2H}], \tag{8}$$

where c is a prefactor that does not depend on j, and N_j is the number of wavelet coefficients at scale j. By plotting $\log_2 E_j$ against j (in practice, $j = 1$ is the finest scale and

[1] $\mathcal{N}(m, \sigma)$ denotes a normal law with mean m and variance σ^2.

[2] In many studies, $X(t)$ is directly considered self-similar. In our study we need to differenciate the concept of being self-simialr (e.g. fBm) from that of being the incremental process of a self-similar process (e.g. fractional Gaussian noise).

$j = \lfloor \log_2(length(X(t)) \rfloor$ is the coarsest), we can see that for the incremental process of an exactly self-similar process, it will yield a linear fit, and for a white noise process the fit line will be horizontal. If we can not fit the plot by a line for all time scales, that means the underlying process is not the incremental process of an exactly self-similar process. It demonstrates a more complicated scaling behavior. If the linear relationship between $\log_2 E_j$ and scale j hold for large time scales, it suggests the incremental process of an asymptotically self-similar process. Performing a least-squares fit on the linear portion of the energy-scale plot[3], we can estimate the Hurst parameter H, so the scaling analysis also serves as an H estimator.

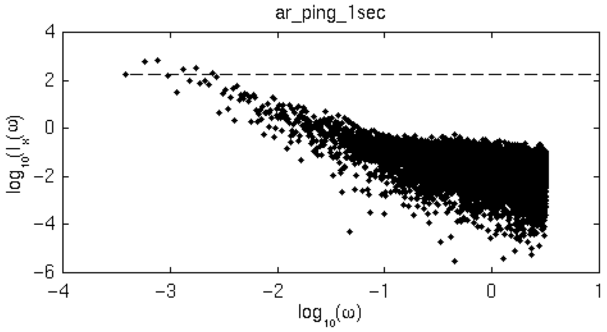

Figure 1: Periodogram of trace ar_ping_1sec.

3.3 Multifractal Analysis

The distinction between *multifractal* and *monofractal* can be clarified by the follow informal definition [6]: $X(t)$ is a random signal, we say that $X(t)$ has a local scaling exponent $\alpha(t_0)$ at time t_0 if $X(t)$ behaves like $(\delta t)^{\alpha(t_0)}$ in an interval $[t_0, t_0 + \delta t]$ of length δt at t_0 as $\delta t \to 0$. Informally, $X(t)$ with $\alpha(t_0) = H < 1$ at all instants t_0 is called monofractal (and includes the incremental process of an exactly self-similar process), while $X(t)$ with non-constant scaling exponent $\alpha(t_0)$ called multifractal.

A wavelet-based multifractal analysis technique is proposed in [6]. It is shown that, if $X(t)$ has a local scaling exponent $\alpha(t_0)$ at t_0, then for large negative j-values (small scales), the wavelet coefficients affected by $X(t_0)$ behave like $d_{j,k}(t_0, j) \approx 2^{j(\alpha(t_0)+1/2)}$ [4], where for two functions f and g, $f(j) \approx g(j)$ means that $\lim_{j \to -\infty} f(j)/g(j) = const.$. The core of the technique is first to construct the so called wavelet-based *partition function* $S(q, j)$ as

$$S(q, j) = \sum_k |\tilde{d}_{j,k}|^q, \tag{9}$$

where $\tilde{d}_{j,k} = 2^{-j/2} d_{j,k}$, the normalized wavelet coefficient. It can be shown that $S(q, j) \approx 2^{-j(1-\alpha q)}$ for the case of a monofractal signal with a constant scaling exponent α. Then it defines a *structure function* $\tau(q)$ as the scaling exponent of $S(q, j)$, as $j \to -\infty$ for a general case, we wish that the partition function would behave like $S(q, j) \approx 2^{j\tau(q)}$, that is,

$$\tau(q) = \lim_{j \to -\infty} \frac{\log S(q, j)}{j \log 2} \tag{10}$$

For a monofractal signal $\tau(q) = \alpha q - 1$ (or $Hq - 1$), $\tau(q)$ changes linearly with q. In the multifractal case, $\tau(q)$ is no longer linear in q but instead a concave function of q.

3.4 Robustness of the Wavelet-based Analyses

If delay traces span over hours or days it may raise a question whether or not the LRD-like effect revealed by certain statistical techniques is actually caused by non-stationarity in the delay process. We address this question by two steps.

First, in Figure 1 we draw the periodogram of trace ar_ping_1sec, that is

$$I_X(\omega) = (2\pi n)^{-1} |\sum_j X_j e^{ij\omega}|^2, 0 \le \omega \le \pi.$$

The expected value of $I_X(\omega)$, which is the spectral density of the trace, is obviously different from that of a white noise process, which would take the dashed line shown in the log-log scale figure as its expectation value. That the spectral density obeys a power-law near the origin suggests a possible $1/f$ process.

Secondly, if the diurnal changing pattern can be modeled by several low frequency components[14], then the wavelet-based techniques can perfectly suppress the low frequency diurnal changing pattern disturbing the scaling analysis results. To prove this, we add several low frequency components to the re-shuffled version of trace ar_ping_1sec. After completely re-shuffling, the reshuffled trace is simply a white noise with the same first order statistics as the original trace. The modified trace after adding the diurnal components looks like

$$X_j^{'} = X_j + \frac{\overline{X}}{4} \sum_{i=1}^{4} \sin(\frac{2\pi i j}{n}),$$

where X_j is the re-shuffled time series, $\overline{X}$ is the mean of X_j, and n is the length of the trace. Then, we can use the wavelet-based scaling analysis technique implemented in Matlab to explore the scaling behaviors of both the re-shuffled trace X_j and the modified trace $X_j^{'}$. We find that, with a Daubechies2 [4] wavelet (which has a vanish moment N = 2), the technique can not completely suppress the sinusoidal components. But, changing to a Daubechies4 wavelet (N = 4), the technique can perfectly suppress the sinusoidal components, so it can detrend the artificial non-stationarity and reveal a scaling behavior similar to that of the re-shuffled trace. The results are shown in Figure 2. It is not surprising to find that both of them are horizontal, implying a Hurst parameter $H = 0.5$.

[3] In this study we call the $\log_2 E_j$ vs. j plot as the energy-scale plot.

Generally, with a larger N, the wavelet would have a better suppressing ability, so the detrending ability of the scaling analysis technique. Besides the detrending capability we discuss here, the wavelet-based technique can actually suppress a broader range of deterministic trends such as linear trends, polynomial trends and arbitrary smooth trends[1, 4]. Throughout this study we choose the Daubechies4 wavelet to detrend the diurnal changing pattern presented in the delay traces.

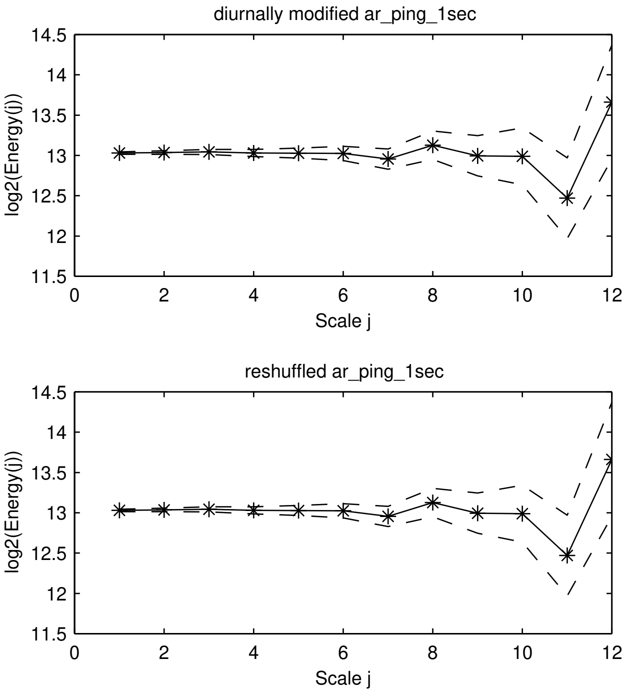

Figure 2: Scaling analyses of both the re-shuffled and modified version of trace ar_ping_1sec.

The above experiments help us conclude that, with a carefully chosen wavelet, the white noise like process (such as the re-shuffled trace X_j) even with strong deterministic trends can not fake a false scaling behavior to confuse the analysis result. So if the wavelet-based analysis really picks up some obviously scaling features, they must be caused by the underlying random process. We believe that the wavelet-based analysis is robust against the diurnal trends as long as they are smoothly changing trends.

4 Wavelet-based Analyses of the Delay Traces

In this section we present the analysis results by applying the wavelet-based techniques to the delay traces.

4.1 Scaling Analysis of the Delay Traces

We investigated the nine ping traces and the NTP trace. Figure 3 shows the scaling analysis results of the nine ping traces. Visually, all the nine energy-scale plots show increasing trends against scales, which suggest some kind of long-range dependence phenomena exist in the traces. Roughly speaking, traces ar_ping_20ms, ar_ping_1sec, and uk_ping_20ms may be characterized as the incremental time series of self-similar process, and traces au_ping_1sec and uk_ping_1sec the incremental time series of asymptotically self-similar process. Trace au_ping_20ms, ar_ping_1min, and au_ping_1min show strong long-range dependency, but hard to fit into either category. Trace uk_ping_1min shows a very weak long-range dependence, so we may treat it as a white noise. Obviously, several of the energy-scale plots do not show a strict linear relationship between $\log_2 E_j$ and j in ranges of either large or small scales, implying a more complicated underlying process. Table 2 gives the estimated Hurst parameter calculated by using a weighted least-squares fit on the roughly linear portions of the energy-scale plots. Generally, H decreases as the sampling interval of a trace increases.

For the purpose of comparison and generalizing the results of scaling behavior found in the ping traces, we also applied the wavelet-based scaling analysis technique to the NTP trace. The result is shown in Figure 4. The energy-scale plot of the NTP trace clearly indicates that it is the incremental time series of an asymptotically self-similar process. This analysis confirms the conclusion obtained in one study [10] which uses the variance-time plot to investigate NTP traces. So, we may conclude that the scaling behavior observed in the ping traces can also be observed in delay traces measured from the application layer.

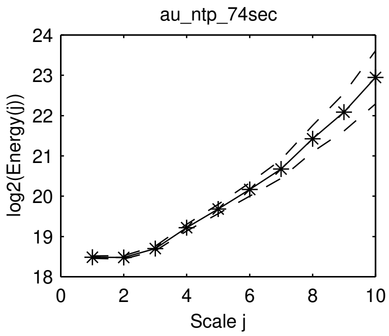

Figure 4: Scaling analysis of the NTP trace.

4.2 Multifractal Analysis of the Delay Traces

Using the methodology proposed in [6], we took trace au_ping_1sec and uk_ping_1sec as examples, calculated their wavelet-based partition function $S(q, j)$'s respectively, and their corresponding wavelet-based structure function $\tau(q)$.

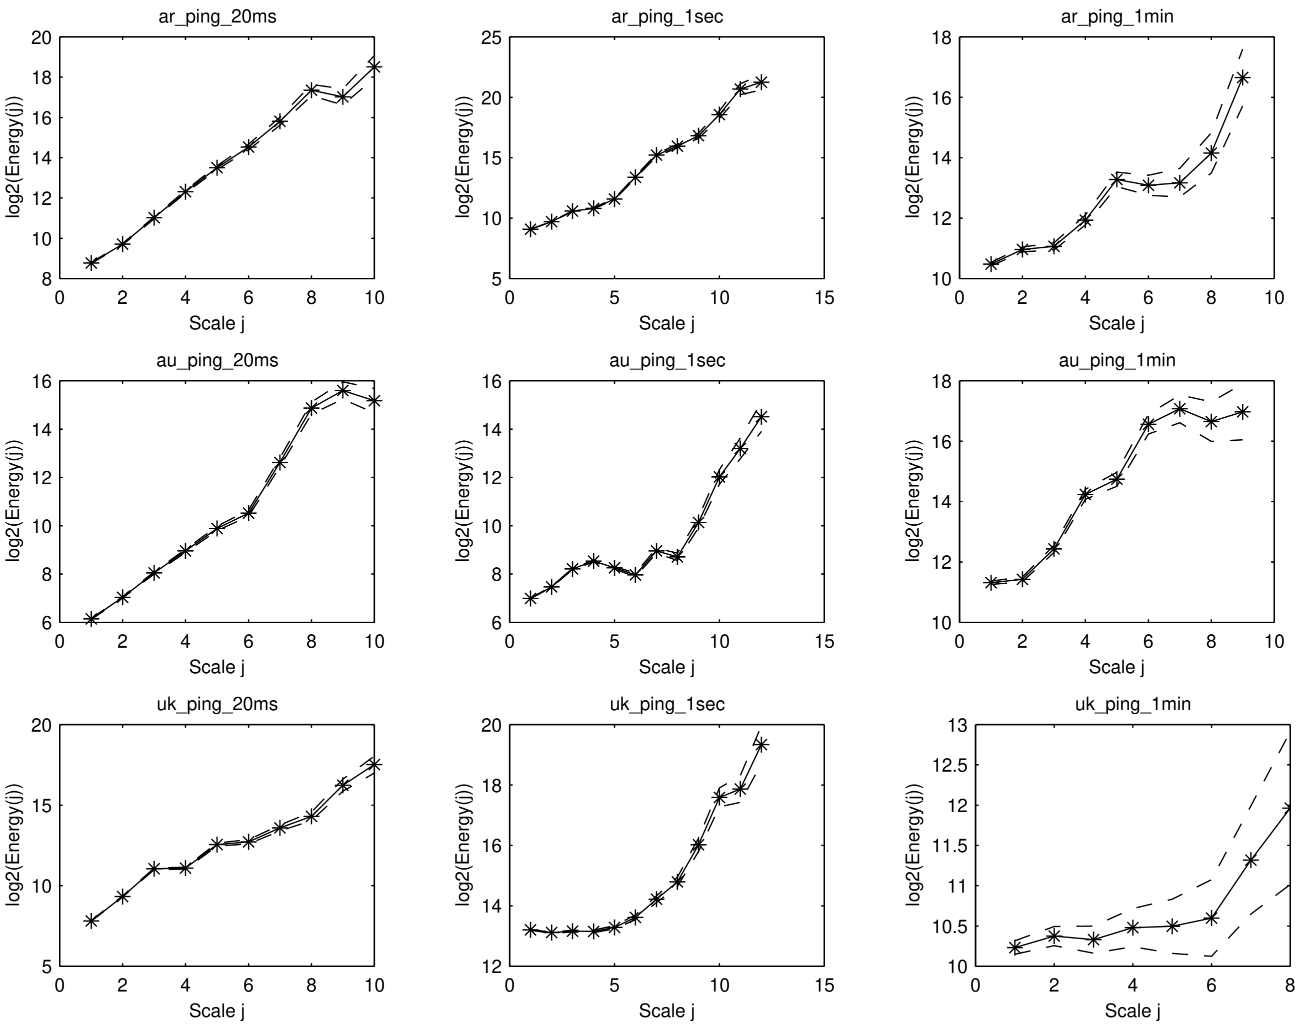

Figure 3: Scaling analyses of the nine ping traces.

See the introduction of this technique in Section 3.3.

Figure 5 shows the results of applying the wavelet-based multifractal analysis method to the two traces. We implemented the algorithm in Matlab and took advantage of the wavelet analysis package already developed in Matlab, which always takes $j = 1$ as the smallest scale index. In our case $j = 1$ is the smallest scale in the $S(q, j)$ plots[4]. Examining the scaling behavior of the partition function $S(q, j)$ in small scales of the two traces in Figure 5 (left), we do not see well-defined scaling regions for large q's. This situation makes the calculation of the structure function $\tau(q)$ difficult and ambiguous, which is normally computed for each q by a linear least-squares fit on the fine-scale portion of $S(q, j)$, and equals the slope of the linear fit. Despite this situation, we still calculated the $\tau(q)$ based on the portion of the $S(q, j)$ in the scales ranging from 1 to 4, and showed the result in the middle of Figure 5. As discussed in [6], if the underlying process is monofractal, $\tau(q)$ will behave like a linear function, where the slope of the $\tau(q)$ is the scaling exponent equal to the Hurst parameter. But for a multifractal process, $\tau(q)$ will diverge from a strict linear function. In order to help visually check the linearity of the $\tau(q)$ plots, we calculated the local slope of $\tau(q)$, and identified the slope as α. The results are shown in Figure 5 (right). The α functions clearly indicate that the $\tau(q)$ functions are not linear, which suggest that the examined traces may be more multifractal in nature than monofractal.

[4] We have adjusted the calculation of the $S(q, j)$ accordingly, which becomes

$$S(q, j) = 2^{-(J+1)} \sum_k |\tilde{d}_{j,k}|^q,$$

where $J = \log_2 N_j$, N_j is the total number of the DWT coefficients $d_{j,k}$, the same meaning as in the scaling analysis. The purpose of adjusting the calculation of $S(q, j)$ is simply to shift the scale index j in (9) from $j = -J$ to $j = 1$. So, the validity of (10) is preserved with the Matlab's wavelet analysis package.

interval	_20ms	_1sec	_1min
ar_ping	1.05 ± 0.03	0.84 ± 0.03	0.72 ± 0.05
au_ping	0.97 ± 0.03	0.76 ± 0.03	0.78 ± 0.05
uk_ping	1.20 ± 0.03	0.63 ± 0.03	0.54 ± 0.05

Table 2: Hurst parameters H for the ping traces

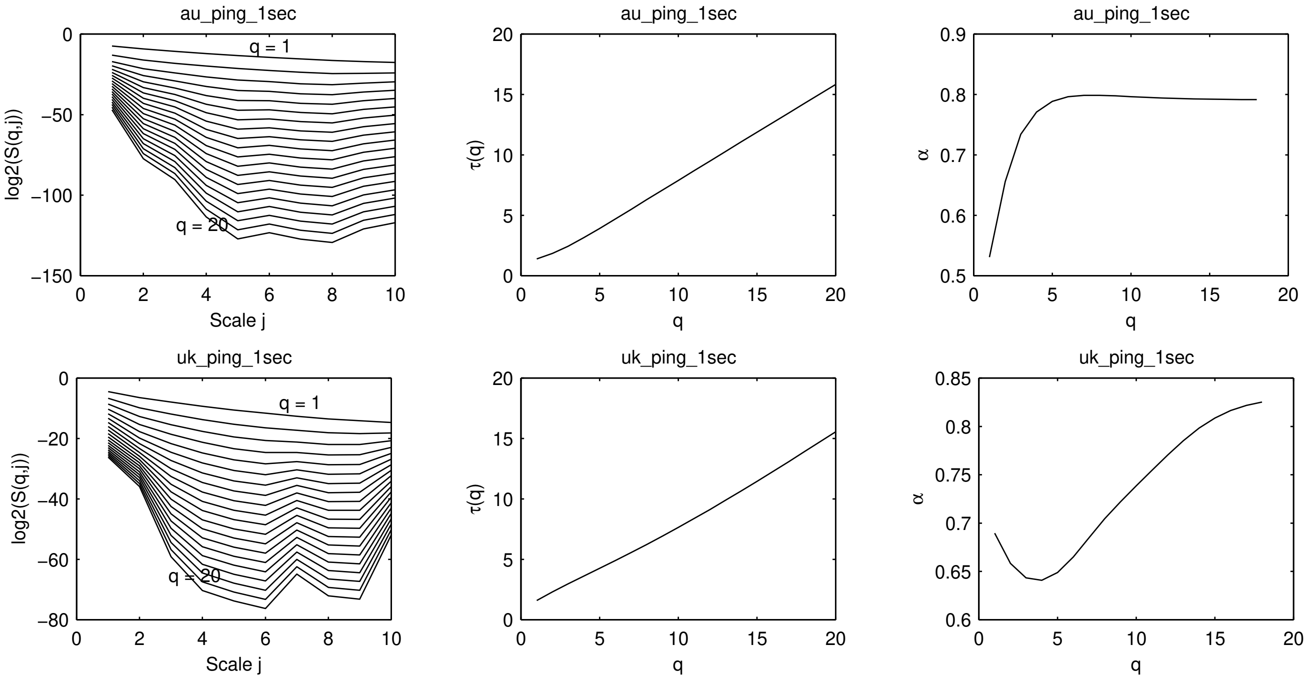

Figure 5: Multifractal analyses of trace au_ping_1sec and uk_ping_1sec.

5 Crossover

Since the values of H in Table 2 for trace ar_ping_20ms and uk_ping_20ms are > 1.0, we began to suspect that in these two cases modeling the delay traces as the fractional Gaussian noises may be incorrect. By directly modeling a delay trace as a fBm, we found the *crossover* phenomena. Crossover simply means here that the characteristics of delay time series changes from fractional Gaussian noise to fBm when the sampling interval decreases to a certain limit. In this section we first introduce the techniques we used in conducting the study and then show the analysis results.

5.1 Height Difference Correlation Function

One of the most robust and direct approaches to characterizing the self-similar curves (or signals) is by defining the *height difference correlation functions* as

$$C_q(\mathbf{t}) = \{\mathbf{E}[(\delta X(t))^q]\}^{\frac{1}{q}} = \{\mathbf{E}[(X(\mathbf{t}_0 + \mathbf{t}) - X(\mathbf{t}_0))^q]\}^{\frac{1}{q}} \tag{11}$$

where $X(t)$ is a fractal curve or signal and usually $q = 2$. It has been shown [12] that for a self-similar signal (e.g., fBm), $C_q(\mathbf{t})$ has the form

$$C_q(\mathbf{t}) \sim \mathbf{t}^H, \tag{12}$$

where H is the Hurst parameter of the signal. In practice, in many case (12) only holds within certain scales of $\mathbf{t}$, or in other words, the scaling behavior of $X(t)$ described by (12) is only valid in a regime bounded by lower and upper correlation lengths as $\xi^- < \mathbf{t} < \xi^+$. In this case we would say $X(t)$ is a self-similar (or precisely self-affine) signal with an *inner cutoff scale* ξ^- and an *outer cutoff scale* ξ^+. By plotting $\log C_2(\mathbf{t})$ against $\log(\mathbf{t})$, we can estimate the Hurst exponent H for $X(t)$. Also, by examining the scaling behavior of $\log C_2(\mathbf{t})$, we can find the inner cutoff scale and the outer cutoff scale of $X(t)$.

5.2 Crossover Analysis and Anti-persistence

Figure 6 shows the results of applying the height difference correlation function method to trace ar_ping_20ms, au_ping_20ms, and uk_ping_20ms. Surprisingly all the plots show a piecewise linear relationship between $\log C_2(\mathbf{t})$ and $\log(\mathbf{t})$, which suggest that the delay traces under question are self-similar time series (modeled by fBm) with outer cutoff scale. In Figure 6 we can easily identify the break points on the $\log C_2(\mathbf{t})$ vs. $\log(\mathbf{t})$ plots. These break points correspond to the cutoff scales, where $\mathbf{t} = 1.5$s, 1.5s, and 100ms separately for each trace. From each plot we can estimate two Hurst exponents based on the portions of the plot before and after the break point. The values are listed in Table 3.

In Table 3 we see that the Hurst exponents of all three traces have big step-style changes before and after the first break points. Before the break points, the three traces all have a Hurst exponent in the range of $0 < H < 0.5$, which suggests the three traces can be characterized by self-similar time serieses(e.g. modeled by fBm) with the break points as the outer cutoff scales. The inner cutoff scales remain an open question. After the break point, the three traces all have a Hurst exponent $H \approx 0.0$, which means that if we take a sub-trace from an original trace with an interval between the two consecutive samples in the sub-trace larger than the break point (or the cutoff scale) of the original trace, then we will see a more noise-like time series. So a delay time series with a sampling interval smaller than the break point of an end-to-end Internet path, can be modeled by a self-similar process (e.g., fBm) with an outer cutoff scale equal to the break point. If the sampling interval is larger than the break point, the delay trace can then be modeled by an LRD process (e.g., fractional Gaussian noise) which is an incremental process of a self-similar process. The changing nature of the delay trace around the break point is referred as the *crossover* phenomena.

Since a delay trace with a small sampling interval can itself be modeled by a self-similar process, it implies that the delay jitter time series derived from the delay trace can be modeled by a LRD process (e.g. fractional Gaussian process) with a Hurst parameter equal to the Hurst exponent of the delay trace. From Table 3 we see that all the Hurst exponents are less than 0.5, suggesting that the derived delay jitter time series are all anti-persistence LRD processes (with $H < 0.5$, the value of the correlation function of a fractional Gaussian noise will be < 0; see (3)). The physical meaning of the anti-persistence is that, if the delay jitter in the past are positive, then the delay jitter in the future trend to be negative, and vice versa. This further implies that, if the delays increase in the past, then the delays will trend to decrease in the future.

6 Discussion

The results in the previous sections provide a new and improved understanding of the Internet packet delay dynamics. But, these results also raise fundamental questions like: "Why the delay dynamics has a crossover" and "Why the anti-persistence exists in the delay dynamics". In this section we present our preliminary thoughts on these questions, and discuss some issues related to traffic modeling, networking protocol design, monitor and diagnosis.

6.1 On the Reason of Crossover and Anti-persistence

In order to understand why the delay trace shows an anti-persistence, self-similar nature in small-scale, while a LRD noise nature in large-scale, we use an analog to explain the reason. Hurst, who pioneered LRD phenomena, spent his lifetime studying the Nile and the problems related to water storage [7]. He found in the case of a reservoir that the volume of the water stored in the reservoir $X(t)$, which is decided by the influx $\xi(t)$ and the discharge flow $\theta(t)$ as

$$X(t) = \int_{-\infty}^{t} (\xi(t) - \theta(t))dt \tag{13}$$

If the reservoir was never emptied, Hurst found $X(t)$ showing a self-similar nature with a $H > 0.5$. The reason for this is believed to be that there exists a long "memory" in the influx $\xi(t)$ over time. The influx of a reservoir is not simply determined by the recent average precipitation, but by the water content in a large drainage area. The precipitation is first stored in the drainage area by many means. The large drainage area serves as a long memory filter to the rainfall, and the influx of the reservoir is the output of the filter. A reservoir can serve as a perfect model for a node in Internet: The influx to a reservoir corresponds to the incoming traffic to a node, the discharge flow of a reservoir corresponds to the output traffic served by a node, the volume of water stored in a reservoir corresponds to the queue length in a node. So, if the incoming traffic has a long memory, such as a LRD traffic, then it is not surprising that the queue length is self-similar in the small-time scale during which the queue in a node has less chance to be completely drained. This analog helps us understanding why the delay traces with a small sampling interval can be modeled by a self-similar process.

This analogy also suggests that the crossover is mainly caused by the events of queue empty and queue overflow. When these events happen, the reservoir model is broken, so the exactly self-similar nature of the queue length is disregarded. This explains the crossover in the delay traces. When a delay trace has a sampling interval larger than the cutoff scale where the crossover occurs, the Internet node's reservoir-behavior will have little effect on the dynamics of the delay trace. The dynamics of the delay trace will be determined directly by the large scale, low frequency dynamics of incoming traffic. The incoming traffic is well-known to be LRD. This can explain why the delay trace with a large sampling interval, even as large as 1 min, still shows the LRD nature.

In the case of the reservoir, the water volume shows a persistence self-similar, but in the case of the Internet node, the queue length shows an anti-persistence self-similar. We believe the reason for this is that, in the case of a reservoir, there is no feedback mechanism to control the precipitation based on the volume of water in the reservoir. But, in the case of Internet node, the incoming traffic is controlled by feedback via

break point	ar_ping_20ms	au_ping_20ms	uk_ping_20ms
before	0.05	0.00	0.07
after	0.23	0.18	0.37

Table 3: Hurst parameters H around the break points of the ping traces

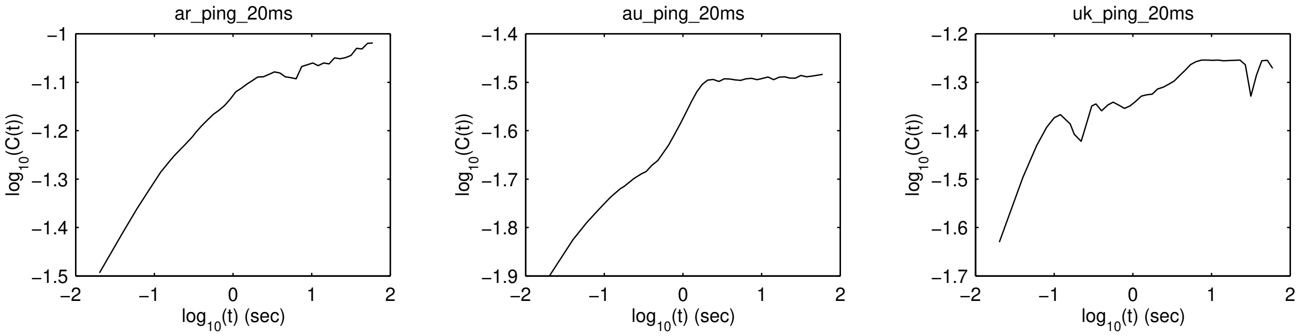

Figure 6: Height difference correlation function analyses of trace ar_ping_20ms (a), au_ping_20ms (b), uk_ping_20ms (c).

the transport layer. TCP uses a window flow control mechanism to regulate traffic. The window is shifted by acks which are paced by the round-trip delays, and essentially paced by the queue lengths of the nodes in the path. So the feedback flow control features of transport protocols are very important to network performance, it is the fundamental reason for the anti-persistence in today's Internet delay traces. The anti-persistence in the delay trace is a good indication that the current Internet is working toward a stable state.

6.2 Implications on the Traffic Model

We think that the scaling behavior of WAN traffic in small-time scales should be disciplined and reflected by the scaling behavior of the delay traces. With TCP as an example, we know that it uses acks pacing the traffic, which is referred as *self-clocking*. The incoming process (or the inter-arrival process) of acks is essentially modulated by the delay jitter process with a sampling interval corresponding to the round-trip delay. It seems that the round-trip delay of a path is still within the cutoff time scale of that path, so the incoming process of acks must be a fractal-like process, which means TCP uses a clock ticking in a fractal way to strobe new packets into the network. So the dynamics of TCP traffic in small-time scale is determined by a *fractal-clock*. This fractal-clock presented may explain the scaling behavior and multifractality observed in the WAN traffic in small-time scales.

The delay jitter process actually reflects the changing nature of Internet workload [2], since the delay jitter process is anti-persistent and the anti-persistence phenomena is robust in delay traces. We predict that the WAN traffic dynamics in small-time scale is actually anti-persistent, too. So, anti-persistence may be another network traffic invariant. This prediction needs to be confirmed by future study.

6.3 Outlook

The findings of crossover and anti-persistence not only can improve our understandings about the current Internet dynamics, but also can be used in many aspects of network engineering. One direct application of our findings is the prediction of delay and delay jitter for Internet protocols and real-time applications. Since the height difference correlation function of a delay trace has a break point, after the break point the function becomes flat (see Figure 6). Here, the so called "height difference" is actually the delay jitter of two delays separated by the correlation lag time. The fact that the height difference correlation function becomes flat implies the standard deviation of the delay jitter between any two delays has an upper boundary. We need to investigate how robust the upper boundary is. If it is robust against the time of day, the upper boundary of the standard deviation of the delay jitter will be very important to the designs of the protocols and applications running on the Internet. It opens a new opportunity to refine the delay prediction algorithms implemented by many protocols and applications. More study needs to be done in this area.

It seems a simple ping trace can tell us more than we thought previously. But, we need further understand what aspects of the network decide the location of the break point, the Hurst exponent, and the boundary of the height difference correlation function. We think all these parameters depend on the network state, workload pattern, loss rate, total buffer storage and the bottleneck bandwidth of a path. The height difference correlation function can servers as a kind of fingerprint for a path. By analyzing the height difference correlation function based on the ping traces, network administrators can understand more about the network state. Further study is needed to relate the network aspects to the parameters of height difference correlation functions.

The anti-persistence shown in the delay traces suggest the network is under control. But it raises a question:"Does there exist an optimum control point?" If $H = 0.5$, the network has no control, the delays in small-time scale behave like a *Brownian motion*, but if $H = 0$, the network is either over controlled, or the load in a path is too light to support a fBm in the nodes. In the case of over control, the network must have enforced a harsh policy on the upper bound of the standard deviation of the delay jitter. For a network provides best-effort services, there may be a reasonable range of the H value. So how to judge the network's performance in the macroscopic level based on the H value is an interesting topic.

7 Conclusions

The long-range dependence (LRD) nature of Internet end-to-end packet delay dynamics was confirmed by this study. By using the wavelet-based scaling analysis technique, partition function method and the height different correlation function method, we conducted an extensive investigation on the dynamic properties of packet delay traces measured using ICMP ping and NTP. Our main observations are: (i) the scaling analysis shows a very complicated scaling behavior of the delay traces and demonstrates that the LRD nature of the delay traces is robust, even for traces with sampling interval as large as 1 min; (ii) some delay traces show possible multifractality, (iii)a delay trace with a large sampling interval can be modeled as an LRD noise-like time series, while with a small sampling interval a trace should better be modeled as a self-similar time series with an outer cutoff time scale. The changing of the nature of the delay traces is referred to as crossover; (iv) by applying the height difference correlation function method to the delay traces with small sampling interval, we found the anti-persistence phenomena, which suggests that the current Internet is working in a controlled state.

References

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

[2] J.-C. Bolot. Characterizing end-to-end packet delay and loss in the internet. *Journal of High Speed Networks*, 2:305–323, 1993.

[3] M. S. Borella, S. Uludag, G. B. Brewster, and I. Sidhu. Self-similarity of internet packet delay. In *Proc. IEEE ICC'97*, August 1997.

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

[5] J. Feder. *Fractals*. Plenum Press, New York, 1988.

[6] A. Feldmann, A. C. Gilbert, and W. Willinger. Data networks as cascades: Investigating the multifractal nature of internet wan traffic. In *Proceedings of ACM SIGCOMM'98*, pages 42–55, Vancouver, British Columbia, Canada, September 1998.

[7] H. E. Hurst, R. P. Black, and Y. M. Simaika. *Long-term storage: an experimental study*. Constable, London, 1965.

[8] V. Jacobson. congestion avoidance and control. In *Proceedings of ACM SIGCOMM'88*, pages 314–329, Stanford, CA, 1988.

[9] W. E. Leland, M. S. Taqqu, and D. V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.

[10] Q. Li and D. L. Mills. On the long-range dependence of packet round-trip delays in internet. In *Proceedings of IEEE ICC'98*, volume 2, pages 1185–1191, Atlanta, GA, 1998.

[11] B. B. Mandelbrot and J. W. Van Ness. Fractional brownian motions, fractional noises and applications. *SIAM Review*, 10(4):422–437, October 1968.

[12] P. Meakin. *Fractal, scaling and growth far from equilibrium*. Cambridge University Press, 1998.

[13] D. L. Mills. Internet time synchronization: the network time protocol. *IEEE Transactions on Communication*, 39(1):1482–1493, 1991.

[14] A. Mukherjee. On the dynamics and significance of low frequency components of internet load. *Internetworking: research and experience*, 5:163–205, 1994.

[15] V. Paxson. End-to-end internet packet dynamics. In *Proceedings of ACM SIGCOMM'97*, pages 139–152, Cannes, France, 1997.

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

[17] W. Richard Stevens. *TCP/IP Illustrated*, volume 1. Addison Wesley, 1994.

[18] 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 and applications*. Birkhauser, Boston, 1998.