

Small-Time Scaling Behaviors of Internet Backbone Traffic: An Empirical Study

Zhi-Li Zhang,
University of Minnesota
Email: zh Zhang@cs.umn.edu

Vinay J. Ribeiro,
Rice University
Email: vinay@rice.edu

Sue Moon, and Christophe Diot
Sprint ATL
Email: {sbmoon,cdiot}@sprintlabs.com

Abstract— We study the *small-time* (sub-seconds) scaling behaviors of Internet backbone traffic, based on traces collected from OC3/12/48 links in a tier-1 ISP. We observe that for a majority of these traces, the (second-order) scaling exponents at small time scales (1ms - 100ms) are fairly close to 0.5, indicating that *traffic fluctuations* at these time scales are (nearly) uncorrelated. In addition, the traces manifest mostly *monofractal* behaviors at small time scales. The objective of the paper is to understand the potential causes or factors that influence the small-time scalings of Internet backbone traffic via empirical data analysis. We analyze the traffic composition of the traces along two dimensions – flow size and flow density. Our study uncovers *dense* flows (i.e., flows with bursts of densely clustered packets) as the correlation-causing factor in small time scales, and reveals that the traffic composition in terms of proportions of *dense* vs. *sparse* flows plays a major role in influencing the small-time scalings of aggregate traffic.

I. INTRODUCTION

Scaling behaviors of the Internet traffic have a significant impact on network performance and engineering, and thus have been the focus of much research (see [1] and references therein). Self-similar scaling over large time scales (e.g., 1 second and above) has been observed in a variety of network traffic (see, e.g., [2], [3], [4]). More recently, more complex, perhaps multifractal-like, scaling behaviors below 1 second time scales have also been reported [5], [6], [7], [8], [9], [10]. Since queuing inside routers and network congestion are strongly influenced by traffic fluctuations at sub-second small time scales, understanding of *small-time* scaling behaviors of network traffic is critical to many network engineering problems, e.g., router buffer dimensioning, delay-sensitive service provisioning, and congestion control [11], [12], [13]. The need for such an understanding is particularly acute in the backbone Internet with high-capacity links and growing traffic volumes.

In this paper we study the *small-time* (sub-seconds) scaling behaviors of Internet backbone traffic. Our analysis is based on day-long packet traces collected from OC3/12/48 links on a tier-1 ISP (Internet Service

Provider). The high-precision time stamping as well as the high link capacity allow us to zoom into finer time scales (say, 1 ms time scale) and perform reliable data analysis at these time scales. In addition, the large set of long packet traces enables us to make observations over time and compare scaling behaviors of traffic carried across links of various types (e.g., links connecting to other ISPs, or big corporate customers). In contrast, majority of the previous traffic data traces used in the study of traffic scaling behaviors have less than 100 Mb/sec bandwidth, and typically they are relatively short in duration. Our study, we believe, is the first effort to use an extensive amount of data from a commercial tier-1 carrier to study the traffic scaling behaviors *inside the Internet backbone*.

As is previously observed on the Internet WAN traffic, all data traces exhibit a dichotomy of scaling behaviors: above 1 second or so, the traffic has a clear-cut single self-similar scaling, while below it, the scaling behaviors appear to be more *varied*. The transition occurs between 100 ms and 1 second, regardless of link speed, link utilization, link type and time. This observation is not surprising; it merely confirms that what was previously observed on relatively low-speed links also holds on high-speed links. What is striking, however, are the observations regarding the *small-time* (sub-second) scaling behaviors on these Internet *backbone* links. We find that over a range of small time scales, a majority of packet traces manifest (*nearly*) *uncorrelated* scalings with a (second-order) scaling exponent below 0.6, often fairly close to 0.5. This seems to indicate that *traffic fluctuation* on these backbone links appear to be *almost independent* at these small time scales! A small number of traces do exhibit some correlation (with scaling exponents within the range of 0.6-0.7). In addition, via multifractal analysis, we show that all traces manifest mostly *monofractal* behaviors at small time scales.

This paper is devoted to the understanding of potential factors that influence the small-time scalings of Internet backbone traffic via empirical data analysis. We analyze the traffic composition of the traces along two dimensions – flow size (defined in terms of bytes in a flow) and flow *density* (defined in terms of dominant packet inter-arrival times within a flow). We show that

⁰This work was conducted while the first two authors were visiting Sprint ATL. Zhi-Li Zhang was on leave from University of Minnesota, and Vinay Ribeiro was a graduate intern. Zhi-Li Zhang was supported in part by NSF grants ITR-0085824 and CAREER Award NCR-9734428, and by the University of Minnesota McKnight Land-grant professorship.

aggregation of flows with small sizes (*small flows*) generally exhibits mostly uncorrelated small time scaling, while aggregation of large flows in itself does not induce correlated or uncorrelated small-time scaling. Hence *flow size alone* does not determine the small-time scaling of aggregate traffic, unlike traffic scaling behavior at large time scale (above 1 second), which has been shown to be determined by heavy-tailed flow size distribution. However, by analyzing traffic composition along the dimension of flow density, we uncover *dense* flows (i.e., flows with bursts of densely clustered packets) as the primary correlation-causing factor in small time scales. To further investigate how correlation in small time scales is affected by traffic composition, namely, the relative proportions of dense and sparse flows, we perform *semi-experiments* using the traffic traces by varying the proportions of dense and sparse flows. Our results demonstrate that traffic composition in terms of proportions of *dense vs. sparse* flows plays a major role in influencing the small-time scalings of aggregate traffic.

The remainder of this paper is structured as follows. In Section II we describe packet traces used in the study and present the data analysis methodology we use. Observations of the small-time scaling behaviors of the traces are reported in Section III. In Section IV we propose two flow classification schemes and use them to analyze traffic compositions of the traces. The impact of traffic composition on small-time scalings of aggregate traffic is analyzed in Section V. The related work is briefly discussed in Section VI, and the paper is concluded in Section VII.

II. DATA ANALYSIS METHODOLOGY

We first provide a short description of the packet traces and how they are collected. We then present the data analysis methodology used in our study, the second-order scaling analysis and multifractal analysis using wavelets.

A. Packet Traces

For this work, we consider traces collected on OC-3 (155 Mbps) and OC-12 (622 Mbps) links within PoPs (Points-of-Presence) and OC-48 links (2.5 Gbps) between PoPs from a tier-1 backbone network. The tier-1 network covers a wide geographical area and has a variety of link types. Links within a PoP have OC-3 (155 Mbps) and OC-12 (622 Mbps) speed, and inter-PoP links have OC-48 (2.5 Gbps) or OC-192 (10 Gbps) speed. The customers are tier-2 or lower ISPs (Internet Service Providers), corporations, and international ISPs. The packet traces were collected at three PoPs: two in the east coast, and the other in the west coast. Each trace is a sequence of packet records containing the first 40 bytes of a packet, and a GPS-synchronized timestamp (with an accuracy down to 5 μ s) which indicates when

the packet was observed. For further details about the measurement system, see [14].

Table I describes *part of* the traces used in our study that led to our observations highlighted in the introduction. The listed traces are used in this paper as examples to illustrate our observations and analyses.

B. Wavelet Analysis, Energy Plot and Scaling Exponent

We employ wavelet analysis as our scaling analysis tool, which provides robust estimators of the Hurst parameter H while eliminating polynomial nonstationarities in traffic [15]. To introduce the essential terminology for our discussion of observed scaling behaviors in Section III, we provide a quick primer on wavelet analysis using the simplest form of wavelets, the *Haar* wavelets.

Fixing a reference time scale T_0 , define time scale T_j as $T_j = 2^{-j}T_0$. At each time scale T_j (scale j for short), let $t_{j,k} = (kT_j, (k+1)T_j)$ denote the k -th time unit (of length T_j). Consider a (stationary) traffic process Y . At time scale j , $Y_{j,k}$ denotes the amount of traffic (i.e., the total bytes) arriving in time unit $t_{j,k}$. Then $Y_j = \{Y_{j,k}, k \in \mathbb{Z}\}$ represents the traffic process that is observed at time scale j . The Haar wavelet coefficients $W_{j,k}$ of the traffic process at time scale j are defined as follows:

$$W_{j,k} = 2^{j/2}(Y_{j+1,2k} - Y_{j+1,2k+1}). \quad (1)$$

The *energy* function of the traffic process at time scale j is then given by

$$\mathcal{E}_j = E[W_{j,k}^2] \quad (2)$$

where $E[\cdot]$ denotes expectation. The energy function at time scale j is related to the spectral density function (SDF), $\Gamma(\nu)$, of the original signal, i.e., the process Y , via the following relation [15]:

$$\mathcal{E}_j = \int \Gamma(\nu)2^{-j}|\Psi_0(2^{-j}\nu)|^2 d\nu \quad (3)$$

where Ψ_0 denotes the Fourier transform of the mother (Haar) wavelet ψ_0 (at time scale 0). Since Ψ_0 corresponds to a bandpass filter centered at the frequency $2^j\nu_0$, where ν_0 is the frequency corresponding to the reference time scale T_0 , \mathcal{E}_j measures the amount of energy in the process Y in a frequency band around $2^j\nu_0$. The energy function captures the second-order statistics of a traffic process, and its scalings as a function of time scale.

Consider a self-similar fractional Gaussian noise (fGn) process with a *Hurst parameter* H ($0.5 < H < 1$): $Y_{j,2k} \stackrel{D}{=} (1/2)^H Y_{j-1,k}$, $k \in \mathbb{Z}$. Here $\stackrel{D}{=}$ denotes *equality in distribution*. The self-similar scaling behavior of Y implies that its SDF has the following form:

$$\Gamma(\nu) = c|\nu|^{1-2H} \quad (4)$$

where c is some constant. It can be shown [16] that

$$\mathcal{E}_j = 2^{-j(2H-1)}\mathcal{E}_0. \quad (5)$$

| Trace | Link Type/Direction | Start Time | Duration | Speed | Avg. Bandwidth |
|----------------|--------------------------------|----------------|----------|-------|----------------|
| OC3-tier1-dom | to domestic tier-1 peer | Aug. 8, 2000 | 15hrs | OC-3 | 42Mbps |
| OC3-tier2-dom | from domestic tier-2 peer | Aug. 8, 2000 | 23hrs | OC-3 | 44Mbps |
| OC3-corp-dom | from corporate | Aug. 8, 2000 | 19hrs | OC-3 | 28Mbps |
| OC12-tier1-dom | from domestic tier-1 peer | Sept. 5, 2001 | 8hrs | OC-12 | 234Mbps |
| OC12-tier2-int | from international tier-2 peer | Sept. 5, 2001 | 6hrs | OC-12 | 228Mbps |
| OC12-tier2-dom | to domestic tier 2 peer | Sept. 5, 2001 | 7hrs | OC-12 | 187Mbps |
| OC12-corp-dom | from corporate | Sept. 5, 2001 | 20hrs | OC-12 | 122Mbps |
| OC48-bb-1 | inter-PoP link | April 19, 2002 | 1 hr | OC48 | 746Mbps |
| OC48-bb-2 | inter-PoP link | April 19, 2002 | 1 hr | OC48 | 844Mbps |
| OC48-bb-3 | inter-PoP link | April 19, 2002 | 1 hr | OC48 | 1356Mbps |

TABLE I
TRACE DESCRIPTION

Hence $\log \mathcal{E}_j = -j\alpha + \log \mathcal{E}_0$, where $\alpha = 2H - 1$. Namely, $\log \mathcal{E}_j$ scales linearly with time scale $-j$.¹ The slope of a simple plot of $\log \mathcal{E}_j$ against $-j$, called the *energy plot*, provides an estimate of H . Clearly if $H = 0.5$, $\alpha = 0$ (a flat slope) while $H > 0.5$ yields $\alpha > 0$ (a positive slope).

Fractional Gaussian noise is an example of stochastic processes that exhibit a *monofractal* scaling over *all* time scales. In particular, a fGn process with $H > 0.5$ is a long-range dependent (LRD) process; whereas fGn with $H = 0.5$ becomes the so-called “white noise,” as the spectrum is constant at all frequencies. In many practical applications, empirical processes such as network traffic processes may manifest different scaling behaviors at different time scales. To better model such processes, stochastic processes that exhibit multiple scalings have also been developed, one example of which is the *multi-scale* fractional Brownian motion proposed in [17], [18]. Such (second-order) multi-scale processes can be characterized by their spectral density function (SDF). We say that the process has a (*local*) *scaling exponent* h in the time scale range $[j_1, j_2]$ if

$$\Gamma(\nu) \sim |\nu|^{1-2h} \text{ for } \nu \in [\nu_1, \nu_2] \quad (6)$$

where $\nu_1 = 2^{j_1}\nu_0$, $\nu_2 = 2^{j_2}\nu_0$, and $[\nu_1, \nu_2]$ is the frequency band corresponding to the time scale range $[j_1, j_2]$. Using (3), it can be shown that

$$\log \mathcal{E}_j = -(2h - 1)j + c(j_1, j_2), \quad j \in [j_1, j_2], \quad (7)$$

where $c(j_1, j_2)$ is a constant depending on $[j_1, j_2]$. Hence the slope of $\log \mathcal{E}_j$ against $-j$ over the range $[j_1, j_2]$ in the energy plot provides an estimate of the (*local*) scaling exponent h . The same estimator using linear regression of $\log \mathcal{E}_j$ developed for Hurst parameter estimation in [15] can be used to estimate h over the range $[j_1, j_2]$. Following the same argument, such an estimator for h is (asymptotically) unbiased and efficient. As an example, consider a multi-scale fractional Brownian motion with two scaling exponents: $h^{(1)}$ for $j \leq j_1$

¹Here the negative sign indicates coarser time scales with increasing values on the x-axis.

and $h^{(2)}$ for $j > j_1$. (Note that $T_j = T_0/2^j$. Hence smaller j indicate larger time scale.) Hence $h^{(1)}$ governs the scaling behavior of the process over the large time scales ($j \leq j_1$), whereas $h^{(2)}$ governs the scaling behavior of the process over the small time scales ($j > j_1$). In particular, if $h^{(1)} > 0.5$, the process is an asymptotically self-similar process (thus LRD) with the Hurst parameter $H = h^{(1)}$. Furthermore, if $h^{(2)} = 0.5$, this process at the smaller time scales ($j > j_1$) behaves like a “white noise,” i.e., the (high-frequency) fluctuations at these time scales are uncorrelated.

C. Multifractal Analysis

Multifractal analysis compares the scaling of *different* wavelet moments to estimate the local regularity in processes [19], [6], [10]. Its main tool is the wavelet *partition function* $S_j(q)$, a generalization of the wavelet energy function (2), defined as

$$S_j(q) = E|W_{j,k}|^q. \quad (8)$$

For certain multifractal processes the partition function scales asymptotically ($j \rightarrow \infty$) as

$$\log_2 S_j(q) \sim q \cdot \text{constant} + j\alpha_q. \quad (9)$$

Thus the slope of $\log_2 S_j(q)$ against j provides an estimate of α_q . Processes with trivial *monofractal* scaling have α_q varying *linearly* with q while for *multifractals* the variation is *non-linear*.

Stationary Gaussian processes necessarily display trivial monofractal scaling. It is easily shown that for two Gaussian random variables X and Y ,

$$\log(E|X|^q/E|Y|^q) = (q/2) \log(E|X|^2/E|Y|^2), \quad (10)$$

thus implying that for Gaussian processes

$$\log S_j(q) - \log S_{j+1}(q) = (q/2)(\log S_j(2) - \log S_{j+1}(2)), \quad \forall j. \quad (11)$$

From (9), it follows that $\alpha_q = (q/2)\alpha_2$, i.e., α_q is a linear function of q . For example, the fGn process with Hurst parameter H has $\alpha_q = q(H + 1/2)$. Non-Gaussian processes, however, can possess *multifractal* scaling.

To test whether a process possesses multifractal scaling, two quantities $\zeta_q = \alpha_q - q/2$ and $h_q = \zeta_q/q$ are introduced [20]. Note that for fGn, $\zeta_q = qH$ and $h_q = H$. Thus a *straight* line ζ_q plot or a *constant* h_q plot characterize *monofractal* processes, while a *non-linear* ζ_q plot or a *non-constant* h_q plot characterize *multifractal* processes.




Fig. 1. Energy plots of OC12 traces.

III. MAJOR OBSERVATIONS

In this section we analyze the scaling behavior of the Internet backbone traffic using wavelets. For all traces we form a time series by counting the number of bytes every $10\mu\text{s}$. We then perform wavelet analysis on the time series using a Daubechies wavelet with 3 vanishing moments.²

A. Energy Plots

As representative examples to illustrate our observations, Figs. 1, 2 and 3 show the energy plots for a one-hour segment³ from the backbone packet traces listed in Table 1. We see that all plots show a dichotomy of scaling behaviors, the “knee” point (the transition region of scaling behaviors) occurring around 100ms - 1sec time scales, typically with a (slight) dip of energy in this region. Above 1sec or so, a strongly self-similar scaling (with Hurst parameter $H > 0.75$) is apparent in all plots. Below 100ms most links have a rather flat slope with local scaling exponent $h < 0.6$. In contrast, the OC12 corporate link (OC12-corp-dom) in Fig. 1(d) has a moderately correlated scaling ($h \approx 0.7$).

The Fourier transform (frequency) plots in Fig. 4 corroborate our energy plot correlation inferences. Recall

²The wavelet analysis code is based on the programs from [20].

³The small-time scalings we observed from the traces are fairly consistent over different time segments of the traces. Due to space limitation, we do not show these results, which can be found in [21].

from Section II-B that the wavelet energy at different scales correspond to the energy of the signal in different frequency bands. For example, the wavelet energy at small time scales 2-100ms correspond to the signal energy at the frequency range 10-500Hz. While the frequency plot of OC12-tier2-dom (Fig. 4(a)) is fairly flat in this region, signifying “white noise” or uncorrelated behavior, that of OC12-corp-dom (Fig. 4(b)) varies in a $1/f$ fashion, signifying stronger correlations.




Fig. 4. Fourier transform (frequency) plots.

B. Multifractal Analysis

To test whether the traces exhibit multifractal scaling at small time scales, we perform the multifractal analysis using the code from [20]. We plot the estimates of ζ_q and h_q at small time scales and find that all traces show trivial *monofractal* scaling. Some examples are shown in Fig. 5, which reveal that ζ_q is linear in q and h_q is a nearly constant, typical of monofractal processes. This finding is in contrast to earlier findings on traffic multifractality [5], [8], [10], [6]. This apparent contradiction is easily de-mystified through an analysis of the *Gaussianity* of the marginal distribution of traffic traces concerned.

| Traces | Kurtosis | Skewness |
|----------------|----------|----------|
| OC3-tier2-dom | 2.7 | 0.26 |
| OC12-tier2-dom | 3.04 | 0.2 |
| OC12-corp-dom | 2.86 | 0.24 |
| OC48-bb-1 | 3.06 | 0.20 |

TABLE II

ESTIMATED KURTOSIS AND SKEWNESS AT 4MS TIME SCALE.

Kurtosis and *skewness* are two metrics commonly used for Gaussianity test. A Gaussian distribution has kurtosis 3 and skewness 0, while heavier tailed distributions have higher values. Many of the earlier studied traces are from relatively low capacity (1-10Mbps) links, which transport less aggregated traffic with non-Gaussian marginals. For example, [22] reported kurtosis values greater than 5 at a time scale as large as 500ms for traces collected at Auckland university. This non-Gaussianity gives rise to non-trivial multifractal scaling. In contrast, the backbone traces have Gaussian-like marginals at small time scales (see Table II and Fig. 6 for marginals at 4ms time scale),




Fig. 2. Energy plots of OC3 traces.




Fig. 3. Energy plots of OC48 traces.

leading to trivial monofractal scaling (as explained in Section II-C). This is actually to be expected, as the backbone traces are of highly aggregated traffic (see Table I for average bandwidth of the traces). Since the backbone traces we study do not possess multifractality, which necessitates an analysis of higher-order statistics, in the remainder of this paper we will use energy plots only to study the small-time (second-order) scaling behaviors of the Internet backbone traffic.

IV. FLOW CLASSIFICATION AND TRAFFIC COMPOSITION

Clearly an aggregate traffic trace observed on an Internet backbone link comprises of packets from many individual traffic *flows* (defined in whatever appropriate terms). In this section, we classify flows in a traffic aggregate along two dimensions: *flow sizes* – the number of bytes contributed by a flow in a given time span, and *flow density* – the distribution of inter-packet arrival times of a flow in a given time span. Using these two flow classification schemes, we analyze the traffic composition of the packet traces, i.e., identifying flow components with different characteristics. In the next section we will use the proposed flow classification schemes to understand the scaling behaviors of various flow components, and investigate the impact of traffic composition on small-time scalings of aggregate traffic.

We use the standard *5-tuple*⁴ flow as our basic flow definition. Since we are primarily interested in *small-time* scaling behaviors in the range of 1-100 ms, we zoom into 1-minute segments of the packets traces to analyze the characteristics of their flows in detail. A 1-minute segment contains enough information to give good estimates of sub-second time scale (1ms-100ms) statistics, while being small enough to enable a feasible detailed flow-level analysis. Moreover, we have observed little deviation of the energy plot of a 1-minute segment of backbone traffic from the energy plot obtained from a 1-hour segment. In the remainder of this paper, we will focus primarily on two packet traces, *OC12-tier2-dom* and *OC12-corp-dom*, as they manifest strikingly different small-time scalings. The goal of our study is to *uncover the causes and factors that influence the small-time scalings of Internet backbone traffic*.

One way to classify flows is based on *flow size*, namely, the total number of bytes belonging to a flow. The distribution of flow size in bytes is largely application-layer dependent, and has proved sufficient to explain the correlation or long-range dependence at *large* (seconds to minutes) time scales [23], [1]. To understand whether flow size distribution also influence the traffic scaling

⁴A 5-tuple flow is a stream of packets that have the same source and destination IP addresses, source and destination port numbers, and protocol field in a given time span. It is the *finest-grain* flow one can identify using only IP/TCP header packet traces.




Fig. 5. Multifractal analysis.




Fig. 6. Histograms of bytes per 4ms.

behavior at sub-second time scales, we thus also classify flows based on their size. We introduce two parameters, $B_L \geq B_S$ for large vs. small flow classification. We term a flow as *large* if it transmits more than B_L bytes in a given 1-minute interval, and a flow as *small* if it transmits less than B_S bytes. Flows of sizes in between are termed as of *intermediate* size. If $B_L = B_S$, then flows are either large or small. In our study we typically choose $B_L = 10^6$ bytes and $B_S = 10^4$ bytes.

Fig. 7 shows (a portion of) the *complementary* cumulative distribution function (CCDF) of flow sizes (in bytes) in five different 1-minute segments that are one hour apart for both the OC12-tier2-dom and OC12-corp-dom traces. Note that the x-axis is given in logscale, and the distribution has a long tail. Clearly in both traces, a great majority of flows are *small*. For example, with $B_S = 10^4$ bytes, OC12-tier2-dom has more than 95% small flows, while OC12-corp-dom has close to 90% flows are small. With $B_L = 10^6$ bytes, only a tiny percentage of flows in both traces are large: OC12-tier2-dom has about 0.1% large flows, and OC12-corp-dom has about 0.5%. On the other hand, the tiny percentage of large flows contribute a significant proportion (around 20-40%) of total bytes in the aggregate traffic, as shown in Fig. 8. Figs.7 and 8 also reveal that the flow size distribution is fairly consistent over time, with little variations in five different 1-minute segments, each an hour apart.

Our second flow classification is based on the *distribution of packet inter-arrival times within a flow*. Before we introduce the formal definition, we show an example of two different flows in Fig. 9(a) and (b). The vertical bars in the top plots of the figures correspond to packet arrivals. We see that packet arrivals of the flow *a* are more sparsely spaced than those of flow *b*. The bottom

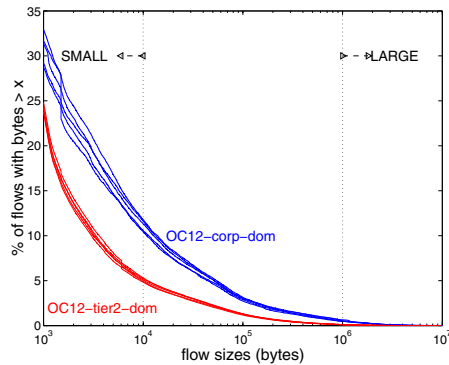


Fig. 7. CCDF of flow size for five 1-minute segments, each separated by 1 hour.

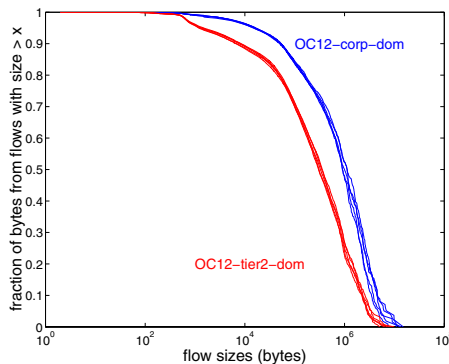


Fig. 8. Fraction of bytes belonging to flows larger than size x vs. x for five 1-min segments, each separated by 1 hour.

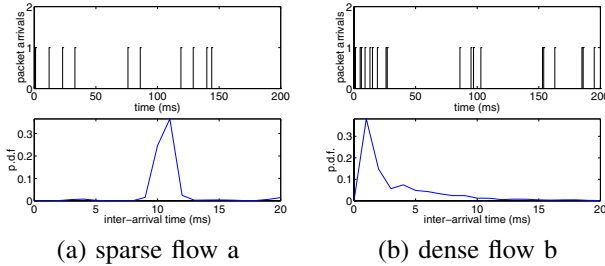


Fig. 9. Packet inter-arrival times and arrival process.

plots show the distribution of packet inter-arrival times of the two flows. Clearly, the *dominant* packet inter-arrival times of flow *a* are around 10 ms or so, while the *dominant* packet inter-arrival times of flow *b* are around 2 ms or so.

The above example motivates us to classify flows based on their *density*, namely, the *dominant packet inter-arrival times* of a flow. Given an appropriately chosen parameter T , we classify⁵ flows with 50% of packet inter-arrival times less than time T *dense*, and the remaining *sparse*. For $T = 3\text{ms}$, the flow *a* in Fig. 9 will be classified as sparse and flow *b* as dense.

To gain an intuitive understanding of the “density” distribution of flows in the OC12-tier2-dom and the OC12-corp-dom traces, we compute the histograms of packet inter-arrival times for individual flows in each trace in a given 1-minute segment, and then sum the corresponding bins of all the histograms to obtain an “aggregate” histogram. The aggregate histogram tells us the relative occurrence of dominant packet inter-arrival times among the flows in the trace. The aggregate histogram is shown in Fig. 10. It is clear that the dominant packet inter-arrival times of the flows in the OC12-corp-dom trace are much smaller than those of the flows in the OC12-tier2-dom. Setting⁶ $T = 2\text{ms}$, Table III shows that the traffic composition of sparse and dense flows in 5 different 1-minute segments, each one hour apart, for both OC12-tier2-dom and OC12-corp-dom traces. We see that the *dense vs. sparse* traffic compositions of the two traces remain fairly consistent over time, and the difference between the OC12-tier2-dom and OC12-corp-dom traces is obvious. In the OC12-corp-dom trace about 2.1-2.5% of the total flows are dense, compared to the less than 1% dense flows in the OC12-tier2-dom. Furthermore, the dense flows in the OC12-corp-dom trace contribute 15-20% of the total bytes, while the dense flows in the OC12-tier2-dom trace contribute fewer than 4% of the total bytes.

⁵Of course, we can define a spectrum of flows with different density, from sparse to dense. For simplicity, we choose a dichotomy of *sparse* and *dense* flows.

⁶As is evident in Fig. 10, the results reported here hold *qualitatively* as long as we classify dense/sparse flows in a reasonable fashion, e.g., setting $T = 4\text{ms}$ instead of 2ms . Due to space limitation, we do not show such results here.

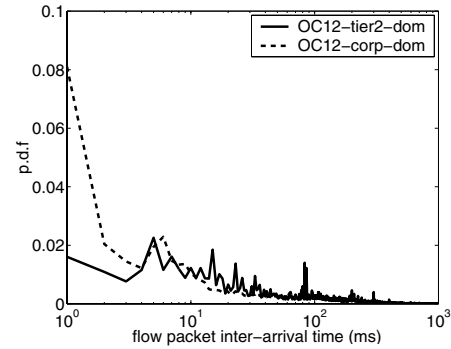


Fig. 10. P.d.f. of intra-flow packet inter-arrival times.

Before we leave this section, we would like to make a few comments. As mentioned earlier, 5-tuple flow size distribution is generally *application-dependent*. The majority of the flows use TCP, with web and file sharing (e.g., Kazaa) applications predominant⁷. On the other hand, the dominant packet inter-arrival times of flows are likely a result of the bottleneck link speed of the underlying networks, coupled with the feedback control mechanism of TCP. We note that many users behind the OC12-corp-dom link have high-speed Internet connections, whereas flows traversing the OC12-tier2-dom link come from users with more diverse Internet connections. This may explain why the OC12-corp-dom trace has more dense flows percentage-wise than the OC12-tier2-dom trace. We also note that both the file size distribution as well as the bottleneck speeds within the Internet backbone are unlikely to change substantially over time, at least not within a time span of several hours. Thus both taxonomies are in accordance with our goal of finding time invariant causes for small-time scalings.

V. TRAFFIC COMPOSITION AND SMALL-TIME SCALINGS

We investigate how the small-time scalings of aggregate traffic depend on their traffic compositions. More specifically, we first illustrate how small-time scalings are affected by the *large vs. small* and *dense vs. sparse* flow decomposition; we then examine how small-time scalings of aggregate can be altered by varying its traffic composition of dense and sparse flows. We focus primarily on the *OC12-corp-dom* and *OC12-tier2-dom* traces. Our results point to the (relatively) large proportion of dense flows in the *OC12-corp-dom* trace as the key cause for its strong correlation in small time scales.

A. Small-Time Scalings of Large vs. Small Flows

We begin by studying the dependence of small time scalings on the *flow size*. We use a 1-minute segment of the two OC-12 traces (the “Min 1” segment in Table III) as an example. The energy plots for the *large* ($B_L = 10^6$

⁷See [21] for a breakdown of various applications in the traces.

| Trace Name | Flow Type | Min 1 | | Min 2 | | Min 3 | | Min 4 | | Min 5 | |
|----------------|-----------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| | | %flows | %bytes | %flows | %bytes | %flows | %bytes | %flows | %bytes | %flows | %bytes |
| OC12-corp-dom | dense | 2.1 | 17.1 | 2.2 | 14.9 | 2.4 | 15.8 | 2.74 | 22.1 | 2.4 | 15.1 |
| | sparse | 97.8 | 82.8 | 97.7 | 85.04 | 97.6 | 84.1 | 97.2 | 77.8 | 97.6 | 84.8 |
| OC12-tier2-dom | dense | 0.7 | 2.0 | 0.7 | 3.9 | 0.9 | 3.3 | 0.8 | 3.0 | 0.8 | 3.6 |
| | sparse | 99.2 | 97.9 | 99.2 | 96.1 | 99.1 | 96.6 | 99.1 | 96.9 | 99.1 | 96.4 |

TABLE III

DENSE VS. SPARSE TRAFFIC COMPOSITIONS IN 5 1-MINUTE SEGMENTS OVER 5 HOURS ($T = 2\text{ms}$).




Fig. 11. Energy plots.

bytes) and *small* ($B_S = 10^4$ bytes) flows of the two traces are displayed in Fig.11. The aggregate of *OC12-corp-com* large flows has a small-time scaling exponent $h = 0.69$, similar to that of the total traffic aggregate in Fig. 1(d). In contrast, the aggregate of *OC12-corp-com* small flows has a small-time scaling exponent $h = 0.54$, considerably smaller than that of the total traffic aggregate. In comparison, both aggregates of *OC12-tier2-dom* large and small flows exhibit an uncorrelated small-time scaling with h fairly close to 0.5, similar to the total traffic aggregate (Fig.1(c)).

The above results show that the aggregate of small flows appear like “white noise” with uncorrelated scaling at small time scales. Unlike small flows, large flows in themselves do necessarily not induce correlated or uncorrelated small-time scalings, as is evident in Fig.11(a). Hence *flow size alone does not determine the small-time scaling of aggregate traffic*. This phenomenon is quite different from what we know about the scaling behavior of aggregate traffic at large (above seconds) time scales: it has been shown that the (asymptotically) *self-similar scaling* of aggregate traffic is caused by the *heavy-tailed* file size distribution.

B. Small-Time Scalings of Dense vs. Sparse Flows

We now study the dependence of small time scalings on the *flow density*. Using the same 1-minute segment from both traces as before, we decompose the total traffic aggregate into aggregates of dense and sparse flows by setting $T = 2\text{ms}$. The corresponding energy plots are shown in Fig. 12. Clearly, the aggregate of dense flows shows strong correlation at small time scales, with $h \geq 0.7$. In particular, the aggregate of dense flows from the *OC12-tier2-dom* has a small-time scaling exponent

$h \approx 0.7$, significantly larger than that of the total traffic aggregate. The aggregate of sparse flows has a much smaller scaling exponent ($h \leq 0.6$) at the small time scales. Note that although the aggregate of sparse flows of the *OC12-corp-dom* trace possesses some correlation ($h \approx 0.6$), it is still far smaller than that of the dense flows ($h \approx 0.76$). It is also smaller than that of the total traffic aggregate ($h \approx 0.65$, see Fig. 1(d)). These observations also hold for other packet traces. As an example, Fig. 13 shows the small time scalings of *dense vs. small dense* flows for two OC48 traces.

To further investigate the impact of flow density on small-time scalings, we separate the total traffic aggregate of two traces into *large* and *small* (with $B_L = B_S = 10^5$ bytes) components, and analyze the impact of *dense* and *sparse* flows (with $T = 2\text{ms}$) on each components. The resulting energy plots are shown in Figs. 14 and 15. We observe that the *sparse large* flows have a much smaller scaling exponent than the *dense large* flows at small time scales. Even the *dense small* flows manifest correlated scalings at small time scales. These results clearly point to dense flows as the correlation-causing factor at small time scales.

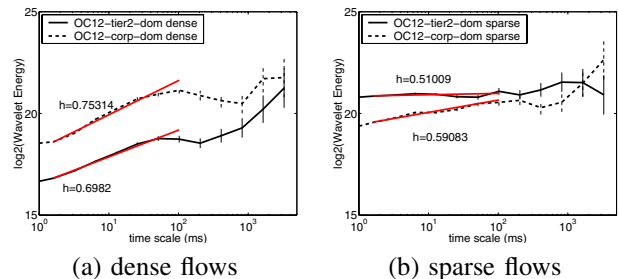


Fig. 12. Small-time scalings of dense vs. sparse flows.

C. Impact of Traffic Composition on Small-Time Scalings

So far we have demonstrated that the key factor in influencing the small-time scaling is the *flow density* – the *dense* flow component produces strong correlation at small time scales, while *sparse* flow component does not. Since both the *OC12-tier2-dom* and *OC12-corp-dom* traces contain dense flow component, why do they exhibit strikingly different small-time scalings? The answer to this question is obvious when we examine




Fig. 13. Small-time scalings of dense vs. sparse flows for OC48 traces.

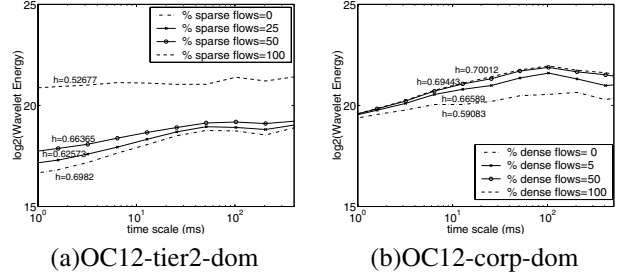


Fig. 16. Impact of *dense vs. sparse* traffic composition on small-time scalings.

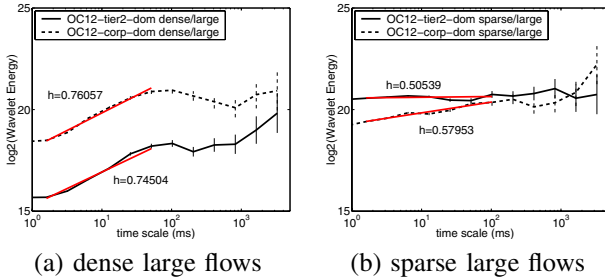


Fig. 14. Small-time scalings of dense vs. sparse large flows.

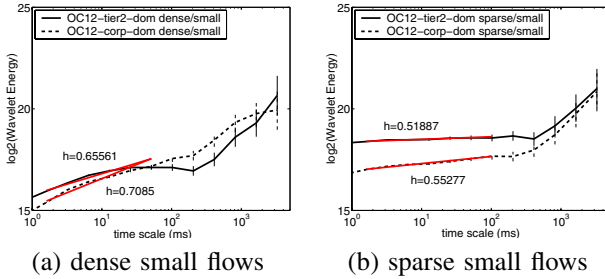


Fig. 15. Small-time scalings of dense vs. sparse small flows.

their traffic compositions in terms of dense and sparse flows. Fig. 10 in Section IV shows that the flow density distribution of the two traces are significantly different: the *OC12-corp-dom* trace has more flows with shorter inter-packet arrival time (e.g., 4 ms), namely, more flows with bursts of densely clustered packet arrivals. Using $T = 2ms$ to classify dense and sparse flows, Table III shows that in the *OC12-corp-dom* trace about 2.1-2.5% of the total flows are dense, which contribute 15-20% of the total bytes; whereas in the *OC12-tier2-dom* trace less than 1% of the total flows are dense, contributing fewer than 4% of the total bytes. Consequently, the dense flows in the *OC12-corp-dom* trace are more dominant, causing the aggregate traffic to exhibit correlated small-time scaling, while the dense flows in the *OC12-tier2-dom* have an almost negligible impact on the small time scaling of the aggregate traffic.

To further illustrate how correlation in small time scales is affected by traffic composition, namely, the relative proportions of dense and sparse flows, we perform *semi-*

experiments using the traffic traces. We mix different proportions of dense and sparse flows and demonstrate how small-time scalings changes with the traffic composition. We first start with the aggregate of all dense large flows in the *OC12-tier2-dom* trace, and observe the change in small-time scaling behaviors by progressively adding in sparse large flows: from 0% to 100% of all sparse flows in the trace. The results are shown in Fig. 16(a), where we see that the scaling exponent h in the time scale range of 1-10 ms decreases from around 0.69 to around 0.52, as the proportion of sparse flows increases. We now show how small-time scaling exponent can change in the other direction by starting with the aggregate of all sparse flows in the *OC12-corp-dom* trace, gradually adding more dense large flows into the traffic aggregate: from 0% to 100% of all *sparse* flows in the trace. The results are shown in Fig. 16(b), where we see that the scaling exponent h in the time scale range of 1-10 ms increases from 0.59 to 0.70, as the proportion of dense large flows increases. These two semi-experiments clearly demonstrate the impact of dense vs. sparse traffic composition on the small-time scaling behaviors.

D. Effect of RTT on Small-Time Scaling

We conclude this section by examining the effect of RTT on small-time scalings of the traces. Through simulations, in [24], [25] it shows that strong periodicities at the RTT time scale caused by TCP's feedback control can lead to a decrease a "dip") in wavelet energy at that scale, which might be attributed as the cause of the "flat" (uncorrelated-like) scaling we observe at the small time scales. To address this issue, we obtain RTT estimates for a large sample of TCP flows using the method in [26]. In both traces, the RTTs for both dense and sparse flows are within the range of around 10^2 ms to 1 sec or so. In the case of *OC12-tier2-dom*, the dense and sparse flows have a very similar RTT distribution. In contrast, the dense flows of the *OC12-corp-dom* have a significantly large number of RTTs at about 10ms, while the sparse flows have a wider RTT distribution. The results suggest that we *cannot* attribute the different small-time scalings of sparse and dense flows we observe in the traces to

the difference in the RTT characteristics of the flows. Direct inspection of the flow characteristics at the frequency domain (see Fig.18) also corroborate the RTT analysis. Note the striking $1/f$ -type decay for the dense flows in the frequency range 10-200Hz (5-100ms time scales) and the (relatively) flat spectrum in the same range for the sparse flows. They more explicitly account for difference in the small-time scaling behaviors (recall the relationship between wavelet energy and spectrum density function in Section II). Note that for the sparse flows of the OC12-corp-dom trace, the spectrum in frequency range 10-200Hz is not as flat as that of the sparse flows of the OC12-tier2-dom, which explain the slightly higher scaling exponent we observe (see Fig.12(b)). In summary, our analysis does not indicate a strong impact of RTTs on small-time scalings we observe in the traces. The difference in the small-time scaling behaviors can be better explained by the difference in the spectrum density functions of the dense and sparse flows in the traces. Intuitively, dense flows, by definition, inject bursts of densely clustered packets into the network, causing strong correlation in small time scales. We believe that flow density is more plausibly determined by the bottleneck link speed of the underlying network (coupled with the TCP feedback control mechanism), rather than the RTT effect of TCP feedback control mechanism alone.




Fig. 17. RTT estimates from 3-way TCP handshake packets.

VI. RELATED WORK

Recently multifractal small-time scalings of various network traffic statistics have been reported in a number of studies [5], [8], [10], most notably [6], where a random multifractal cascade is used for modeling small-time scaling behaviors. In [24], through simulation, the authors prescribe TCP (especially the resulting strong RTT periodicities) as a possible cause for the diverse small-time scaling behaviors of network traffic. By performing detailed flow-level analysis of packet traces from high-speed links, we attribute traffic composition in terms of flow density (a likely product of both bottleneck link speed and TCP feedback control mechanism) as a key factor in influencing small-time scalings. Closely related to our work is the connection-level analysis of low-speed packet traces (3 Mbps) conducted in [22], which proposes the *alpha vs. beta* flow dichotomy in

terms of flow marginals at 500ms time scale (similar to our small/large flow classification), and study the impact of alpha/beta flows on queuing performance. It also points to the bottleneck link speed as a potential factor in producing alpha flows. Our work differs from [22] in both the time scales (1-100 ms vs. 500 ms) of analysis and the fact that we study the impact of flow properties on correlations (small-time scalings) in addition to marginals. Note that on low-speed links, alpha or large flows would behave like dense flows, due to relatively low degree of aggregation. In another set of studies [27], [28], based on the analysis of inter-packet arrival times, the authors conclude that packet arrivals and sizes can go *locally* to independence with increased statistical multiplexing. Our study differs from theirs in two important aspects: (1) we focus on the second-order statistics of traffic volume fluctuations at small time scales, not packet arrivals or sizes; and (2) we find that high degree of aggregation *in itself* does not produce *uncorrelated* scalings, but instead distribution of flow density plays a major role.

VII. CONCLUSIONS

Using packet traces collected from OC3/12/48 links on a tier-1 ISP we studied the small-time scaling behaviors of Internet backbone traffic. We observed that for a majority of these traces, the (second-order) scaling parameters at small time scales (1ms - 100ms) are fairly close to 0.5. Hence for these traces the traffic fluctuations at these time scales are nearly uncorrelated. A small number of traces do exhibit some correlation (with scaling exponents within the range of 0.6 - 0.7). In addition, the traces manifest mostly *monofractal* (instead of multifractal) behaviors at small time scales. This can be attributed to the high degree of traffic aggregation on the Internet backbone links, which results in more Gaussian-like marginal distributions at small time scales. To understand the factors that influence the small-time scalings, we analyzed the traffic composition of the traces along two dimensions – flow size and flow density. We found that flow size alone cannot determine the small-time scaling of aggregate traffic. Our results point to dense flows as the correlation-causing factor at small time scales. They reveal that traffic composition in terms of proportions of *dense vs. sparse* flows plays a major role in influencing the small-time scalings of aggregate traffic. We are currently investigating mathematical models to provide explanation for and gain theoretical insights into our findings. Our observations and results have significant implications in networking modeling, service provisioning and traffic engineering. For example, the discovery of uncorrelated small-time scalings on many Internet backbone links can lead to simpler network models for analyzing network performance at small time scales [11], [13]. Our finding that small-time scalings are determined by dense




Fig. 18. Fourier transform (frequency) plots of dense and sparse flows.

vs. sparse traffic composition also raises many intriguing issues regarding the impact of Internet evolution on traffic behaviors. On the one hand, with increasing speed of Internet backbone links, dense flows appear “sparser” as we zoom into even smaller time scales; therefore we would expect traffic on these links are more likely to appear “independent” at those small time scales that matter to queuing [11]. On the other hand, as broadband access becomes more widely deployed, large files and objects will be transmitted faster into the Internet, with more correlated bursts. In addition, the changing nature of applications and increased use of protocols other than TCP can also effect a change in the small-time scalings of the Internet backbone traffic. These are important questions awaiting to be explored.

REFERENCES

- [1] K. Park and W. Willinger, Eds., *Self-Similar Network Traffic and Performance Evaluation*. Wiley Interscience, 2000.
- [2] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, “On the self-similar nature of Ethernet traffic (extended version),” *IEEE/ACM Trans. Networking*, pp. 1–15, 1994.
- [3] V. Paxson and S. Floyd, “Wide-area traffic: The failure of Poisson modeling,” *IEEE/ACM Trans. on Networking*, vol. 3, pp. 226–244, 1995.
- [4] M. Crovella and A. Bestavros, “Self-similarity in World Wide Web traffic. Evidence and possible causes,” *IEEE/ACM Trans. on Networking*, vol. 5, pp. 835–846, December 1997.
- [5] J. Lévy Véhel and R. Riedi, “Fractional Brownian motion and data traffic modeling: The other end of the spectrum,” *Fractals in Engineering*, pp. 185–202, Springer 1997.
- [6] A. Feldmann, A. C. Gilbert, and W. Willinger, “Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic,” *Proc. ACM/SIGCOMM 98*, vol. 28, pp. 42–55, 1998.
- [7] A. Feldmann, A. C. Gilbert, W. Willinger, and T. Kurtz, “The changing nature of network traffic: Scaling phenomena,” *Computer Communication Review*, vol. 28, pp. 5–19, 1998.
- [8] P. Mannersalo and I. Norros, “Multifractal analysis of real ATM traffic: a first look,” *COST257TD*, 1997.
- [9] R. J. Gibbens and Y. C. Teh, “Critical time and space scales for statistical multiplexing,” *Proc. of the 16th International Teletraffic Congress – ITC 16*, pp. 87–96, 1999.
- [10] R. Riedi, M. S. Crouse, V. Ribeiro, and R. G. Baraniuk, “A multifractal wavelet model with application to TCP network traffic,” *IEEE Trans. Info. Theory, Special issue on multiscale statistical signal analysis and its applications*, vol. 45, pp. 992–1018, April 1999.
- [11] A. L. Neidhardt and J. L. Wang, “The concept of relevant time scales and its application to queuing analysis of self-similar traffic,” in *Proc. SIGMETRICS '98/PERFORMANCE '98*, 1998, pp. 222–232.
- [12] M. Grossglauser and J.-C. Bolot, “On the relevance of long-range dependence in network traffic,” *Computer-Communication-Review*, vol. 26, no. 4, pp. 15–24, October 1996.
- [13] C. Fraleigh, F. Tobagi, and C. Diot, “Provisioning ip backbone networks to support delay-based service level agreements,” in *Proceedings of IEEE INFOCOM 2003*, San Francisco, CA, March 30 - April 2 2003.
- [14] C. Fraleigh, S. Moon, C. Diot, B. Lyles, and F. Tobagi, “Packet-level traffic measurements from a tier-1 IP backbone,” *Sprint ATL Technical Report TR01-ATL-110101*, November 2001.
- [15] P. Abry and D. Veitch, “Wavelet analysis of long-range-dependent traffic,” *IEEE Trans. on Information Theory*, vol. 44, no. 1, pp. 2–15, January 1998.
- [16] L. Kaplan and C.-C. Kuo, “Fractal estimation from noisy data via discrete fractional Gaussian noise (DFGN) and the Haar basis,” *IEEE Trans. Signal Proc.*, vol. 41, no. 12, pp. 3554–3562, Dec. 1993.
- [17] J.-M. Bardet and P. Bertrand, “Detecting abrupt change on the Hurst parameter of a multi-scale fractional Brownian motion with applications,” in *Internation Meeting on Self-Similarity and Applications*, Clermonts-Ferrand, France, May 27-30 2002.
- [18] A. Benassi and S. Deguy, “Multi-scale fractional Brownian motion: definition and identification,” Preprint, 1999.
- [19] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, and D. Veitch, “The multiscale nature of network traffic,” *IEEE Signal Processing Magazine*, 2002.
- [20] D. Veitch, “D. Veitch’s Matlab code for the estimation of scaling exponents,” <http://www.emulab.ee.mu.oz.au/~darryl/>.
- [21] V. Riberio, Z.-L. Zhang, S. Moon, and C. Diot, “Small-time scaling behaviors of internet backbone traffic: An empirical study,” *Sprint ATL Technical Report*, July 2002.
- [22] S. Sarvotham, R. Riedi, and R. Baraniuk, “Connection-level analysis and modeling of network traffic,” *Proc. IEEE/ACM Internet Measurement Workshop*, 2001.
- [23] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, “Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level,” *Computer-Communication-Review*, vol. 25, no. 4, pp. 100–113, October 1995.
- [24] A. Feldmann, A. C. G. P. Huang, and W. Willinger, “Dynamics of IP traffic: A study of the role of variability and the impact of control,” *Proc. ACM/SIGCOMM*, 1999.
- [25] P. Huang, A. Feldmann, and W. Willinger, “A non-intrusive, wavelet-based approach to detecting network performance problems,” *Proc. IEEE/ACM Internet Measurement Workshop*, 2001.
- [26] H. Jiang and C. Dovrolis, “Passive estimation of tcp round-trip times,” *To appear in the ACM Computer Communications Review*, 2002.
- [27] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun, “On the nonstationarity of Internet traffic,” *Proc. ACM SIGMETRICS '01*, pp. 102–112, 2001.
- [28] —, “The effect of statistical multiplexing on internet packet traffic: Theory and empirical study,” *Bell Labs Technical Report*, 2001.