



Small-time scaling behavior of Internet backbone traffic [☆]

Vinay J. Ribeiro ^{a,*}, Zhi-Li Zhang ^b, Sue Moon ^c, Christophe Diot ^d

^a Department of Electrical and Computer Engineering, Rice University, 6100 Main Street, Houston, TX 77005, USA

^b Department of Computer Science, University of Minnesota, 200 Union Street S.E., Minneapolis, MN 55455-0159, USA

^c Department of Computer Science, KAIST, Guseong-Dong, Yuseong-Gu, Daejeon 305-701, South Korea

^d Intel Research, 15 JJ Thomson Av., Cambridge CB3 0FD, UK

Available online 31 December 2004

Abstract

We perform an extensive wavelet analysis of Internet backbone traffic traces to observe and understand the causes of small-time scaling phenomena present in them. We observe that for a majority of the traces, the second-order scaling exponents at small time scales (1–100 ms) are fairly close to 0.5, indicating that traffic fluctuations at these time scales are nearly uncorrelated. Some traces, however, do exhibit moderately large scaling exponents (≈ 0.7) at small time scales. In addition, the traces manifest mostly monofractal behaviors at small time scales. To identify the network causes of the observed scaling behavior, we analyze the flow composition of the traffic along two dimensions—flow byte contribution and flow density. Our study points to the dense flows (i.e., flows with 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. Since queuing inside routers is influenced by traffic fluctuations at small time-scales, our observations and results have important implications for networking modeling, service provisioning and traffic engineering.

© 2004 Elsevier B.V. All rights reserved.

Keywords: Traffic; Self-similarity; Scaling; Long range dependence; Internet; Backbone

1. Introduction

Scaling behaviors of traffic influence network performance and thus have been the focus of much research (see [1] and references therein). Self-similar scaling over large time scales (e.g., 1 s and above) has been observed in a variety of network traffic [2–4]. More recently, more complex, perhaps multifractal-like, scaling behaviors below

[☆] The majority of this work was conducted while the authors were at Sprint Advanced Technologies Laboratories.

* Corresponding author.

E-mail addresses: vinay@rice.edu (V.J. Ribeiro), zhzhang@cs.umn.edu (Z.-L. Zhang), sbmoon@cs.kaist.ac.kr (S. Moon), christophe.diot@intel.com (C. Diot).

the 1 s time scale have also been reported [5–10]. Since traffic fluctuations at sub-second small time scales influence queuing inside routers, understanding of *small-time* scaling behaviors of network traffic is critical to many network design and engineering problems such as router buffer dimensioning and delay-sensitive service provisioning in the Internet backbone [11–13]. In addition to the scaling behavior of traffic, other factors such as TCP's feedback mechanism impact queuing delays at highly congested router queues [14]. However, such routers are more likely to be located in *access* networks than in the Internet backbone.

In this paper we study the small-time (1–100 ms) scaling behaviors of Internet backbone traffic. Our analysis is based on day-long packet traces collected from OC3, OC12 and OC48 links of the Sprint network, a major tier-1 Internet Service Provider (ISP). The high-precision time stamping and the high link capacity allow us to zoom into finer time scales and perform reliable data analysis at these time scales. In addition, the large set of long packet traces enable us to compare scaling behaviors of traffic carried across links of various types such as links connecting to other ISPs and big corporate customers. In contrast, majority of the previous traffic data traces used in the study of traffic scaling behaviors have less than 100 Mbps 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 ISP 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 s or so, the traffic has a clear-cut single self-similar scaling, while below it, the scaling behaviors appear to be more varied [7]. The transition occurs between 100 ms and 1 s, 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 scaling behaviors on these Internet backbone links. We find that over a range of small time scales (typi-

cally 1–100 ms), 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 indicates that traffic fluctuation on these backbone links are almost independent at these small time scales! A small number of traces do exhibit some correlation with scaling exponents within the range of 0.6–0.75. 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 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 byte contribution (defined in bytes of a flow in a given time interval) and flow density (defined in terms of dominant packet inter-arrival times within a flow). We show that aggregation of flows with small byte contributions generally exhibits mostly uncorrelated small time scaling, while aggregation of flows with large byte contributions in itself does not induce correlated or uncorrelated small-time scaling. Hence flow byte contribution alone does not determine the small-time scaling of aggregate traffic. However, by analyzing traffic composition along the dimension of flow density, we uncover flows with densely clustered packets (dense flows) as the primary correlation-causing factor at 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 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.

This work extends an earlier paper presented at INFOCOM 2003 [15]. We enhance our earlier paper with additional trace analysis and more detailed theory sections.

1.1. Paper organization

The remainder of this paper is structured as follows. In Section 2 we describe the packet traces

used in the study. We then rigorously define various scaling concepts in Section 3 and present standard wavelet estimation schemes in Section 4 setting the stage for the wavelet analysis of backbone traffic in Section 5. Combining the power of wavelets with different flow classification schemes, we identify the causes of the observed scaling in Section 6. After a discussion on related work in Section 7, we conclude in Section 8.

2. Internet backbone traces

We use traces collected on the Sprint backbone network. The Sprint network covers a wide geographical area and has a variety of link types. Links within a Point-of-Presence (PoP) have OC3 (155 Mbps) to OC192 (10 Gbps) speeds, and inter-PoP links have OC48 (2.5 Gbps) or OC192 speeds. The customers are tier-2 or ISPs, corporations, and international ISPs. Roughly, the links can be classified into five types: “peering” links with other tier-1 ISPs; domestic “peering” links with tier-2 or lower ISPs; international “peering” links with non-US-based ISPs; and “corporate” links to corporate customers, such as Fortune 500 companies, web hosting sites, and web server farms; and inter-PoP links between backbone routers. By *peering* links, we mean the links carrying traffic between the ISP networks in question; they do *not* reflect the contractual agreements nor BGP policies between the ISPs. Note that classification of link types is somewhat subjective and ambiguous, as it is impossible to precisely pin down the role of each network in these terms.

The packet traces were collected at three different PoPs in the US: two on the East Coast, and the other on 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 of 5 μ s) which indicates when the packet was observed. For further details about the measurement system, see [16].

Table 1 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.

3. Scaling phenomena

This section briefly describes several scaling phenomena and related terminology of relevance to this paper. All the phenomena involve analyzing data at multiple time scales and hence the term “scaling”. For a more detailed discussion of this topic see [1].

For clarity of the discussion, we introduce the signal $Z(t)$, $t \in \mathbb{R}$, its increment process Y defined as

$$Y(k) := Z((k+1)\Delta) - Z(K\Delta), \quad k \in \mathbb{Z}, \quad (1)$$

where Δ is some constant, and its (distributional) derivative $X(t)$ defined as

$$Z(t) = \int_{\tau=0}^t dX(\tau). \quad (2)$$

In the networking context, $Z(t)$ represents the total bytes arriving in the time interval $[0, t]$, $Y(k)$ the bytes-per-time process obtained by binning the traffic in time intervals of size Δ , and $X(t)$ the instantaneous traffic rate at time t .

3.1. Self-similarity

The process $Z(t)$, $t \in \mathbb{R}$, is *self-similar* with *Hurst* parameter $H > 0$ if

$$Z(at) \stackrel{\text{fd}}{=} a^H Z(t), \quad t \in \mathbb{R}, \quad \forall a > 0, \quad (3)$$

where $\stackrel{\text{fd}}{=}$ denotes equality in finite dimensional distributions. Essentially, when scaled in time, self-similar processes maintain their statistical properties except for a constant scaling factor. A prime example of self-similar processes is fractional Brownian motion (fBm).

3.2. LRD and If processes

A stationary process $Y(k)$, $k \in \mathbb{Z}$, with autocovariance function $r_Y(k)$ and spectral density function (SDF) $\Gamma_Y(f)$ is LRD [17], if

$$\sum_k r_Y(k) = \infty \quad \text{or} \quad \Gamma_Y(0) = \infty. \quad (4)$$

The non-negligible correlations at large time lags (or equivalently strong low frequency components) make statistical estimation of LRD processes non-trivial [18].

Table 1
Traffic data description

Trace	Link type/Direction	Start time	Duration (h)	Speed	Average bandwidth (Mbps)
OC3-tier1-dom	To domestic tier-1 peer	Aug. 8, 2000	15	OC3	42
OC3-tier2-dom	From domestic tier-2 peer	Aug. 8, 2000	23	OC3	44
OC3-corp-dom	From corporate	Aug. 8, 2000	19	OC3	28
OC12-tier2-dom	To domestic tier-2 peer	Sept. 5, 2001	7	OC12	187
OC12-corp-dom	From corporate	Sept. 5, 2001	20	OC12	122
OC48-bb-1	Inter-PoP link	April 19, 2002	1	OC48	746
OC48-bb-2	Inter-PoP link	April 19, 2002	1	OC48	844
OC48-bb-3	Inter-PoP link	April 19, 2002	1	OC48	1356

Classic examples of LRD processes are $1/f$ -noise processes whose SDFs behave like

$$\Gamma(f) \sim c|f|^{1-2h}, \quad f \rightarrow 0, \quad (5)$$

where c is a non-zero constant and the scaling exponent $h \in (0.5, 1)$. Here $f(x) \sim l(x)$ as $x \rightarrow 0$ denotes $\lim_{x \rightarrow 0} f(x)/l(x) = 1$. A scaling exponent value of $h = 0.5$ corresponds to white noise or the absence of correlation, while h close to 1 indicates a strongly correlated process.

Several real world time series possess multiple frequency scaling regions, that is,

$$\Gamma(f) \approx c_i |f|^{1-2h_i}, \quad f_1^{(i)} < f < f_2^{(i)}, \quad (6)$$

$$i = 1, \dots, K,$$

where c_i and h_i are constants as in (5). These time series are better modeled by generalizations of fBm such as the recently proposed multiscale fractional Brownian motion [19,20]. We will later use wavelet techniques to estimate the different scaling exponents h_i observed in network traffic traces.

3.3. Multifractality

In contrast to LRD which is a large time-scale (or equivalently, low frequency) second-order property, multifractality is a fine time-scale property characterized by higher-order moments of data [21].

Informally, the process $Z(t)$ has Hölder (or singularity) exponent $\beta(t)$ at time t if

$$|Z(t + \epsilon) - Z(t)| \stackrel{\epsilon \rightarrow 0}{\approx} \epsilon^{\beta(t)}. \quad (7)$$

A monofractal process has the same $\beta(t)$ at all time instants t , while a multifractal process has $\beta(t)$

varying with t . The fBm is an example of a monofractal process with $\beta(t) = H$, for all t . We will outline a technique to determine if $Z(t)$ is multifractal in the next section that uses the wavelet transform of the derivative $X(t)$ of $Z(t)$.

While LRD is defined for asymptotically large time scales and multifractality likewise for infinitesimally small time-scales, in practice we are constrained to analyze data only within a finite range of time scales. We will thus apply the above concepts to a *limited* range of time scales while analyzing the backbone traces.

4. Wavelets and scaling

This section describes wavelets and their use in estimating different scaling parameters. Wavelets provide a multiscale analysis framework that is ideal for identifying scaling phenomena [22,23]. Strongly correlated time signals become decorrelated in the wavelet domain allowing powerful wavelet-based scaling exponent estimators superior to time-domain ones. In addition wavelets eliminate polynomial non-stationary trends in data which can corrupt estimation techniques that rely on stationarity.

4.1. Wavelet transform

The discrete wavelet transform represents the 1-D function $X(t)$ in terms of shifted and dilated versions of a prototype bandpass wavelet function $\psi(t)$ and shifted versions of a low pass scaling function $\phi(t)$. For special choices of $\phi(t)$ and $\psi(t)$ we have the representation

$$X(t) = \sum_k 2^{J/2} a_X(J, k) \phi(2^J t - k) + \sum_{j=J}^{\infty} 2^{j/2} w_X(j, k) \psi(2^j t - k), \quad (8)$$

where the approximation coefficients $a_X(J, k)$ are defined as

$$a_X(J, k) := \int 2^{J/2} \phi(2^J t - k) dX(t), \quad (9)$$

and the wavelet coefficients $w_X(j, k)$ as

$$w_X(j, k) := \int 2^{j/2} \psi(2^j t - k) dX(t) \quad (10)$$

for some $J \in \mathbb{Z}$.

4.2. Wavelet logscale diagrams

Wavelet logscale diagrams provide simple estimates of different scaling parameters. They are plots of the logarithm of estimated wavelet coefficient moments against time scale j . Defining

$$S_j(q) := \frac{1}{n_j} \sum_{k=0}^{n_j-1} |w_X(j, k)|^q, \quad q \in \mathbb{R}, \quad (11)$$

where n_j is the number of wavelet coefficients at scale j , the q th-order logscale diagrams are plots of $\log S_j(q)$ vs. scale j . The second-order logscale diagrams are also sometimes called *wavelet energy* plots.

4.3. Second-order scaling exponent estimation

The second moment of wavelet coefficients is intimately related to the SDF of a wide-sense stationary process $X(t)$. If X has SDF

$$\Gamma_X(f) = c|f|^{1-2h}, \quad f_1 \leq |f| \leq f_2; \quad f_1, f_2 > 0, \quad (12)$$

and the Fourier transform of wavelet $\psi(t)$ has compact support in $f_a \leq |f| \leq f_b$, then for all j such that

$$f_1 \leq 2^j f_a \leq 2^j f_b \leq f_2, \quad (13)$$

we have

$$\mathbb{E}|w_X(j, k)|^2 = \kappa 2^{j(-1-2h)},$$

where κ is a constant. Thus the slope of $\log \mathbb{E}|w_X(j, k)|^2$ vs. j in the region where (13) holds gives us the scaling exponent h . In case X has multiple frequency scaling regions as in (6), the same technique gives us all scaling parameters h_i . Based on the second-order logscale diagram, Abry et al. presented an estimator of the scaling exponent h that is unbiased and efficient for Gaussian data assuming perfect decorrelation of $w_X(j, k)$ within scale j [23].

4.4. Multifractal analysis

Consider the process $Z(t)$ and its derivative $X(t)$ as defined in (2). Define

$$\zeta_q := -\lim_{j \rightarrow \infty} \frac{\log_2 \mathbb{E}|w_X(j, k)|^q}{j} - q/2, \quad q \in \mathbb{R}. \quad (14)$$

If ζ_q is a non-linear function of q then $Z(t)$ corresponds to a multifractal process. However, if ζ_q is linear then $Z(t)$ corresponds to a monofractal process [22].

Stationary Gaussian processes necessarily display trivial monofractal scaling. It is easily shown that for two zero-mean Gaussian random variables A and B ,

$$\log(\mathbb{E}|A|^q / \mathbb{E}|B|^q) = (q/2)(\log \mathbb{E}|A|^2 / \mathbb{E}|B|^2), \quad (15)$$

thus implying that for stationary Gaussian processes

$$\begin{aligned} \log \mathbb{E}|w_X(j, k)|^q - \log \mathbb{E}|w_X(0, k)|^q &= (q/2)(\log \mathbb{E}|w_X(j, k)|^2 \\ &\quad - \log \mathbb{E}|w_X(0, k)|^2), \quad \forall j. \end{aligned} \quad (16)$$

From (14), it follows that $\zeta_q = (q/2)\zeta_2$, i.e., ζ_q is a linear function of q . Non-Gaussian stationary processes, however, can possess multifractal scaling.

In practice we estimate ζ_q from the slope of the q th-order logscale diagram in a limited range of time scales j . For convenience we also plot the quantity $h_q = \zeta_q/q$. Note that for the monofractal fBm, $h_q = H$ for all q .

5. Scaling analysis of backbone traffic

In this section we analyze the scaling behavior of the Internet backbone traffic traces using

wavelets. We first form time series by counting the number of bytes arriving at a link every $10 \mu\text{s}$. We perform a wavelet analysis on this time series using a Daubechies wavelet with three vanishing moments. The wavelet analysis code is based on the programs from [24]. We follow the common practice of equating the finest time scale approximation coefficients with the time series values.

We proceed with the analysis assuming stationarity. While backbone traffic has non-stationarities caused by diurnal trends, they are well-modeled as stationary processes over periods of about an hour or less, periods over which the number of users accessing the network and the mean traffic rates are fairly constant. Moreover, our assumption is further justified by the fact that

the wavelet analysis is robust to certain polynomial non-stationarities.

5.1. Wavelet energy plots

As representative examples to illustrate our observations, Figs. 1–3 show the energy plots for 1-h segments from the backbone packet traces listed in Table 1. We use time-scale ranges where the energy plots are approximately linear to compute the scaling exponents. For the scaling exponent at small time scales we employ the 1.28–163.84 ms range for the OC12 and OC48 traces and the 10.24–163.84 ms range for the OC3 traces. Below 10 ms the OC3 traces show a different scaling to that between 10 and 100 ms.

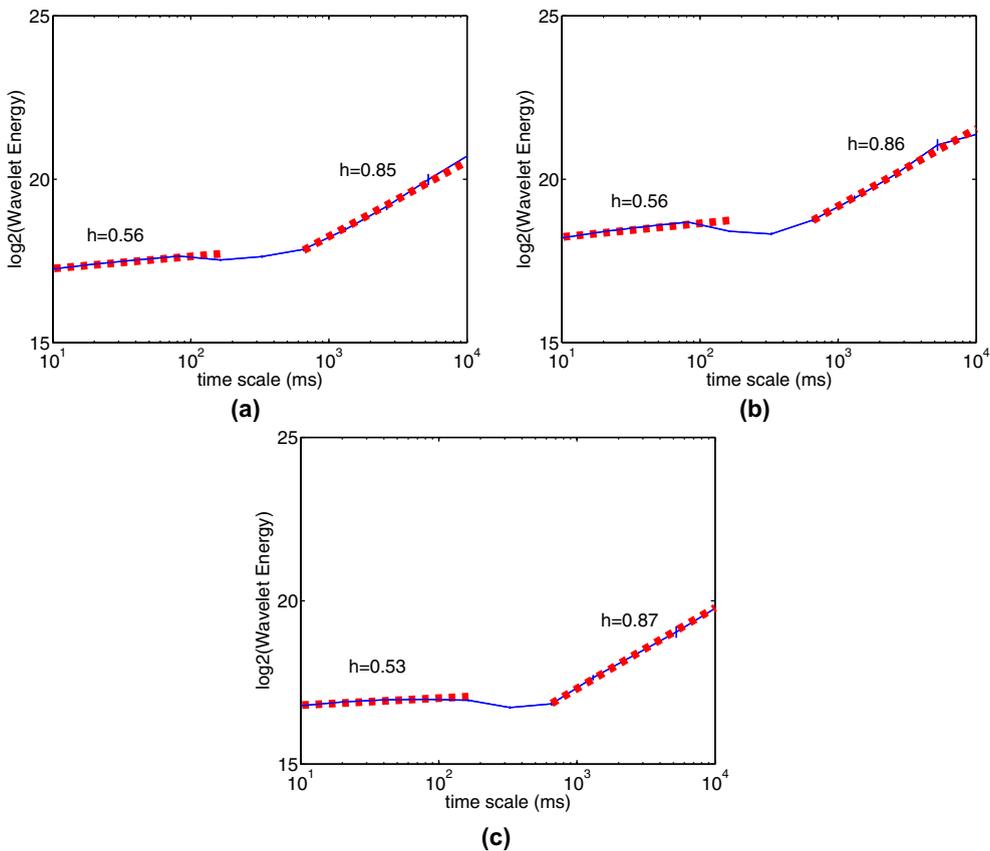


Fig. 1. Wavelet energy plots of OC3 traffic. Observe uncorrelated scaling between 10 and 100 ms: (a) OC3-tier1-dom, (b) OC3-tier2-dom, (c) OC3-corp-dom.

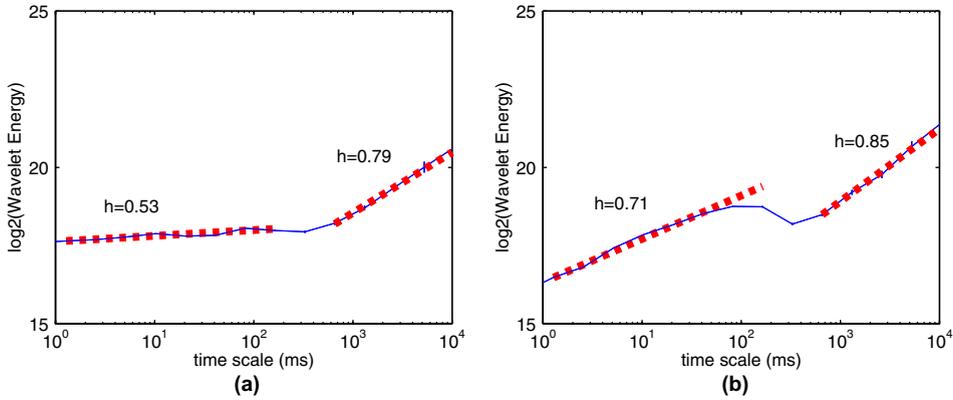


Fig. 2. Wavelet energy plots of OC12 traffic. Observe nearly uncorrelated scaling between 1 and 100 ms in (a) while in (b) observe strongly correlated scaling. (a) OC12-tier2-dom, (b) OC12-corp-dom.

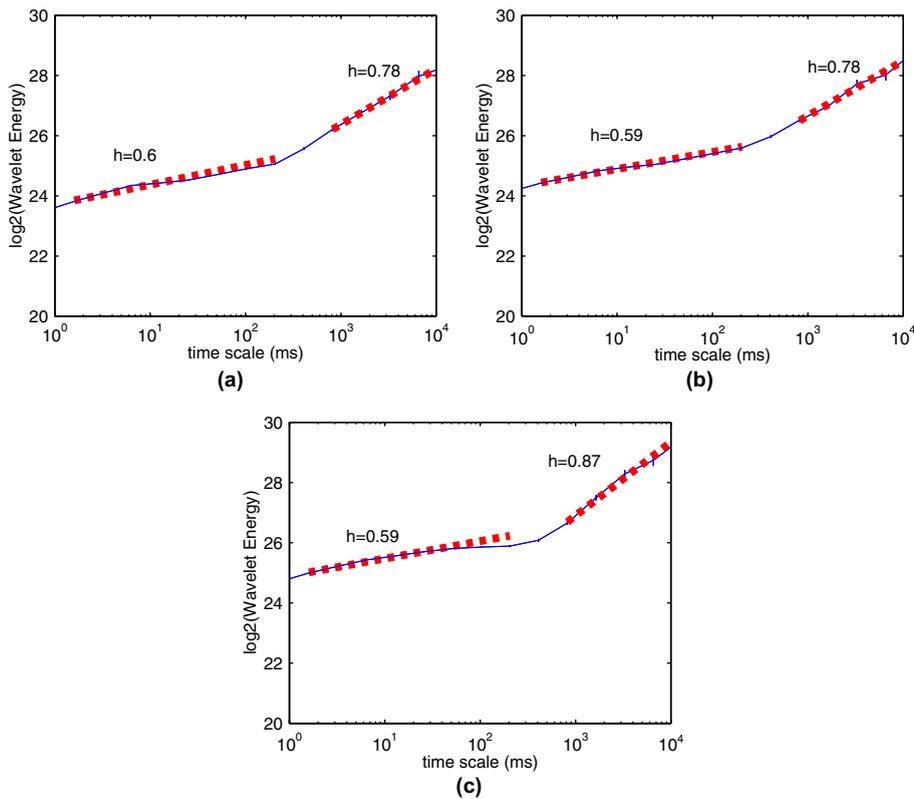


Fig. 3. Wavelet energy plots of OC48 traffic. All traffic traces show nearly uncorrelated scaling at the 1–100 ms time-scales: (a) OC48-bb-1, (b) OC48-bb-2, (c) OC48-bb-3.

Due to the low aggregation of traffic at the 1–10 ms time scales the packet size variation and

other factors could potentially influence the scaling behavior. A detailed analysis to determine the

causes of this difference in scaling is beyond the scope of this paper. For the scaling exponent at large time scales we employ the 0.655–10.48 s range. The dashed lines in the various plots correspond to the estimated scaling exponents.

We see that all plots show a dichotomy of scaling behaviors, the “knee” point (the transition region of scaling behaviors) occurring around the 100 ms–1 s time scales, typically with a (slight) dip of energy in this region. Above 1 s or so, a single linear strongly self-similar scaling is apparent in all plots. Below 100 ms 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. 2(b) has a moderately correlated scaling ($h \approx 0.7$). We also observe from Table 2 that the scaling exponents are fairly consistent over different time segments. The short 1-h duration of the OC48 traces prevents their inclusion in Table 2.

The Fourier transform plots of Fig. 4 corroborate our energy plot correlation inferences. Recall from Section 4.3 that the wavelet energy at different scales correspond to the energy in different frequency bands. For example, the wavelet energy at small time scales 1–100 ms correspond to the energy at the frequency range 10–1000 Hz. 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.

5.2. Multifractal analysis

To test whether the traces exhibit multifractal scaling at small time scales, we perform the multifractal analysis using the code from [24].

We obtain estimates of ζ_q and h_q using time scale ranges where the logscale diagrams are

Table 2
Estimated second-order scaling exponents for five different 1-h segments

Trace	Scaling exponent (small time scales)					Scaling exponent (large time scales)				
	Hour 1	Hour 2	Hour 3	Hour 4	Hour 5	Hour 1	Hour 2	Hour 3	Hour 4	Hour 5
OC3-tier1-dom	0.55	0.56	0.58	0.55	0.55	0.84	0.94	0.91	0.93	0.92
OC3-tier2-dom	0.56	0.57	0.53	0.52	0.55	0.85	0.91	0.89	0.85	0.89
OC3-corp-dom	0.53	0.53	0.53	0.53	0.53	0.87	0.86	0.88	0.84	0.87
OC12-tier2-dom	0.52	0.51	0.53	0.52	0.53	0.79	0.81	0.80	0.81	0.78
OC12-corp-dom	0.70	0.70	0.70	0.72	0.77	0.85	0.79	0.82	0.75	0.67

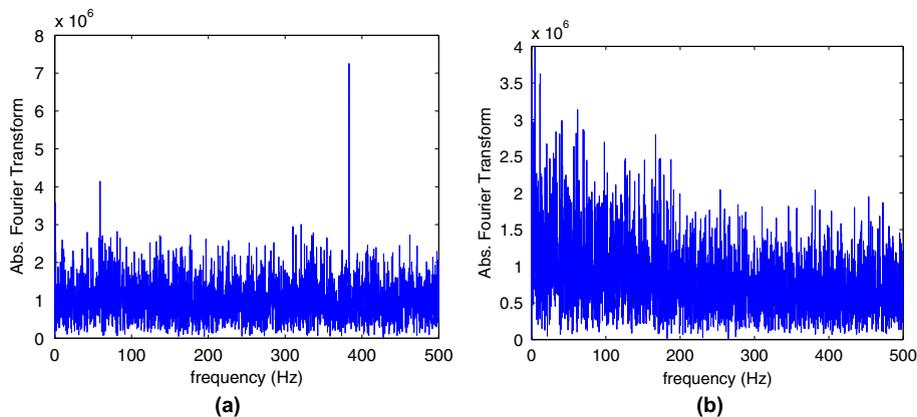


Fig. 4. Fourier transform plots. In (a) the spectrum is white, while in (b) it is like colored $1/f$ noise: (a) OC12-tier2-dom, (b) OC2-corp-dom.

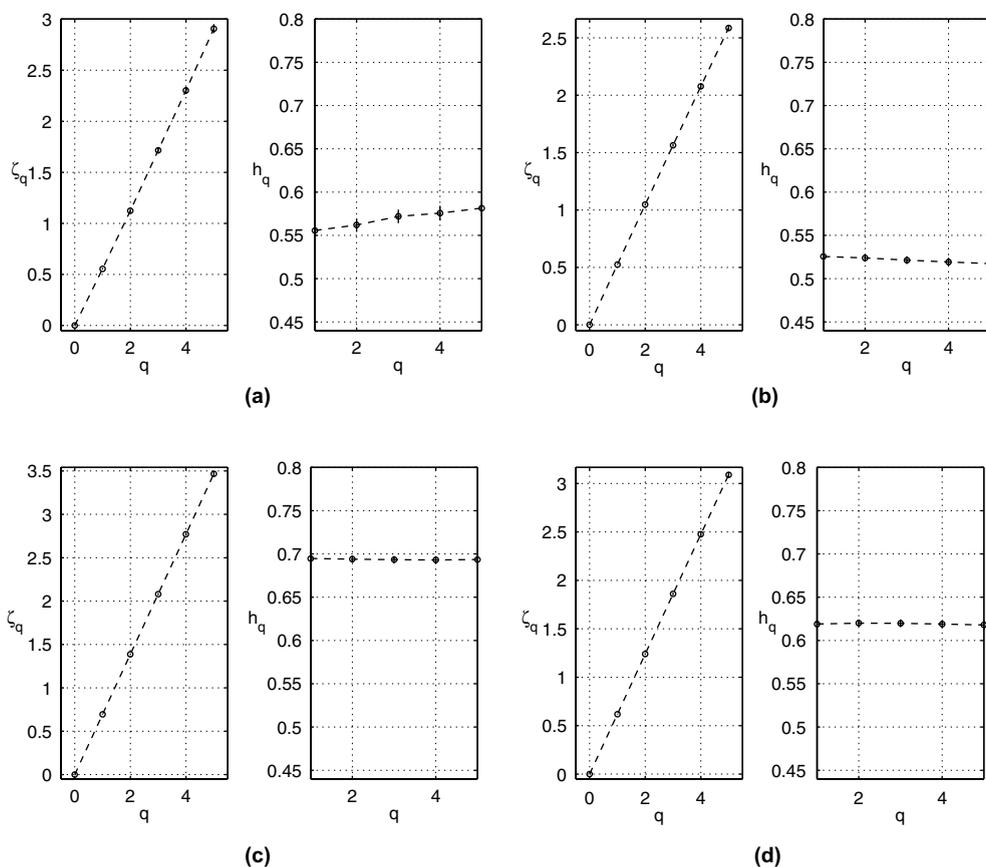


Fig. 5. Multifractal analysis: (a) OC3-tier2-dom, (b) OC12-tier2-dom, (c) OC12-corp-dom, (d) OC48-bb-1.

approximately linear. We employ the range 0.8–102.4 ms for the OC12 and OC48 traces and the 6.4 ms–25.6 ms range for the OC3 traces. We 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 nearly a constant, typical of monofractal processes. This finding is in contrast to earlier findings on traffic multifractality [5,6,8,10]. This apparent contradiction is easily demystified through an analysis of the Gaussianity of the marginal distribution of the traces concerned.

Kurtosis and *skewness* are two metrics commonly used to test for Gaussianity. 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 (1–10 Mbps) links, which transport less

aggregated traffic with non-Gaussian marginals. For example, one study reported kurtosis values greater than 5 at a time scale as large as 500 ms for traces collected at Auckland university [25]. 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 3 and Fig. 6 for marginals at the 4 ms time scale), leading to a trivial monofractal scaling as explained in Section 4.4. This is actually to be

Table 3
Estimated kurtosis and skewness at the 4 ms time scale

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

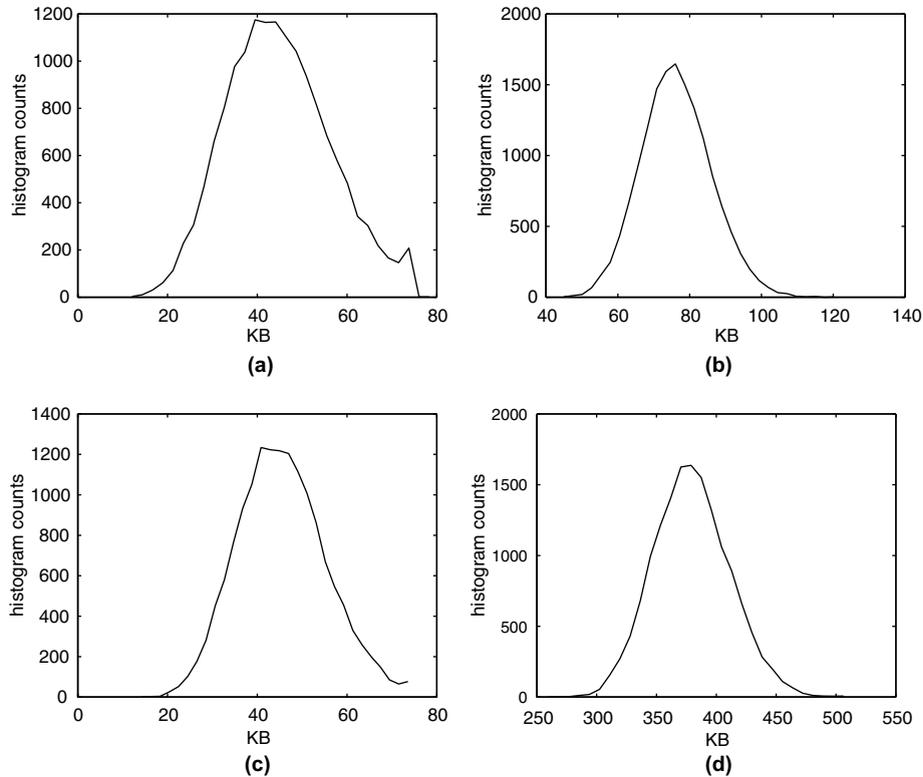


Fig. 6. Histograms of bytes per 4 ms: (a) OC3-tier2-dom, (b) OC12-tier2-dom, (c) OC12-corp-dom, (d) OC48-bb-1.

expected, as the analyzed backbone traffic is highly aggregated (see Table 1 for average bandwidth of the traffic). 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 use wavelet energy plots only to study the small-time (second-order) scaling behaviors of the Internet backbone traffic.

6. Search for the physical causes of scaling

The goal of this section is to uncover the probable causes and factors that influence the small-time scalings of Internet backbone traffic. We focus on two data sets that manifest strikingly different scaling behaviors: the OC12-tier2-dom data (Fig. 2(a)) that displays uncorrelated, or “white noise”-like, small-time scaling and the OC12-corp-dom (Fig. 2(b)) that possesses strongly corre-

lated scaling. Our strategy is to decompose both data sets in different ways to identify the small time scale correlation causing components.

6.1. Flow classification and traffic composition

Clearly, an aggregate 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 byte contributions—the number of bytes contributed by a flow in a given time segment, and flow density—the dominant inter-packet arrival times of a flow in a given time segment. These two classification schemes are qualitatively very different. While flow byte contribution only captures an aggregate or “global” property of the flow, namely the total bytes it contributes to a time segment, flow density captures a “local” property

of the flow, that is how packets are spaced in the flow relative to each other. Using these two flow classification schemes, we analyze the traffic composition of the packet traces, i.e., identifying flow components with different scaling characteristics.

We use the standard 5-tuple flow as our basic flow definition. A 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 segment. Since we are primarily interested in small-time scaling behaviors in the range of 1–100 ms, we zoom into 1-min segments of the packets traces to analyze the characteristics of their flows in detail. A 1-min segment contains enough information to give good estimates of sub-second time scale (1–100 ms) statistics, while being small enough to enable a detailed flow-level analysis. Moreover, we have observed little deviation of the energy plot of a 1-min segment of backbone traffic from the energy plot obtained from a 1-h segment at small time scales.

One way to classify flows is based on flow byte contribution, namely, the total number of bytes belonging to a flow in the 1-min time segment of analysis. We note that flow byte contribution differs from the notion of *flow size*, that is the total bytes transmitted by a flow in its entire lifetime, which was used to explain the correlation or LRD of traffic at large (seconds to minutes) time scales [1,26]. We exclude packets that arrive out-

side the analysis segment because they do not influence the energy plot corresponding to the segment. We introduce two parameters, $B_L \geq B_S$ in order to enable a large vs. small flow classification. We term a flow as large if it transmits more than B_L bytes in a given 1-min segment, and a flow as small if it transmits less than B_S bytes. In our study we typically choose $B_L = 10^6$ bytes and $B_S = 10^4$ bytes.

Fig. 7(a) shows the complementary cumulative distribution function (CCDF) of flow byte contributions in five different 1-min 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 of total bytes in the aggregate traffic, as shown in Fig. 7(b). Fig. 7(b) also reveals that the flow byte contribution distribution is fairly consistent over time, with little variations in five different 1-min segments, each an hour apart.

Our second flow classification is based on the distribution of packet inter-arrival times within a

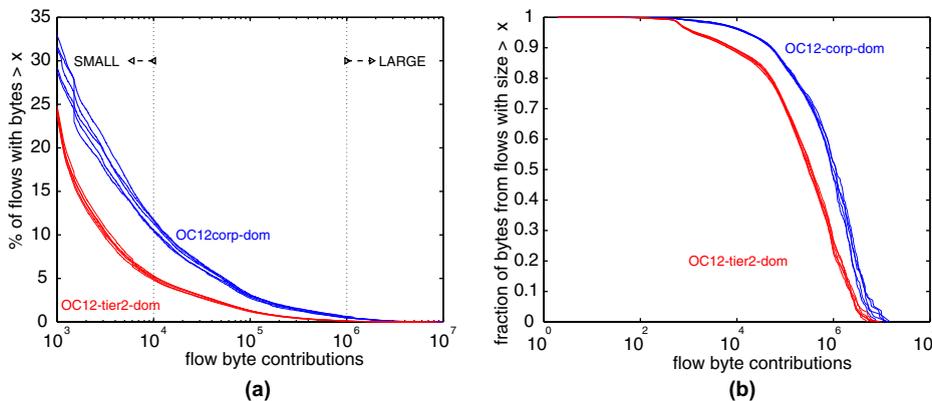


Fig. 7. (a) CCDF of flow byte contribution for a five 1-min time segments separated by 1 h. (b) Fraction of bytes belonging to flows with byte contribution larger than x vs. x for the same time segments.

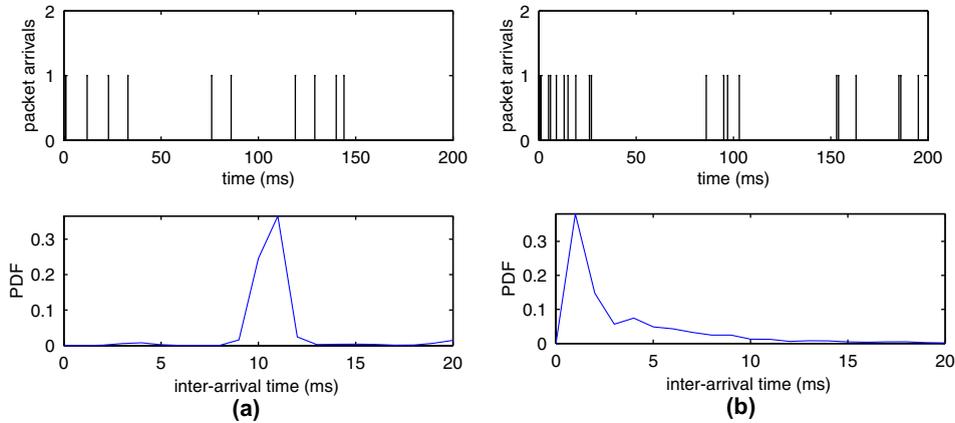


Fig. 8. Estimated probability density function (PDF) of packet inter-arrival times (top) and arrival process (bottom). The sparse flow in (a) has packets spaced farther apart than the dense flow in (b).

flow. Before we introduce the formal definition, we show an example of two different flows in Fig. 8(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 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 call flows with *median* of packet inter-arrival times less than time T *dense*, and the remaining *sparse*. Intuitively dense flows have more closely spaced packets than sparse flows. If we choose $T = 3$ ms for example, then flow a in Fig. 8 will be classified as sparse and flow b as dense. Of 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.

6.2. Small-time scalings of large vs. small flows

We begin by studying how the flow byte contribution affects the small-time scaling behavior. We use a 1-min segment of the two OC12 traces (the “Min 1” segment in Table 4) as an example. The energy plots for the large ($B_L = 10^6$ bytes) and small ($B_S = 10^4$ bytes) flows of the two traces are displayed in Fig. 9. The aggregate of OC12-corp-dom large flows has a small-time scaling exponent $h = 0.69$, similar to that of the total traffic aggregate in Fig. 2(b). In contrast, the aggregate of OC12-corp-dom small flows has a small-time scaling exponent $h = 0.57$, 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

Table 4

Dense vs. sparse traffic compositions in five 1-min segments, each segment belonging to a different hour of the traces ($T = 2$ ms)

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

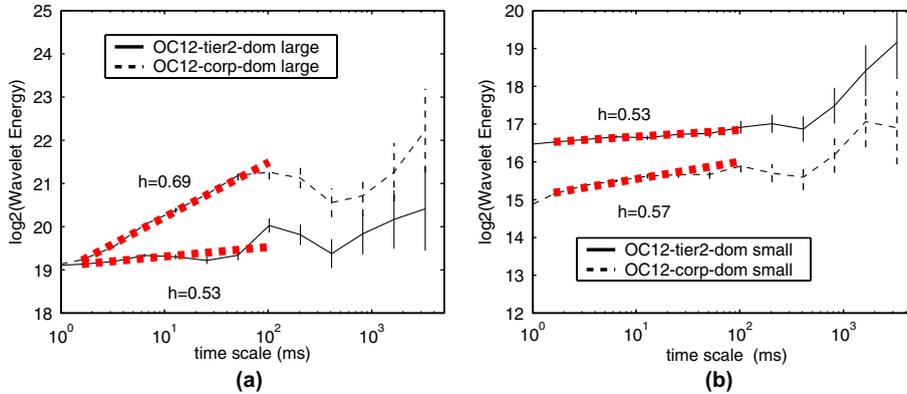


Fig. 9. Wavelet energy plots of (a) large and (b) small flows. The small flows consistently display uncorrelated small-time scaling while the large flows show correlated scaling for one trace and uncorrelated scaling for the other.

scaling with h fairly close to 0.5, similar to the total traffic aggregate (Fig. 2(a)).

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 not necessarily induce correlated or uncorrelated small-time scalings, as is evident in Fig. 9. Hence flow byte contribution alone does not determine the small-time scaling of aggregate traffic.

6.3. 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-min segment from both traces as before, we decompose the total traffic into aggregates of dense and sparse flows by setting $T = 2$ ms. The corresponding energy plots are shown in Fig. 10. 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. 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

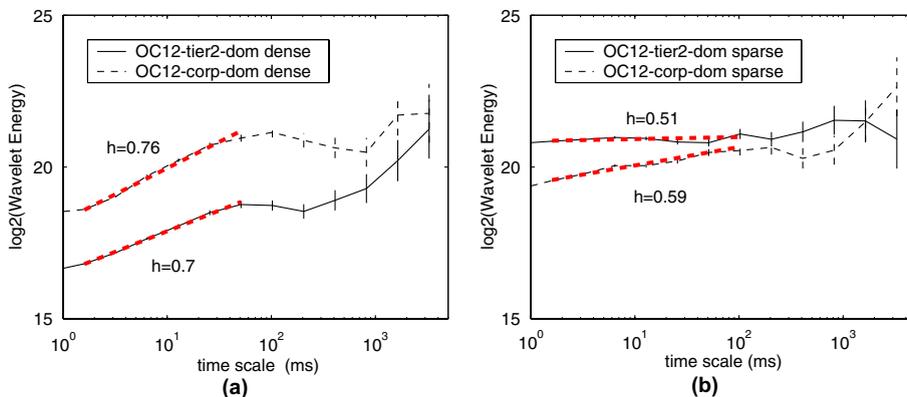


Fig. 10. Small-time scaling of (a) dense vs. (b) sparse flows. The dense flows are strongly correlated at small-time scales while sparse flows are nearly uncorrelated.

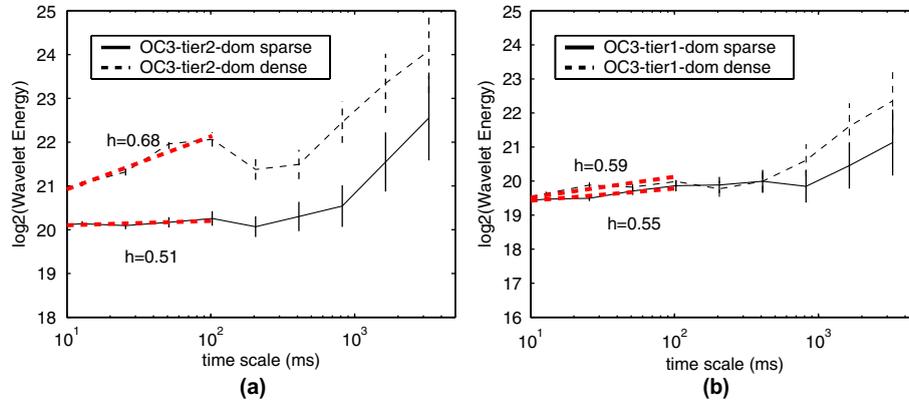


Fig. 11. Small-time scaling of dense vs. sparse flows for OC3 traces. The dense flows are more strongly correlated at small-time scales than the sparse flows: (a) OC3-tier2-dom, (b) OC3-tier1-dom.

that of the total traffic aggregate ($h \approx 0.71$, see Fig. 2(b)). We note that the results reported here hold qualitatively as long as we classify dense/sparse flows in a reasonable fashion, e.g., setting $T = 4$ ms instead of 2 ms. Due to space limitations, we do not show such results here.

Analysis of several other traces confirm the above finding. From the scaling of OC3 (Fig. 11) and OC48 traces (Fig. 12) we observe that dense flows indeed possess scaling exponents larger than those of the sparse flows. While the difference in scaling between dense and sparse flows is more pronounced for some traces (e.g. Figs. 11(a) and 12(a) and (b)) it is less pronounced for others (e.g. Figs. 11(b) and 12(c)). Clearly further analysis and mathematical models beyond that presented in this paper are required to determine what influences the strength of correlation of the dense flows.

6.4. Impact of traffic composition on small-time scalings

So far we have demonstrated that a 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 a dense flow component, why do they exhibit strikingly different small-time scalings? The answer to this question is obvious when we examine their

traffic compositions in terms of dense and sparse flows.

Using $T = 2$ ms to classify dense and sparse flows, Table 4 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.

Table 4 also explains why the scaling of different traces does not change over time. Since the percentage of dense and sparse traffic does not change by much over time their relative contribution to the scaling behavior is time invariant.

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 the small-time scalings change with the traffic composition. We start with the aggregate of all dense flows in the OC12-tier2-dom trace, and observe the change in small-time scaling behaviors by progressively adding in sparse flows: from 0% to 100%

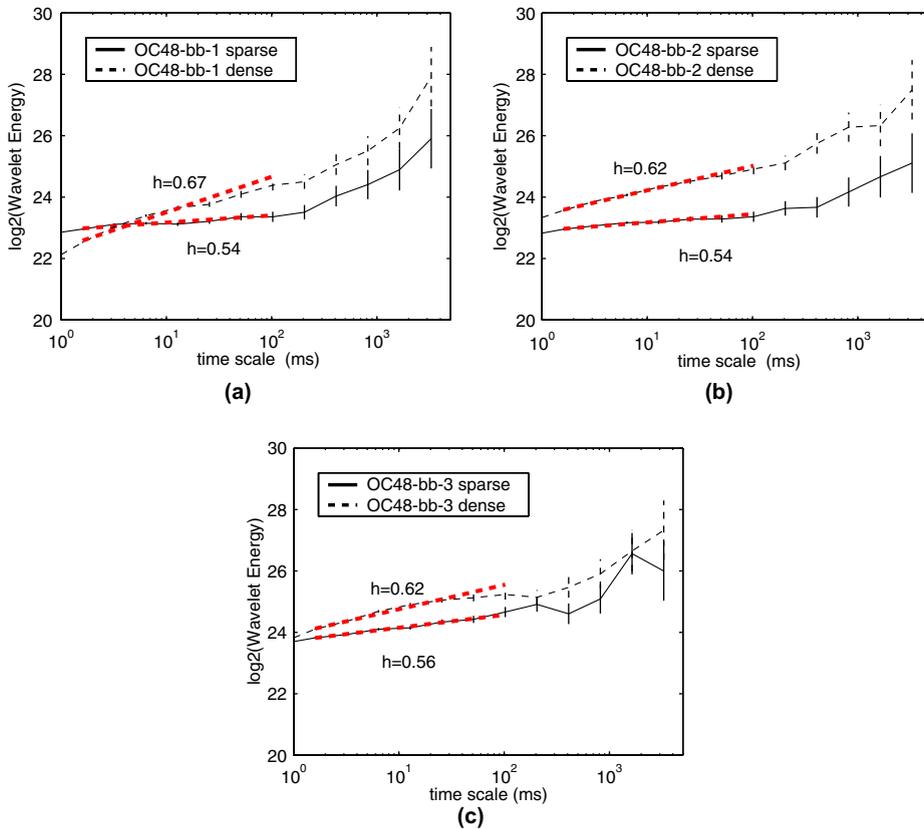


Fig. 12. Small-time scaling of dense vs. sparse flows for OC-48 traces. The dense flows are more strongly correlated at small-time scales than the sparse flows: (a) OC48-bb-1, (b) OC48-bb-2, (c) OC48-bb-3.

of all sparse flows in the trace. The results are shown in Fig. 13(a), where we see that the scaling exponent h in the time-scale range of 1–10 ms decreases from 0.75 to 0.54, as the proportion of sparse flows increases.

We now show how the 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 and gradually adding more dense flows into the traffic aggregate: from 0% to 100% of all sparse flows in the trace. The results are shown in Fig. 13(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 flows increases. These two semi-experiments clearly demonstrate the impact of dense vs. sparse traffic composition on the small-time scaling behaviors.

6.5. Effect of RTT on small-time scaling

We conclude this section by examining the effect of the RTT on small-time scalings of the traces. Simulations have shown that strong periodicities at the RTT time scale caused by TCP's feedback control can lead to a decrease or 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 [27,28]. To address this issue, we obtain RTT estimates for a large sample of TCP flows using the method described in [29]. In both traces, the RTTs for both dense and sparse flows are within the range of around 10 ms–1 s or so (Fig. 14). 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

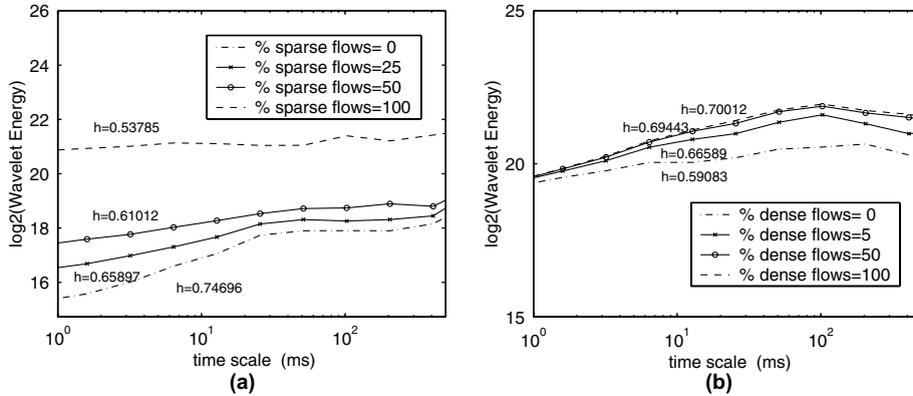


Fig. 13. Impact of dense vs. sparse traffic composition on small-time scalings. In (a) the scaling changes from strongly correlated to nearly uncorrelated with increasing percentage of sparse flows. In (b) the scaling becomes more correlated with increasing percentage of dense flows. (a) OC12-tier2-dom, (b) OC12-corp-dom.

large number of RTTs at about 10 ms, 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 in the frequency domain (see Fig. 15) corroborates the RTT analysis. We do not observe any strong periodicities that may be caused by the RTT. Instead we observe the striking $1/f$ -type decay for the dense flows in the frequency range 10–200 Hz (5–100 ms time scales) and the relatively flat spectrum in the same range for the sparse flows. This

more explicitly accounts for difference in the small-time scaling behaviors (recall the relationship between wavelet energy and spectrum density function in Section 3). Note that for the sparse flows of the OC12-corp-dom trace, the spectrum in frequency range 10–200 Hz is not as flat as that of the sparse flows of the OC12-tier2-dom, which explains the slightly higher scaling exponent we observe (see Fig. 10(b)).

Based on our analysis we now speculate on the physical causes of flow density. Intuitively, dense flows, by definition, inject closely spaced packets into the network, causing strong correlation in small time scales. We believe that flow density is

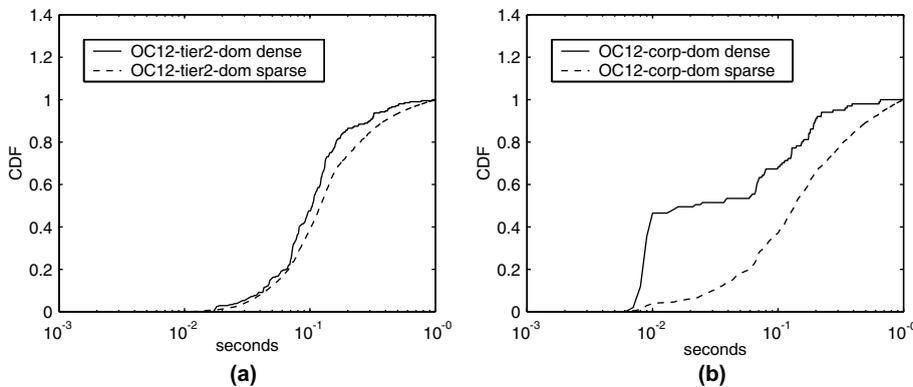


Fig. 14. RTT estimates from 3-way TCP handshake packets. The dense flows of the two traces have dissimilar RTT distributions though their scaling behavior is similar. Also the RTTs of the dense and sparse flows of OC12-tier2-dom are similar though their scaling is dissimilar: (a) OC12-tier2-dom, (b) OC12-corp-dom.

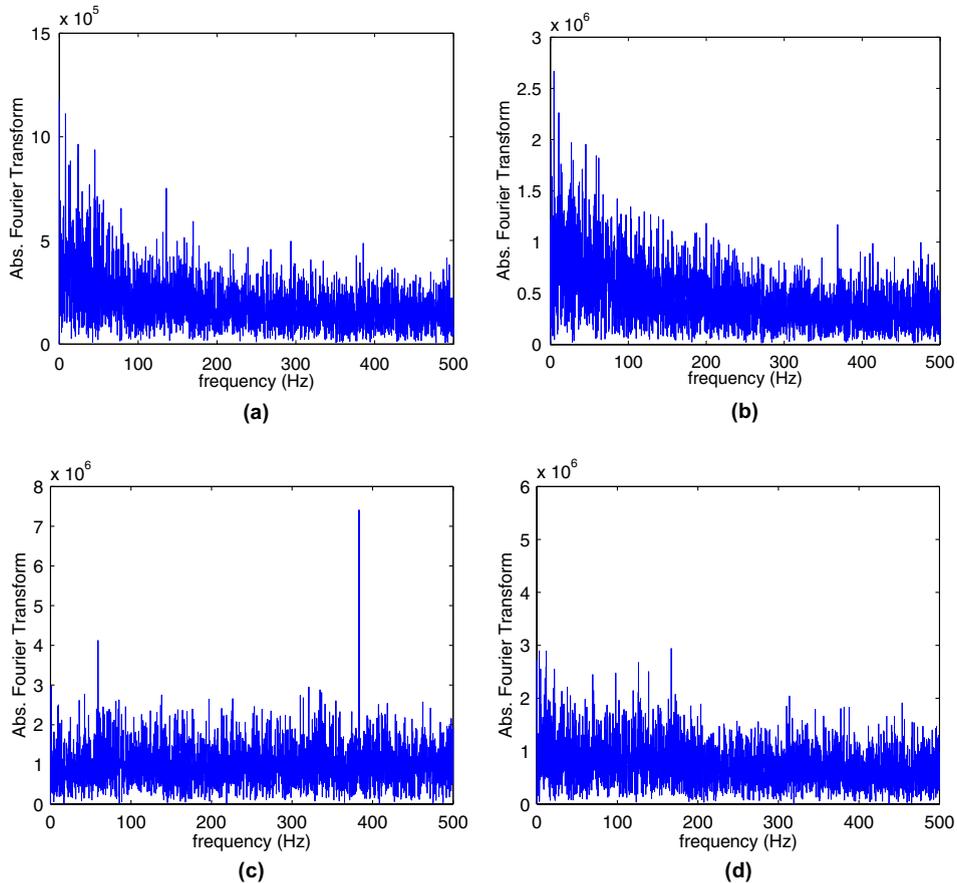


Fig. 15. Fourier transform of dense and sparse flows. The dense flows show a marked $1/f$ frequency decay while the sparse flows do not. (a) OC12-tier2-dom dense, (b) OC12-corp-dom dense, (c) OC12-tier2-dom sparse, (d) OC12-corp-dom sparse.

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 the TCP feedback control mechanism alone. Since the packet spacing on a link is inversely proportional to the link bandwidth, we can expect higher bottleneck bandwidths to allow more dense packet clustering.

7. Related work

The dichotomy of scaling in the energy plots of WAN traces has been observed in an earlier study [7]. Recently multifractal small-time scalings of various network traffic statistics have been re-

ported 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 [27], 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 [25], which proposes the alpha vs.

beta flow dichotomy in terms of flow marginals at the 500 ms time scale, and study the impact of alpha/beta flows on queuing performance. It also points to the bottle-neck link speed as a potential factor in producing alpha and beta flows. Our work differs from [25] 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.

Another set of studies based on the analysis of inter-packet arrival times concludes that packet arrivals and sizes can go *locally* to independence with increased statistical multiplexing while maintaining the correlation structure of the time series of byte counts [30,31]. Unlike these studies which focused on the packet interarrivals of the *aggregated* traffic, we focus on the packet interarrivals within *individual* flows and demonstrate their influence on the time series of byte counts through the concept of flow density.

By modeling inter-packet arrival times within flows as independent random variables, recent analysis shows that traffic displays uncorrelated small-time scaling behavior [32]. This work, however, does not address the scenarios where traffic can indeed be correlated at small time scales.

A recent study corroborates our finding that the scaling structure at small scales has its origin in the packet patterns within flows [33]. With the help of an analytical traffic model it was shown that increasing the coefficient of variation of inter-packet arrival times of individual flows leads to larger scaling exponents at small time scales. Our work differs in that we study the effect of the *median* of inter-packet arrival times on scaling behavior rather than the coefficient of variation. In addition our work is purely empirical and makes no analytical modeling assumptions. Further modeling and analysis is required in order to make a more detailed comparison between the two studies.

8. Conclusions

From a detailed wavelet analysis of traffic collected at OC3/12/48 links on the Sprint network, we have observed that for a majority of links,

the (second-order) scaling parameters at small time scales (1–100 ms) are fairly close to 0.5. Hence for these links simple white noise models suffice in describing the small time scale traffic fluctuations. In addition, the traces manifest mostly monofractal (instead of multifractal) behaviors at the 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 the small time scales.

However, a small number of data sets do exhibit correlation at small time scales (with scaling exponents within the range of 0.6–0.7), thus raising the question as to when to expect white noise behavior and when not to. To understand the potential causes or factors that influence the small-time scalings, we analyze the traffic composition of the traces along two dimensions—flow byte contribution and flow density. We find that flow byte contribution alone cannot determine the small-time scaling of aggregate traffic. Instead our results point to dense flows as the correlation-causing factor in small time scales and reveal that the traffic composition in terms of proportions of dense vs. sparse flows decides the overall small-time scalings of aggregate traffic. Further investigations and mathematical models are required to provide an explanation for and shed theoretical insights into our findings.

Our observations and results have significant implications in network modeling, service provisioning and traffic engineering. For example, the discovery of uncorrelated small-time scaling behaviors on many Internet backbone links can lead to simpler network models for analyzing network performance at small time scales [11,13]. Our discovery that small-time scalings are determined by the dense 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 to more likely appear “independent” at those small time scales that impact queuing [11]. On the other hand, as broadband access becomes more widely deployed, large files and objects can be transmitted faster

into the Internet, thus causing flows that are more dense. In addition, the changing nature of applications and increased use of protocols other than TCP can also effect a change in the small-time scaling behaviors of the Internet backbone traffic. These are important questions awaiting to be explored.

Acknowledgments

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. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

References

- [1] K. Park, W. Willinger (Eds.), *Self-Similar Network Traffic and Performance Evaluation*, Wiley Interscience, New York, 2000.
- [2] W. Leland, M. Taqqu, W. Willinger, D. Wilson, On the self-similar nature of Ethernet traffic (extended version), *IEEE/ACM Transactions on Networking* 2 (1994) 1–15.
- [3] V. Paxson, S. Floyd, Wide-area traffic: the failure of Poisson modeling, *IEEE/ACM Transactions on Networking* 3 (1995) 226–244.
- [4] M. Crovella, A. Bestavros, Self-similarity in World Wide Web traffic. Evidence and possible causes, *IEEE/ACM Transactions on Networking* 5 (1997) 835–846.
- [5] J. Lévy Véhel, R. Riedi, Fractional Brownian motion and data traffic modeling: the other end of the spectrum, in: *Fractals in Engineering*, Springer, 1997, pp. 185–202.
- [6] A. Feldmann, A.C. Gilbert, W. Willinger, Data networks as cascades: investigating the multifractal nature of Internet WAN traffic, *ACM SIGCOMM*, 1998, pp. 42–55.
- [7] A. Feldmann, A.C. Gilbert, W. Willinger, T. Kurtz, The changing nature of network traffic: scaling phenomena, *Computer Communication Review* 28 (1998) 5–19.
- [8] P. Mannersalo, I. Norros, Multifractal analysis of real ATM traffic: a first look, *COST257TD*, 1997.
- [9] R.J. Gibbens, Y.C. Teh, Critical time and space scales for statistical multiplexing, in: *16th International Teletraffic Congress-ITC16*, 1999, pp. 87–96.
- [10] R. Riedi, M.S. Crouse, V. Ribeiro, R.G. Baraniuk, A multifractal wavelet model with application to TCP network traffic, *IEEE Transactions on Information Theory*, Special Issue on Multiscale Statistical Signal Analysis and its Applications 45 (1999) 992–1018.
- [11] A.L. Neidhardt, J.L. Wang, The concept of relevant time scales and its application to queuing analysis of self-similar traffic, *ACM SIGMETRICS*, 1998, pp. 222–232.
- [12] M. Grossglauser, J.-C. Bolot, On the relevance of long-range dependence in network traffic, *ACM Computer Communication Review* 26 (1996) 15–24.
- [13] C. Fraleigh, F. Tobagi, C. Diot, Provisioning IP backbone networks to support delay-based service level agreements, *IEEE INFOCOM*, vol. 1, 2003, pp. 375–385.
- [14] A. Arvidsson, P. Karlsson, On traffic models for TCP/IP, in: P. Key, D. Smith (Eds.), *Teletraffic Engineering in a Competitive World*, Proc. ITC-16, Edinburgh, UK, North-Holland, Amsterdam, 1999, pp. 457–466.
- [15] Z.-L. Zhang, V. Ribeiro, S. Moon, C. Diot, Small-time scaling behaviors of Internet backbone traffic: an empirical study, *IEEE INFOCOM*, vol. 3, 2003, pp. 1826–1836.
- [16] C. Fraleigh, S. Moon, C. Diot, B. Lyles, F. Tobagi, Packet-level traffic measurements from a tier-1 IP backbone, *IEEE Network* 17 (2003) 6–16.
- [17] D.R. Cox, Long-range dependence: a review, *Statistics: An Appraisal* (1984) 55–74.
- [18] J. Beran, *Statistics for Long-Memory Processes*, Chapman and Hall, London, 1994.
- [19] A. Benassi, S. Deguy, Multi-scale fractional Brownian motion: definition and identification, Preprint, 1999.
- [20] J.M. Bardet, P. Bertrand, Identification of the multiscale fractional Brownian motion with biomechanical applications, Preprint of L.S.P., University Toulouse III, 2003.
- [21] R. Riedi, Multifractals: an introduction, Research report, Rice University, 1997. Available from: <<http://www.dsp.rice.edu/1996>>.
- [22] P. Abry, P. Flandrin, M.S. Taqqu, D. Veitch, Wavelets for the analysis, estimation, and synthesis of scaling data, in: K. Park, W. Willinger (Eds.), *Self-Similar Network Traffic and Performance Evaluation*, Wiley Interscience, New York, 2000.
- [23] P. Abry, D. Veitch, Wavelet analysis of long-range-dependent traffic, *IEEE Transactions on Information Theory* 44 (1998) 2–15.
- [24] D. Veitch's home page. Available from: <<http://www.emu-lab.ee.mu.oz.au/~darryl/>>.
- [25] S. Sarvotham, R. Riedi, R. Baraniuk, Connection-level analysis and modeling of network traffic, *IEEE/ACM Internet Measurement Workshop*, 2001, pp. 99–103.
- [26] W. Willinger, M.S. Taqqu, R. Sherman, D.V. Wilson, Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level, *ACM Computer Communication Review* 25 (1995) 100–113.
- [27] A. Feldmann, A.C.G.P. Huang, W. Willinger, Dynamics of IP traffic: a study of the role of variability and the impact of control, *ACM SIGCOMM*, 1999, pp. 301–313.
- [28] P. Huang, A. Feldmann, W. Willinger, A non-intrusive, wavelet-based approach to detecting network performance problems, *IEEE/ACM Internet Measurement Workshop*, 2001, pp. 213–227.

- [29] H. Jiang, C. Dovrolis, Passive estimation of TCP round-trip times, *ACM Computer Communications Review* 32 (2002) 75–88.
- [30] J. Cao, W.S. Cleveland, D. Lin, D.X. Sun, On the nonstationarity of Internet traffic, *ACM SIGMETRICS*, 2001, pp. 102–112.
- [31] J. Cao, W.S. Cleveland, D. Lin, D.X. Sun, The effect of statistical multiplexing on Internet packet traffic: theory and empirical study, *Bell Labs Technical Report*, 2002.
- [32] K. Drakakis, A detailed mathematical study of several aspects of the Internet, Ph.D. Thesis, Princeton University, 2003.
- [33] N. Hohn, D. Veitch, P. Abry, Cluster processes a natural language for network traffic, *IEEE Transactions on Signal Processing*, Special Issue on Signal Processing in Networking 51 (2003) 2229–2244.



Vinay Ribeiro received his B.Tech. degree in electrical engineering from the Indian Institute of Technology Madras, India, in 1997 and his M.S. degree in electrical engineering from Rice University in 1999 where he is currently pursuing his Ph.D. He was an invited fellow at the Institut Mittag-Leffler in Sweden in Fall 2004 and has held research internship positions at AT&T Labs and Sprint Advanced Technology Labs. He received the best student paper award at the Passive and Active Measurement Workshop 2003, a Texas Instruments Fellowship from Rice University, and is a member of Eta Kappa Nu.



Zhi-Li Zhang received the B.S. degree in computer science from Nanjing University, China, in 1986 and his M.S. and Ph.D. degrees in computer science from the University of Massachusetts in 1992 and 1997. In 1997 he joined the Computer Science and Engineering faculty at the University of Minnesota, where he is currently an Associate Professor. From 1987 to 1990, he conducted research in Computer Science Department at Aarhus University, Denmark, under a fellowship from the Chinese National Committee for Education. He has held visiting positions at Sprint Advanced Technology Labs; IBM T.J. Watson Research Center; Fujitsu Labs of

America, Microsoft Research China, and INRIA, Sophia-Antipolis, France. His research interests include computer communication and networks, especially the QoS guarantee issues in high-speed networks, multimedia and real-time systems, and modeling and performance evaluation of computer and communication systems.

He currently serves on the Editorial boards of *IEEE/ACM Transactions on Networking* and *Computer Network*, an International Journal. He is co-chair of IEEE/IFIP International Workshop on QoS 2004, and is co-chair of IEEE INFOCOM 2006 in Barcelona Spain. He was a Program Co-chair of the SPIE ITCOM 2002 conference on Scalability and Traffic Control in IP Networks, served on the Executive Committee for IEEE Infocom 2001 and Infocom 2003, and on the Technical Program Committees of various conferences and workshops including IEEE Infocom and ACM SIGMETRICS. He received the National Science Foundation CAREER Award in 1997. He has also awarded the prestigious McKnight Land-Grant Professorship at the University of Minnesota. He is a co-recipient of an ACM SIGMETRICS best paper award and an IEEE ICNP best paper award. He is a member of IEEE, ACM and INFORMS Telecommunication Section.



Sue Moon received her B.S. and M.S. from Seoul National University, Seoul, Korea, in 1988 and 1990, respectively, all in computer engineering. She received a Ph.D. degree in computer science from the University of Massachusetts at Amherst in 2000. From 1999 to 2003, she worked in the IPMON project at Sprint ATL in Burlingame, California. In August of 2003, she joined KAIST as an assistant professor and now teaches in Daejeon, Korea. Her research interests are in network performance measurement and monitoring of diverse network types and their security and anomalous aspects.



Christophe Diot received a Ph.D. degree in Computer Science from INP Grenoble in 1991. From 1993 to 1998, he was a research scientist at INRIA Sophia Antipolis, working on new Internet architecture and protocols. From 1998 to 2003, he created and managed the IP research group at Sprint Advanced Technology Labs. Diot recently moved to INTEL research in Cambridge, UK. He is still active in the measurement community. However, he has a growing interest on understanding how the Internet is gonna survive mobility and wireless technology.