



HAL
open science

A Longitudinal Study of Small-Time Scaling Behavior of Internet Traffic

Himanshu Gupta, Vinay J. Ribeiro, Anirban Mahanti

► **To cite this version:**

Himanshu Gupta, Vinay J. Ribeiro, Anirban Mahanti. A Longitudinal Study of Small-Time Scaling Behavior of Internet Traffic. 9th International IFIP TC 6 Networking Conference (NETWORKING), May 2010, Chennai, India. pp.83-95, 10.1007/978-3-642-12963-6_7 . hal-01056299

HAL Id: hal-01056299

<https://inria.hal.science/hal-01056299v1>

Submitted on 18 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A Longitudinal Study of Small-Time Scaling Behavior of Internet Traffic

Himanshu Gupta^{1,2}, Vinay J. Ribeiro², and Anirban Mahanti³

¹ IBM Research Laboratory, New Delhi, India
higupta8@in.ibm.com

² Department of Computer Science and Engineering,
Indian Institute of Technology, New Delhi, India
vinay@cse.iitd.ac.in

³ NICTA, Locked Bag 9013, Alexandria, NSW 1435, Australia
anirban.mahanti@nicta.com.au

Abstract. During the last decade, many new Web applications have emerged and become extremely popular. Together, these new “Web 2” applications have changed how people use the Web and the Internet. In light of these changes, we conduct a longitudinal study of the small-time scaling behavior of Internet traffic using network traffic traces, available from the MAWI repository, that span a period of eight years. The MAWI traces are affected by anomalies; these anomalies make correct identification of scaling behavior difficult. To mitigate influence of anomalies, we apply a sketch-based procedure for robust estimation of the scaling exponent. Our longitudinal study finds tiny to moderate correlations at small-time scales, with scaling parameter in the range $[0.5, 0.75]$, across the traces examined. We also find that recent traces show larger correlations at small-time scales than older traces. Our analysis shows that this increased correlation is due to the increase in the fraction of aggregate traffic volume carried by dense flows.

Key words: Traffic Analysis, Small-time scaling, Dense Flows, Robust Estimation

1 Introduction

Scaling behavior of Internet traffic has received much attention from networking researchers. It is known that Internet traffic displays two scaling regimes with a transition occurring in the 100ms to 1s time range [3, 7]. Internet traffic when aggregated to large-time scales ($\geq 1s$) is quite bursty and is modeled using *long-range dependent* (LRD) processes [11, 12]. The scaling parameter (H) in large-time scales is typically in the $(0.8, 1)$ range which is indicative of highly correlated packet arrivals [3, 6, 11, 12].

The focus of this paper is on the small-time scaling behavior of recent Internet traffic. By small-time scale, we refer to timescales smaller than the scaling transition point, which is in the 100 to 1000ms range, wherein factors that cause

correlations at large-time scales may not yet be strong or even present. Traffic models proposed for small-time scales include simple Poisson processes, independent Gamma interarrivals, and non-Gaussian multifractal processes [5,9]. Zhang et al. [16] examined traces from backbone networks (collected in 2000-02) and found most of the examined traces to exhibit little or no correlations at small-time scales with a scaling parameter h below 0.6, and often fairly close to 0.5.⁴ Only a small fraction of the traces examined exhibited some small-time correlations with scaling parameter within the range of 0.6-0.7. They also introduced the concept of “dense flows”, defined as flows with bursts of densely clustered packets, and showed that traces with more dense flows display relatively larger correlations at small-time scales.

In this paper, we present a longitudinal analysis of the small-time scaling phenomena. Our work is driven by the observation that a number of significant changes have occurred since the beginning of 2000 (when prior work in this area was undertaken). In particular, typical backbone network capacities as well as number of Internet connected hosts have increased. Furthermore, composition of traffic has changed with the rapid growth of “Web 2” sites such as YouTube, FaceBook, Flickr, and Twitter, and also with popularity of peer-to-peer file sharing applications [2]. Therefore, we revisit aspects concerning the small-time scaling behavior of Internet traffic. Specific questions that we consider include the following: Do recent traces mostly display no/small/moderate correlations at small-time scales? Do dense flows still drive small-time scaling behavior of Internet traffic? In this paper, we provide answers to these questions by conducting a longitudinal analysis of publicly available Internet traffic traces from the MAWI repository that span 8 years (2001-2009).

The traces in the MAWI repository are known to contain anomalies [4]. In a recent work, Borgnat et al. [3] have shown that anomalies can interfere with the correct identification of scaling behavior. They proposed a robust estimation method based on sketches (random projections) to mitigate the effect of network anomalies and used this estimation method for a longitudinal analysis of the LRD behavior of network traces. We use this robust estimation method for the longitudinal analysis of small-time scaling behavior. We show that in absence of a robust estimation procedure, misleading inferences regarding small-time scaling behavior of Internet traffic may be derived. We believe our work further reinforces the importance of this new method.

Our primary contribution is a study of the evolution of small-time scaling behavior of the MAWI traces throughout this decade. Once the effects of network anomalies have been mitigated, we find the scaling parameter in small-time scales to be consistently between 0.5 and 0.75, thereby, showing the presence of tiny to moderate correlations in small-time scales in the traces examined. We also find that the recent MAWI traces exhibit comparatively larger correlations at small-time scales than the earlier traces. Among the traces examined, traces from 2007 and later years tend to exhibit slightly higher correlation (e.g., h of

⁴ Scaling parameter in small-time scales is represented as h while in large-time scales is represented as H

0.6 or more) than those prior to 2007. This observation is not in consonance with the prediction that Internet traffic will likely be described by simple models (e.g. Poisson) [10]. We further show that dense flows account for a larger fraction of the aggregate traffic in recent MAWI traces vis-a-vis traces from earlier years.

We would like to emphasize here that the results presented in this paper are specific to the MAWI traces. Generalizing the trends regarding Internet traffic by studying packet traces from only one link is difficult as any observations derived from such an analysis run the risk of getting biased by individual events happening at the concerned link. Nonetheless, our results are still valuable as they summarize the evolution of small-time scaling behavior for an Internet backbone for a period of 8 years, and hence can act as a reference for longitudinal studies on other backbone links.

The rest of this paper is organized as follows. Section 2 presents necessary background on the approach used for estimating the scaling parameter, and also summarizes the traces used in this work. Section 3 illustrates the importance of the robust estimation method in the context of scaling at small-time scales. Section 4 focuses on results from our longitudinal analysis. Section 5 analyzes the influence of dense flows on the small-time scaling behavior of the aggregate traffic. Conclusions are presented in Section 6.

2 Background

2.1 The MAWI Dataset

We use publicly available traces from the MAWI repository [1]. This repository provides traces collected from a trans-Pacific backbone. A 15-minute long trace is made public for download every day. We use the traces captured at collection points B and F. The link corresponding to B (100 Mbps) was replaced by F (100 Mbps) in July 2006, and was subsequently upgraded to 150 Mbps in June 2007. For our longitudinal study, we select the traces collected on 1st and 15th of every month, from January 2001 to December 2008. This gives us a set of 180 traces spanning eight years. Each trace is partitioned into two subtraces, one for traffic flowing in each direction referred to as JptoUS and UStoJp. We analyze each subtrace separately.

Overall, the traces exhibit substantial variabilities during the time period we consider. For example, a global increase of throughput from 100 kbps in 2001 to more than 12 Mbps in 2009 is observed. Several long lasting congestion periods are also observed (e.g., UStoJp:2003/04 to 2004/10, UStoJp:2005/09 to 2006/06, JptoUS:2005/09 to 2006/06) [3]. Strong fluctuations in packet number are observed on UStoJp from 2004/07 to 2005/04 due to massive activities of the Sasser worm. Many anomalies such as ping floods and SYN scans have also been observed [3, 4]. Overall, this suite of 180, 15-minute, packet traces is uniformly spread across eight years, and hence likely captures various variations and anomalies one expects to see in Internet traffic traces. We refer the reader to [3, 4] for a detailed statistical characterization of traces available from the MAWI repository.

2.2 Scaling Analysis

The analysis involves studying the data at multiple timescales and hence the term “scaling”. The analysis of the scaling behavior of a process helps to characterize the way the process will aggregate over different time intervals. We say a process Y scales with a scaling parameter h if

$$\text{Variance}(Y^{(m)}) \approx m^h \text{Variance}(Y) \quad (1)$$

where $Y^{(m)}$ is the aggregated process of process Y and is defined as follows:

$$Y^{(m)}(k) = \frac{1}{m} \sum_{i=km-m+1}^{i=km} Y(i), k = 1, 2, \dots, \infty \quad (2)$$

In scaling analysis of network traces, the process Y usually represents the number of bytes or packets sent within a fixed length time interval. The scaling parameter h measures the strength of correlations present in the data. A value of 0.5 represents the absence of correlations; a Poisson process displays a scaling parameter of 0.5. The larger the value of h , the stronger the correlations. We interpret scaling parameter values lying within the ranges $[0.5, 0.6)$, $[0.6, 0.7)$, and $[0.7, 0.75]$ as tiny, small, and moderate correlations, respectively.

Of all the methods for estimating the value of scaling parameter (e.g., Variance-Time plot, R/S plot etc), Wavelet-based method is the most preferred [15]. Other methods are not reliable in the presence of non-stationary trends in the data. An implementation of Wavelet-based method is provided at [13] and we use this software to carry out our analysis. A concise summary of this method follows.

2.3 The Wavelet Estimator for the Scaling Parameter

This method allows us to observe the scaling behavior of a traffic process over a certain range of timescales. The working of this estimator is illustrated using *Haar wavelets* which are the simplest kind of wavelets. In practice any other kind of wavelet can be considered.

Consider a reference timescale T_0 and let $T_j = 2^j T_0$ for $j = 1, 2, \dots$ be increasingly coarser timescales. To form the process at scale j , we partition the trace into consecutive and non-overlapping time intervals of size T_j and count the numbers of bytes (or packets) in these intervals. If t_i^j is the i^{th} time interval at scale $j > 0$ then t_i^j consists of the intervals t_{2i}^{j-1} and t_{2i+1}^{j-1} . Let X_i^j be the amount of traffic in t_i^j , with $X_i^j = X_{2i}^{j-1} + X_{2i+1}^{j-1}$. The *Haar wavelet coefficients* d_i^j at scale j are defined as:

$$d_i^j = 2^{-j/2} (X_{2i}^{j-1} - X_{2i+1}^{j-1}), i = 1, \dots, N_j \quad (3)$$

where N_j is the number of wavelet coefficients at scale j . The *energy* ξ_j at scale j is defined as:

$$\xi_j = E[(d_i^j)^2] \approx \frac{\sum_i (d_i^j)^2}{N_j}. \quad (4)$$

Plotting the logarithm of energy ξ_j as a function of timescale j gives us a *logscale diagram* (LD). The magnitude of ξ_j increases with variability of the traffic process X^{j-1} at scale $j-1$. Variation of energy ξ_j with j captures the scaling behavior of the process. The slope of this energy plot α estimates the *scaling parameter* h through $\alpha=2h-1$. The wavelet energy can only be estimated at a scale because of limited data. Hence the slope of logscale diagram can itself only give an estimate of the scaling parameter.

If the scaling parameter is constant across all timescales, the process is said to show *global scaling* [8,14]. If the slope of the energy plot is roughly constant over a range of timescales j to $j+k$, the traffic process is said to exhibit *local scaling* in the timescales T_j to T_{j+k} [8,14]. This paper is concerned with local scaling at small-time scales (1-100 ms).

2.4 Robust Estimation Technique

Borgnat et al. [3] recently showed that anomalies (e.g., attack, congestion etc.), if present, in traffic traces can radically affect the observed scaling behavior. By looking at such a trace, one can draw misleading inferences regarding the scaling behavior of the traffic. For example, changes in scaling behavior may be incorrectly attributed to network mechanisms at work rather than to network anomalies. Hence, it is important to mitigate the effect of anomalies so that the impact of networking mechanisms can be precisely studied. This is especially important if one is trying to disentangle smooth long-term evolution features from day-to-day fluctuations.

Borgnat et al. [3] presented a sketch-based method for robust estimation of the scaling parameter. Let f denote a hash table of size M . The original collection of packets is split into M sub-collections, each consisting of all packets with identical sketch output $m = f(A)$ where the hashing key A is chosen as one of the packet attributes (IPdst, IPsrc,...). This approach, referred to as “random projections”, preserves flow structures as packets belonging to a given flow are assigned to the same sub-collection. Each sub-collection is aggregated and its scaling parameter is computed. A robust estimate of the scaling parameter is the median of the above values [3].

Statistically, robustness in estimation is achieved by performing averages over independent copies of equivalent data. Finding equivalent traces is a complex problem. Random projection using sketches is one way to achieve independent copies of equivalent traces. The resulting logscale diagrams have the same shape as the original, with a variance which is appropriately scaled down⁵, consistent with an independent and identically distributed (i.i.d.) superposition model [7]. In the presence of anomalies, sketching the original packet stream reduces their impact, possibly restricting them to only some of the sub-collections. The small-time correlation structure of the traffic in the sub-collections containing anomalies would likely differ from normal traffic as well as traffic in other

⁵ If one selects flows with probability 0.7, the resulting LD will have a variance that is approximately equal to 70% of the original.




Fig. 1. Representative Examples: Sketch based Robust Estimation

sub-collections. Taking the median of h over independent sketches achieves robustness. The median is chosen instead of mean as median is a non-linear statistic that provides robustness against outliers. The estimator can still be fooled if anomalies are a dominant component of the trace. Robustness in such cases can be achieved by maintaining multiple sketches and taking the median over estimates computed from them.

3 Representative Examples

Fig. 1(a) shows the logscale diagram (upper circled plot) for an anomaly-free trace (July 11, 2005; UStoJp). The anomaly-free nature of this trace was established (in a prior study) through careful manual inspection and application of an anomaly identification algorithm [3]. The time-series of byte counts every millisecond is used to construct the LD (i.e., reference timescale T_0 is 1 ms and X_i^j represents byte count in interval t_i^j). The scaling parameter estimate is 0.65 which indicates the presence of small correlations at small-time scales.

Next, we estimate the scaling parameter using the sketch-based robust estimation method [3]. We hash the trace into 8 bins using destination IP as the hash-key, and estimate the scaling parameter for each subtrace. Robust estimation is achieved by taking median over this set of 8 estimations of the scaling parameter, one corresponding to each subtrace. Fig. 1(a) plots all the subtrace LDs as well. All sketched subtrace LDs are found to be parallel to the original LD, and therefore, indicate a similar scaling behavior as the original trace. We use notations h_g and h_m to represent global and median estimates, respectively. For this anomaly free trace, both h_g and h_m are found to be identical (≈ 0.65).

Fig. 1(b) shows the logscale diagram (upper circled plot) for a trace (Oct 11, 2005; JptoUS) which is known to contain a low-intensity long-lasting spoofed flooding anomaly [4]. The anomaly consists of source IP addresses being spoofed (source IP is identical to destination IP) and destination port being 0 (which is not normally used) [4]. The global scaling estimate h_g is 0.40 which indicates the presence of a small negative correlation (also called *anti-persistent* behavior).




Fig. 2. Global and Median Scaling Parameter: Byte-Count vs Packet-Count

We also estimate the robust (median) value of the scaling parameter. Fig. 1(b) also shows the LDs for all the sketched subtraces. Except for one LD, possibly containing anomaly, all other subtrace LDs have recovered normal behavior, and each of these subtraces now displays tiny positive correlation ($h \approx 0.55$). Note that the scaling behavior of these sketches is markedly different from the scaling behavior shown by the original trace. As a result, the median estimate h_m over 8 sketches, is found to be 0.55 while the global estimate h_g was 0.40. This shows that h_g being 0.40 is an artifact of network anomalies; otherwise all sketched LDs should have had a similar value of scaling parameter.

We next show the value of robust estimation by comparing the values of scaling parameter computed using byte-counts and packet-counts. A debate exists regarding whether scaling analysis should be measured on byte-counts or packet-counts [3]. We hence compute the small-time scaling parameter, both global and median, for all 180 traces in our MAWI tracesuit. Fig. 2 shows scatter plots comparing the values of the scaling parameter using byte-counts and packet-counts. The top row compares global values of the scaling parameter across all traces examined while the bottom row compares the median values. We find that most of the median values of the scaling parameter lie along a 45 degree line indicating that the parameter values as obtained from byte-counts or packet-counts are nearly identical, an observation consistent with prior studies [3, 7, 11] that predict the same scaling parameter for packet and byte counts. In contrast, note that the global values of the scaling parameter show relatively larger variability and




Fig. 3. Scaling behavior of traffic at small-time scales (2001–2009).

dispersion which erroneously suggests inequality of byte-count and packet-count scaling parameter for many traces.

Unless stated otherwise, all results presented in the remainder of this paper are from application of the robust estimation technique with 8 bins and destination IP address as the hashing key. Also, all scaling analysis is done using byte-counts.

4 Longitudinal Analysis of Scaling Behavior

4.1 Importance of Robust Estimation

Fig. 3 plots the scaling parameter values for the traces across eight years, with and without robust estimation. There are multiple traces for which global estimation h_g is found to be less than 0.5. Specifically, from 2005 to mid-2006 the JptoUS subtrace has h_g values consistently less than 0.5, often close to 0.4, suggesting presence of small negative short range correlations. The median values of the scaling parameter h_m , however, computed using the robust estimation procedure, are markedly different with h_m consistently lying close to 0.55. Similarly, for many traces the global scaling parameter h_g is found to be close to 0.8 (e.g., around 2007 UStoJp) suggesting large correlation in small-time scales, in contrast with our current understanding [16]; however, the values of h_m indicate only tiny to moderate correlations. Clearly, robust estimation of the scaling parameter is important, and should not be overlooked when studying small-time scaling behavior of network traffic.




Fig. 4. Results on 24-hour long traces (2003, 2009).

4.2 Evolution of Short Range Correlations

Fig. 3 shows that the median scaling parameter h_m lies between 0.55 and 0.75, thereby indicating that traces across the eight years display tiny to moderate short range correlations. This observation is consistent with that in the original study by Zhang et al [16]. The traces used in their study were collected in 2000–2002, and since then despite many things having changed, the range of the small-time scaling parameter remains unchanged.

A closer look at Fig. 3 reveals that a majority of the post-2007 traces exhibit small to moderate correlations, with corresponding scaling parameter lying between 0.6 and 0.75, whereas a majority of the pre-2007 traces typically have scaling parameters less than 0.6. Our observations for these traces are consistent with those made by Zhang et al. [16] from their traces from 2000–02. We believe this observation provides an evidence against the predictions by Karagiannis et al. [10] that Internet traffic is moving towards simpler to describe models. We further explore this issue in Section 5 wherein dense flows are discussed.

4.3 Analysis of 24-hour Long Traces

The MAWI repository also makes available some 24-hour long traces. We analyze two such traces, one collected on February 27, 2003 and the other collected on March 30, 2009, to complement our observations on small-time scaling behavior (made using the 15-minute long traces collected at 14:00 hours). As with other MAWI traces, these 24-hour long traces are partitioned into 15-minute long segments. Thus, we have a total of 96 trace segments for each 24-hour trace. We analyzed the small-time scaling behavior for each of these 96, 15-minute, segments.




Fig. 5. Fraction of aggregate bytes contributed by dense flows.

Fig. 4 compares the median small-time scaling parameter values obtained by application of the sketch-based robust estimation procedure. The results obtained are found to be consistent with observations outlined earlier in Section 4. We find all median values of the scaling parameter to be between 0.5 and 0.75. Moreover, the median values of scaling parameter for recent 2009 trace are typically larger than those for the older 2003 trace. For the recent 2009 trace, the values of median scaling parameter are found to be more than 0.6 in both directions. The median scaling parameter values for the older 2003 trace are generally less than 0.6. This provides further evidence that recent traces are more likely to display small to moderate correlations in small-time scales as compared to older traces.

5 Longitudinal Analysis of Dense Flows

This section explains some of the observations in the preceding section in terms of dense flows. Zhang et al. [16] defined a flow as *dense* if at least 50% of the packet interarrival times of the flow are below a certain threshold T ; otherwise, the flow is called *sparse*. Intuitively, dense flows have bursts of densely clustered packet arrivals. Zhang et al. [16] studied the aggregate of dense and sparse flows which are referred to as the dense and sparse components, respectively. Their study showed that the sparse component has a smaller scaling parameter in small-time scales ($h \leq 0.6$) relative to that of the original trace. Their work further illustrated that flow size alone does not play a significant role in shaping the small-time scaling behavior of network traffic. In this section, we revisit this property in the context of recent traces.

Fig. 5 shows, for the traces considered earlier in Section 4.2, the fraction of traffic carried by dense flows (defined using threshold $T = 2$ ms). Overall, we find that the contribution of dense flows to traffic volume has increased in the recent MAWI traces as compared to earlier traces. This increase in the fraction of aggregate traffic owing to dense flows is one factor that is likely responsible for the increase in the scaling parameter values we noted earlier in Section 4. A detailed analysis of the influence of dense flows follows using the “semi-experiments” based approach developed by Zhang et al. [16].




Fig. 6. Scaling behavior of the sparse component.

The process of artificially modifying the packet arrival process is referred to as a semi-experiment in the networking community. A comparison of scaling behavior before and after the semi-experiment leads to conclusions about the importance of the role played by the parameters modified by the semi-experiment. Specifically, for each trace considered, we compute the small-time scaling behavior for the sparse and small components of the trace. The small-time scaling parameter of the sparse component of a trace is computed by analyzing each trace sans the dense flows. Specifically, we removed all dense flows (using threshold $T = 2 \text{ ms}$) from the trace being analyzed and computed the scaling parameter h . To compute the scaling parameter for the small component of a trace, we remove from the trace the largest k flows (in terms of bytes) that contribute as many bytes as the dense flows. The scaling component is estimated using the robust estimation technique discussed earlier in order to weed out the impact of network anomalies.

Fig. 6 shows the median scaling parameter of the original traces and their sparse components. For all the traces exhibiting small to moderate correlations, the removal of dense flows reduces the scaling parameter. The scaling parameter for the sparse component is consistently close to 0.55 for all the traces. For the 2004-2006 JptoUS traces, the aggregate scaling parameter is close to 0.55, and hence no further reduction in scaling parameter is observed for the sparse component. For the 2001-2004 JptoUS traces, the scaling parameter is close to 0.6, and thus only a small decrease in the scaling parameter is observed for the sparse component. The 2006-2009 traces contain small to moderate correlations and in these a clear decrease in the scaling parameter is observed.

Fig. 7 compares the scaling parameters of the original traces and their corresponding small components. We find that, except for a few traces, the scaling parameters for the small component are close to the original traces even though




Fig. 7. Scaling behavior of the small component.

the small component is created by removing large flows that together contribute as many bytes as the dense component with $T=2\text{ms}$. Overall, our analysis reinforces the observation by Zhang et al. [16] that dense flows are one key factor to consider when studying small-time scaling behavior of network traces.

6 Conclusions

This paper presented a longitudinal study of small-time scaling behavior of Internet traffic, using traces spanning a period of eight years obtained from the MAWI trace repository. We believe this study has served multiple purposes. First, this study has re-emphasized the need for robust analysis techniques when studying scaling behavior of network traffic. Second, our study complements previous works on small-time scaling behavior [7, 16]. Specifically, we showed that small-time scaling behavior and many associated properties identified in the literature have more or less remained invariant throughout this decade. Our work is based on traces collected at one trans-pacific link. A longitudinal analysis of traces from other links and a subsequent comparison with our results will aid in understanding the evolution of small-time scaling behavior of Internet traffic.

7 Acknowledgements

We thank Brighten Godfrey, our shepherd, for his detailed feedback and suggestions. The authors are also grateful to the anonymous reviewers for their constructive suggestions, which helped improve the presentation of the paper. Financial support for this research was provided by NICTA. NICTA is funded by the Australian Government as represented by the Department of Broadband,

Communications and the Digital Economy and the Australian Research Council through the ICT Center of Excellence program.

References

1. MAWI working group traffic archive. <http://tracer.csl.sony.co.jp/mawi>.
2. N. Basher, A. Mahanti, A. Mahanti, C. Williamson, and M. Arlitt. A comparative analysis of Web and Peer-to-Peer traffic. In *WWW*, 2008.
3. P. Borgnat, G. Dawaele, K. Fukuda, P. Abry, and K. Cho. Seven years and one day: Sketching the evolution of Internet traffic. In *IEEE INFOCOM*, 2009.
4. G. Dawaele, K. Fukuda, P. Borgnat, P. Abry, and K. Cho. Extracting hidden anomalies using sketch and non-gaussian multiresolution statistical detection procedure. In *SIGCOMM LSAD*, 2007.
5. A. Feldmann, A. C. Gilbert, and W. Willinger. Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic. In *SIGCOMM*, 1998.
6. H. Gupta, A. Mahanti, and V. Ribeiro. Revisiting coexistence of Poissonity and Self-similarity in Internet traffic. In *MASCOTS*, 2009.
7. N. Hohn, D. Veitch, and P. Abry. Does fractal scaling at the IP level depend on TCP flow arrival processes? In *IMW*, 2002.
8. H. Jiang and C. Dovrolis. Source-level IP packet bursts: Causes and effects. In *IMC*, 2003.
9. H. Jiang and C. Dovrolis. Why is the Internet traffic bursty in short time scales? In *SIGMETRICS*, 2005.
10. T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A non-stationary Poisson view of Internet traffic. In *IEEE INFOCOM*, 2004.
11. W. E. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self similar nature of Ethernet traffic. *IEEE/ACM Transaction on Networking*, 2(1), 1994.
12. V. Paxson and S. Floyd. Wide area traffic: the failure of Poisson modelling. *IEEE/ACM Transaction on Networking*, 3(3), 1995.
13. D. Veitch. Code for estimation of scaling exponents, <http://www.cubinlab.ee.mu.oz.au/darryl>.
14. D. Veitch and P. Abry. A Wavelet-based joint estimator of the parameters of long-range dependence. *IEEE Transactions on Information Theory*, 45(3), 1999.
15. D. Veitch and P. Abry. A statistical test for time constancy of scaling exponents. *IEEE Transaction on Signal Processing*, 49(10), 2001.
16. Z. L. Zhang, V. J. Ribeiro, S. Moon, and C. Diot. Small-time scaling behaviors of Internet backbone traffic: An empirical study. In *IEEE INFOCOM*, 2003.