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

Himanshu Gupta[1,2], Vinay J. Ribeiro[2], and Anirban Mahanti[3]

[1] IBM Research Laboratory, New Delhi, India
higupta8@in.ibm.com
[2] Indian Institute of Technology, New Delhi, India
vinay@cse.iitd.ac.in
[3] NICTA, Alexandria, NSW, Australia
anirban.mahanti@nicta.com.au

Abstract. We carry out a longitudinal study of evolution of small-time scaling behavior of Internet traffic on the MAWI dataset spanning 8 years. MAWI dataset contains a number of anomalies which interfere with the correct identification of scaling behavior, and hence to mitigate these effects, we use a sketch-based procedure for robust estimation of scaling exponent. We first show the importance of robust estimation procedure while studying small-time scaling behavior of Internet traffic. We further study the evolution of the following properties concerning the origins of small-time scaling behavior: (1) Scaling at IP level is independent of flow arrivals and (2) Dense flows are primary correlation-causing factor in small time scales. Traditionally these properties have been shown to hold by using a semi-experiments based methodology. We next show that due to network anomalies, semi-experiments can result in misleading inferences. Hence we propose and motivate the use of "robust semi-experiments" i.e., a semi-experiment coupled with the use of a robust estimation procedure for inferring scaling behavior. By making use of robust semi-experiments we find the above properties to be invariant across the entire MAWI dataset. Our other results consist in showing that dense flows form a larger fraction of aggregate traffic for recent traces and hence recent traces show larger short range correlations vis-a-vis earlier traces.

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

1 Introduction

Scaling behavior of Internet traffic has been the focus of much networking research. It is well documented that Internet traffic displays two scaling regimes with transition point lying in 100ms - 1s time range [8, 3]. Internet traffic when aggregated to large time scales (≥1s) is quite bursty and is modeled using *long-range dependent* (LRD) processes [13, 15]. Scaling parameter (H) in large time scales lies within the range (0.8,1) which represents highly correlated packet arrival process. Few studies predicted that LRD may disappear as Internet traffic evolves and backbone links become highly loaded [12, 4]. However some studies [3, 7] showed that recent traces from highly loaded links do clearly exhibit LRD and that LRD is indeed an invariant.

The focus of this paper is on small-time scaling behavior of Internet traffic (≤100ms). Many studies have looked at the nature and mechanisms of small-time scaling behavior of Internet traffic [9, 8, 11, 10, 17]. Unlike in large time scales, Internet traffic displays tiny (h≈0.55) to moderate (h≈0.7) correlations[4] in small time scales [17]. Hohn et al. [8, 9], by using a *semi-experiments*[5] based methodology, showed that scaling at the IP level is independent of flow arrivals and hence small-time scaling structure has its origins in packet patterns within individual flows. Zhang et al. [17] further showed that packet patterns within *dense* flows, flows with bursts of densely clustered packets, are the primary correlation causing factors in small time scales. These two properties taken together explain the origins of small-time scaling behavior of Internet traffic.

Since the traces used in prior studies were collected (2000-02), a number of changes have taken place. Backbone capacities as well as Internet connected hosts have gone up. Composition of Internet traffic is significantly altered (Web-2.0 vs Web-1.0). This raises a number of interesting questions. As today's traffic consists of various modern applications e.g. YouTube, P2P file sharing, VoIP etc with widely different characteristic vis-a-vis traditional Web-1.0 traffic [2], does it still hold true that scaling at IP level is independent of flow arrivals? Do traces still display tiny to moderate short range correlations? Are recent flows more dense? Is the amount of traffic carried out by dense flows increasing or decreasing? Does it impact small-time scaling exponent? Answers to these question are not immediately apparent as with increasing backbone speed flows will appear sparser and the traffic uncorrelated. However with wide deployment of broadband access, large files will be transmitted faster with more correlated bursts, thereby making the flows more dense [17]. A clear understanding of small-time scaling behavior is critical to various network engineering problems e.g., router buffer dimensioning, delay-sensitive service provisioning etc [14, 6].

In light of these questions, this paper conducts a longitudinal analysis of small-time scaling behavior and properties on MAWI dataset [1] spanning 8 years (2001-2009). This dataset is known to contain a number of anomalies [5] which interfere with the reliable computation of scaling exponent and hence pose a problem in disentangling smooth long term evolutions from day-to-day fluctuations. To mitigate the effects of network anomalies, Borgnat et al. [3] developed a sketch-based procedure for robust estimation of scaling exponent and used this method to study the evolution of LRD behavior of Internet traffic across MAWI dataset. In this paper we use this method to study the evolution of small-time scaling behavior of Internet traffic.

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

[5] The approach of artificially modifying the packet arrival process for a trace is referred as a semi-experiment in the networking community.

The primary contribution of this paper is to present a case for application of a robust estimation procedure [3, 5] while carrying out a study of small-time scaling behavior of Internet traffic. Once the effects of network anomalies have been disentangled, we find that scaling parameter in small time scales consistently lies within 0.55-0.7 range, thereby showing the presence of tiny to moderate correlations in small time scales to be an invariant. Without the application of robust estimation procedure [3] we find many instances of traces either showing negative correlation (h <0.5) or large correlation (h >0.8).

We next present a case for coupling the robust estimation procedure with semi-experiments based methodology while studying scaling behavior of Internet traffic. By making use of such *robust semi-experiments* we show that the two properties (1) small-time scaling behavior is independent of flow arrivals and (2) dense (and not large flows) are primary correlation causing factors in small time scales, are invariant. If we do not couple robust estimation procedure with semi-experiments, we find many instances of traces which do not conform to these properties thereby giving misleading results. To the best of our knowledge, this is the first study which motivates the need of robust semi-experiments.

Our other results consist in showing the evolution of small-time scaling behavior of Internet traffic. We find that recent MAWI traces consistently show larger small time correlations ($h\approx0.7$) as compared to earlier traces. This is in contrast to the observations made in [17] that packet traces only occasionally show small correlation (h within 0.6 and 0.7). This finding also provides an evidence against the prediction that Internet traffic will likely be describable by simple models (e.g. Poisson) [12]. We further show that dense flows in recent years are carrying a larger fraction of aggregate traffic vis-a-vis earlier years.

The rest of the paper is organized as followed. Section 2 summarizes MAWI dataset and sketch-based procedure for robust estimation of scaling parameter. Section 3 studies small-time scaling behavior across the years and shows the need of robust analysis for the same. Section 4 motivates the concept of robust semi-experiments and shows that the property of IP level scaling being independent of flow arrivals is invariant. Section 5 shows that dense flows are consistently driving small-time scaling behavior across the years and further highlights the importance of robust semi-experiments. Section 6 concludes the paper.

2 MAWI dataset and Robust Estimation Procedure

MAWI Dataset: We use publicly available traces collected from WIDE, a trans-Pacific backbone [1]. A detailed statistical characterization of this dataset is provided by Borgnat et al. [3]. For each day a 15 minute extract is made public for download. We use data collected across samplepoints B and F. Samplepoint-B was a 100 Mbps link and was replaced in July 2006 by Samplepoint-F, a 150 Mbps link. For our study we select the trace collected on 1^{st} and 15^{th} of every month, from Jan 2001 to Dec 2008. Days for which traces are unavailable are left out. This gives us a set of 180 traces spanning 8 years. Each trace is partitioned into two subt-races, one for traffic flowing in each direction, labeled UStoJp and JptoUS. We analyze each sub-trace separately.

As mentioned earlier, nearly all traces contain some sort of anomaly while some anomalies are severe and last weeks or months(e.g., Sasser worm for 2004/07 to 2005/04; UStoJp, Ping flood for 2003/08 to 2003/12 both directions, severe volume decrease for 2003/05 to 2004/03; JptoUS, flooding attacks in 2001 etc) [5, 3]. The dataset displays wide range of throughput values, a global increase of throughput from 100 kbps in 2001 to more than 12 Mbps in 2008. There are several long lasting congestions (e.g. 2005/09 to 2006/06; JptoUS). To mitigate the effects of anomalies, Borgnat et al. [3] proposed a method for robust estimation of scaling parameter (also called Hurst parameter), described next.

Robust Estimation of Hurst Parameter: We use the Wavelet method [16] to estimate the value of Hurst parameter. This method plots the logarithm of variance of the coefficients obtained after taking a discrete wavelet transform of the process against scale j, the plot known as logscale diagram (LD). The slope of this plot α , gives an estimate of Hurst parameter ($H = \frac{1+\alpha}{2}$). A hurst value of 0.5 implies uncorrelated scaling. In this paper we interpret Hurst values of 0.55, 0.6 and 0.7 to represent tiny, small and moderate correlations respectively.

Robust estimation enables long term analysis without being affected by specific traffic conditions or anomalies. Let f_n denote a hash table of size M. Original collection of packets is split into M sub-collections, each of them consisting of all packets with identical sketch output $m = f_n(A)$ where the hashing key A is chosen as one of the packet attributes (IPdst,IPsrc,...). This amounts to performing random projections, preserving flow structures as packets belonging to a given flow are assigned to same sub-collection. Each sub-collection is aggregated and Hurst parameter computed. Robust estimation of Hurst parameter results by taking median over the values of Hurst parameter estimated by using individual sketch outputs [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. Resulting LDs have the same shape as original, with a variance which is appropriately scaled down[6], consistent with an independent and identically distributed (i.i.d.) superposition model [8]. In presence of anomalies, sketching the original packet stream reduces their impact, possibly mapping them to few bins only. Small time correlation structure of the traffic in the bins containing anomalies differs from normal traffic as well as traffic in other bins. Median over indepedent sketches achieves the robustness. $Median$ is chosen instead of *mean* as median is a non linear procedure providing robustness against outliers. Robust estimator can still be fooled if anomalies are dominant part of the traffic. Robustness in such cases can then be achieved by maintaining multiple sketches and taking median over estimates computed from them. For a detailed description of the method we refer the reader to [3].

[6] if one selects flows with probability 0.7, the resulting LD will have a variance approximately 70% of the original

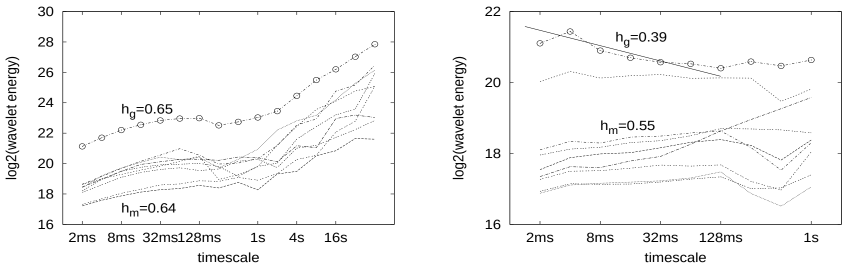

Fig. 1. Robust Estimation (a) Anomaly Free Trace (b) Anomalous Trace

3 Evolution of short range correlations

This section first illustrates the importance of robust estimation of scaling parameter (to remove the effects of anomalies) while studying small time correlations of Internet traffic. Then it shows that, once the effects of network anomalies have been disentangled, MAWI traces consistently display tiny to moderate small time correlations with scaling parameter lying in range (0.5-0.7).

Figure 1(a) shows logscale diagram (upper plot) for a MAWI trace (July 11, 2005; UStoJp) which is free of any anomaly [3]. Stationarity for 15 minutes is well established [3, 7] and hence wavelet estimator can be applied. We construct a time series by counting the number of bytes every millisecond and then use this time series to do scaling analysis using wavelet method. Figure 1(a) plot shows the representative scaling behavior of Internet traffic [9]. Biscaling behavior is clearly evident with the *knee point* (change of scaling behavior) falling within (100ms-1s) range. This trace has small correlations in small time scales ($h = 0.65$). As the focus of this paper is on small-time scaling behavior of Internet traffic, for all experiments we run our analysis only on 5 minutes long extract which is enough to give good estimates of small-time scaling behavior (1-100ms). Also all occurrences of scaling parameter and logscale diagram from hereon mean scaling parameter in small time scales and logscale diagram for small time scales respectively unless specified otherwise.

Figure 1(b) shows logscale diagram for a trace (Oct 11, 2005; JptoUS) which deviates from normal behavior. This trace contains a low-intensive long-lasting spoofed flooding anomaly [5]. Scaling exponent is found to be 0.39 which indicates the presence of small negative correlation (also called *anti-persistent* behavior) in the trace. This is in direct contrast to the fact that Internet traffic consists of tiny to moderate positive correlation in small time scales ($0.5<h<0.7$) [17].

Next we estimate the scaling exponent using sketch based robust estimator [3]. We hash the trace (destination IP as hash-key) into 8 bins and estimate the scaling exponent for each sub-trace. Figure 1(b) plots all sub-traces LDs for anomalous trace. Except for two sub-traces (possibly containing anomaly), all other sub-trace LDs have recovered normal behavior and each of these sub-traces now displays tiny positive correlation. As a result h_m, *median of estimates* over 8 sketches, is found to be 0.55 while the *global estimate* h_g was 0.39. This shows that h_g being 0.39 is an artifact of network anomalies; otherwise all sketched LDs should have had a similar value of scaling exponent. For anomaly free trace, all

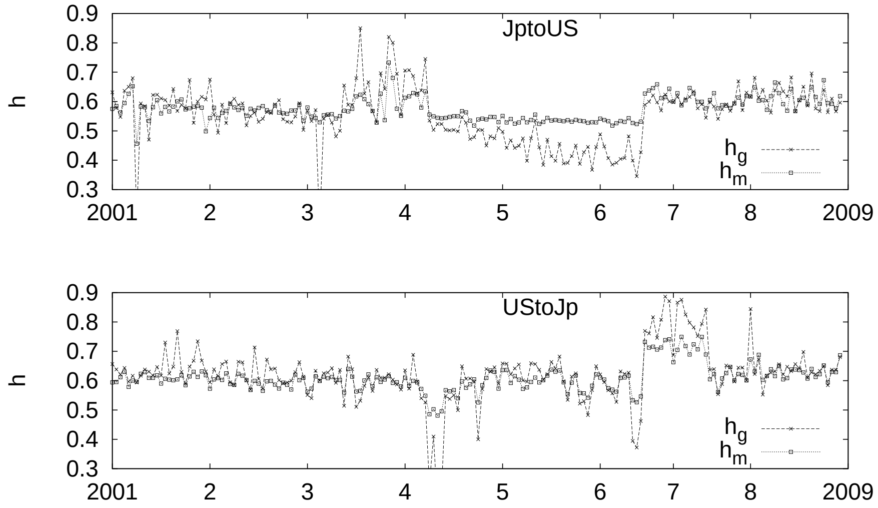

Fig. 2. Short range correlations across the years, both directions

sketched LDs (Figure 1(a)) are found to be parallel to original LD as expected and hence median value h_m matches with h_g computed from whole trace.

Figure 2 plots scaling exponent for MAWI traces across 8 years, with and without robust estimation. There are multiple traces for which h_g is found to be less than 0.5 signifying negative correlations. Specifically, from 2005 to mid-2006; JptoUS, value of h_g is consistently less than 0.5, often close to 0.4, suggesting small negative short range correlations. However the median values of scaling parameter h_m, computed by using robust estimation procedure, are markedly different with h_m consistently lying close to 0.55. Similarly for many traces 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 common knowledge. However the values of h_m revert back to their usual behavior i.e. tiny to moderate correlations. This illustrates the importance of using a robust estimation procedure while studying short range correlations and in absence of such a procedure one may draw faulty conclusions e.g. the presence of negative correlation in small time scales.

This median-sketch based longitudinal analysis of MAWI traces shows that the presence of tiny to moderate short range correlations, with corresponding scaling parameter lying in the range (0.5-0.7), is an invariant. Despite evolution of Internet traffic, presence of congestions and anomalies, variation in bandwidth occupancy rate etc, small-time scaling behavior is found to be stable.

We note that many traces, in both directions, manifest small to moderate correlations (h between 0.6 and 0.7). This is in contrast with the findings that small to moderate correlations are found only in small number of traces [17]. A close look reveals that this phenomenon is more frequent for recent traces (2007 onwards). Most of the pre-2007 traces have scaling exponent less than 0.6 and 0.65 for directions JptoUS and UStoJp respectively while many post-2007 traces cross these bounds. This suggests that the trend is towards increasing

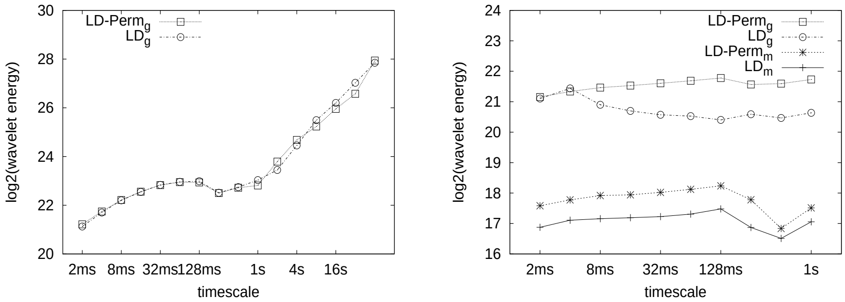

Fig. 3. Logscale Diagrams before and after permuting the flows (a) Anomaly Free Trace (b) Anomalous Trace

short range correlations, although the scaling exponent remains less than 0.7 . Section 5 provides an explanation for recent traces having larger correlations (proliferation of dense flows). However this observation suggests that sub-second Poisson modeling is unsuitable for recent traces and provides an evidence against the prediction that Internet traffic is moving towards simpler to describe models (e.g., Poisson) [12].

4 Independence of scaling at IP level from flow arrivals

This section illustrates the importance of coupling a semi-experimental methodology with a robust estimation procedure. A semi-experiment based methodology has been frequently used to infer various properties regarding scaling behavior of Internet traffic [8, 17, 10]. We argue that in presence of network anomalies, it is important to couple semi-experiments with a robust estimation procedure otherwise semi-experiments may give misleading inferences. By making use of such *robust semi-experiments* we show that the property, scaling at IP level is independent of flow arrivals, is found to be invariant. Hohn. et al. [8] showed this independence while trying to unravel the origins of scaling in small time scales. This property is important as it suggests that for the purpose of modeling the overall process of IP packets, flows can be treated as statistically independent and hence forms the basis of cluster process models [9].

For our analysis we make use of a similar semi-experiments based methodology as used by [8, 9]. We modify the arrival process of flows while maintaining in full the packet arrival patterns within each flow. Specifically, we permute the flows around the original arrival points and then compare the scaling structure before and after the semi-experiment. Any flow which lasts longer than the trace finish-time is wrapped around. If the scaling behavior remains identical, it shows the independence of IP level scaling from flow arrivals process.

Figure 3(a) plots LDs of anomaly-free trace before and after the permuting semi-experiment. As expected, scaling behavior is found to be identical, both in small and large scales. Figure 3(b) shows LDs for anomalous trace which displays different small-time scaling behaviors before and after the permuting semi-experiment (upper two plots)[7]. While the original LD has small negative

[7] Figure 1 and Figure 3 use the same traces

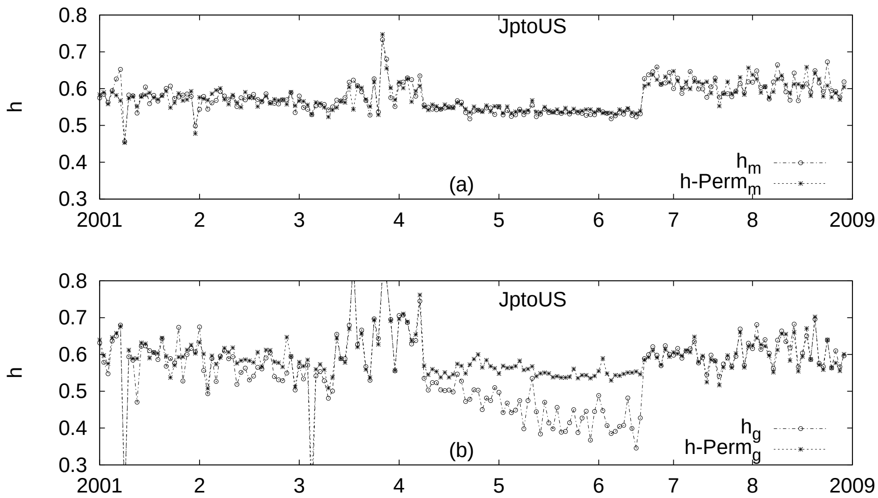

Fig. 4. Short range correlations across the years before and after flow permutations, with and without robust estimation

correlation ($h_g = 0.39$), LD after flow permutation shows tiny positive correlation ($h\text{-}Perm_g = 0.54$). At first glance it suggests that this violates the independence property. However as pointed out in previous section, network anomalies may interfere in identifying the correct scaling behavior, we need to do a robust estimation. Hence for both original and permuted trace, we estimate the scaling exponent using robust estimation procedure and compare the median-LDs. Figure 3(b) plots the median-LDs for both original and permuted trace (lower two plots). Scaling behavior for both median-LDs is found to be identical thereby showing the independence of scaling from flow arrival process for this trace. This shows the importance of using a robust semi-experimental methodology. One can alternatively think of first sketching the original trace, permuting each sketched sub-trace and finally checking whether this independence property holds for majority of sketched sub-traces. Median can then be taken over matching sketched sub-traces. However, the former approach is more strict as the robust value should not depend on the actual hash-mappings taking place during the sketching procedure.

Figure 4 plots the values of scaling exponent for MAWI dataset for direction JptoUS. Both median (a) and global (b) are plotted before and after permuting the flows. Figure 4(a) shows that median values of scaling exponent match quite nicely before and after permuting the flows, thereby showing invariance of this property. Similar results are obtained for direction UStoJp as well.

A comparison of Figure 4(a) and 4(b) again throws light on the importance of using a robust semi-experimentation based methodology. For duration 2005-mid 2006 global estimations of scaling exponent, h_g and $h\text{-}Perm_g$, do not match. For direction UStoJp same observation is made for mid 2001-2002 and 2007 traces (plot not shown here). In absence of a robust estimation procedure one may draw a misleading conclusion. On the other hand for few traces (notably around 2004, JptoUS) global estimations of scaling exponent (h_g and $h\text{-}Perm_g$) are close to 0.8 and are found to be matching before and after permuting semi-experiment. For such traces one will draw the correct conclusion of independence of IP level scaling from flow arrivals but the exact nature of scaling behavior will be misinterpreted (large correlations instead of moderate correlations). A true picture is obtained only after applying the robust estimation procedure.

5 Small-time scaling behavior and Dense flows

Hohn et al.'s [8] result regarding independence of scaling behavior at IP level from flow arrival process implies that dependence between packet processes across different flows are very weak and hence Internet traffic can be considered to be a collection of independent flows layed down in some independent manner. This further suggests that small-time scaling behavior arises out from packet patterns within individual flows. Zhang et al. [17] showed by analyzing backbone traces collected in 2001-02 that small-time scaling behavior is driven by packet patterns within individual dense flows. A flow is defined as dense if 50% of packet interarrival times are less than a threshold T. Moreover large flows do not have much say in small-time scaling behavior of Internet traffic. This was a surprising result as large flows are known to be reason behind LRD. Main objective of this section is to look whether the property that dense flows, and not large flows, are the primary correlation causing factors in small time scales is an invariant or not. In doing so, this section reinforces the importance of coupling robust estimation method with semi-experiments. We further study the evolution of MAWI traces vis-a-vis dense flows.

For our experiments on MAWI traces we carry out a similar semi-experiments based analysis as used by [17]. For a given trace we extract out all dense flows with threshold $T = 2ms$ and remove these dense flows from the trace. This hence leaves us with *sparse* component of the original trace. Next we construct the *small* component as followed: We compute the number of bytes contributed by the dense flows *dense-bytes* and then remove *top-k* flows (in terms of size) which taken together contribute as many bytes as dense flows do i.e. *dense-bytes*. We then compare the scaling behavior of sparse and small components vis-a-vis the original trace in small time scales. Once again we apply the robust estimation procedure to weed out the effects of network anomalies.

Figure 5 plots the scaling parameter of sparse and small components as well as aggregate for MAWI traces across the years. Both directions are shown. Scaling parameters only after application of robust estimation procedure are shown. For all the traces having small to moderate correlations, removal of dense flows brings down the scaling parameter. Scaling parameter for sparse component is consistently close to 0.55 for all the traces (both directions). For traces JptoUS; 2004-2006 aggregate scaling parameter is close to 0.55 (i.e. almost uncorrelated) and hence no further reduction in scaling parameter is observed for sparse component. For traces 2001-04; Jp2US scaling parameter is close to 0.6 and hence only a small decrease in scaling parameter is observed for sparse component. However traces for UStoJp and JptoUS; 2006-09 contain small to moderate correlations and a clear decrease in scaling parameter is observed.

It can be further observed that removing large flows from the traces does not have much of an effect on scaling exponent. Scaling exponent (across all traces) of small component is close to aggregate trace as compared to sparse component even though both sparse and small components have been obtained by removing same number of bytes. We tried out multiple other values of threshold T and a similar result is observed every time[8]. This analysis along with the earlier observation of sparse component consistently displaying almost uncorrelated scaling, shows that the property, small-time scaling behavior is driven by dense flows (and not large flows), is an invariant.

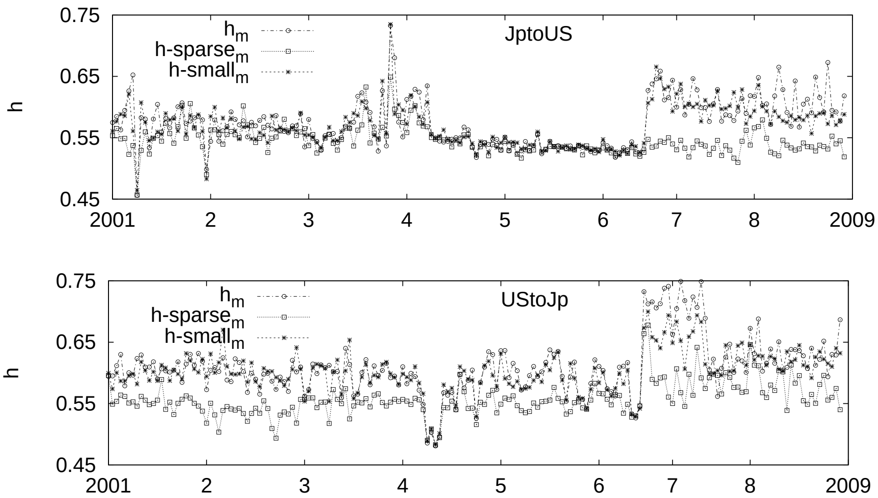

Fig. 5. Comparison of small-time scaling behavior across the years: sparse and small components vis-a-vis aggregate

We next look at global values of scaling parameter (without robust estimation) for sparse and small component. We find that sparse component consistently has a smaller global scaling exponent ($h\text{-}sparse_g$) as compared to small ($h\text{-}small_g$) and aggregate (h_g) component (plot not shown). However we find anomalies clearly interfere with the reliable estimation of scaling parameter. Figure 6 plots the global estimations of scaling parameter for sparse components. Once again we find multiple instances of traces showing negative correlations clearly indicating the effect of anomalies. Moreover many traces display a value

[8] We also tried computing small component by removing all flows with size greater than 1 MB. Scaling exponent for sparse component is again found to be less than that of small component.

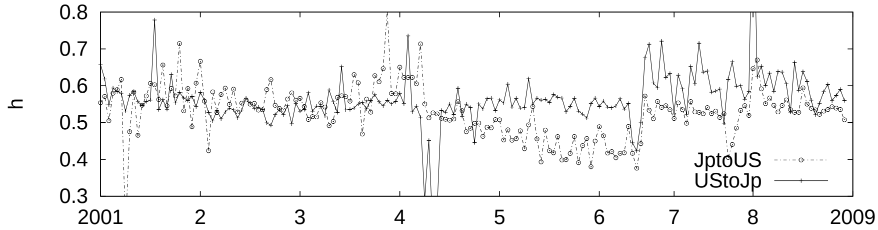

Fig. 6. Global values of scaling parameter for sparse component ($h\text{-}sparse_g$)

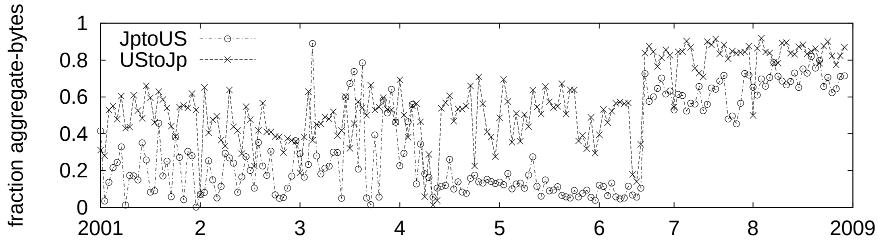

Fig. 7. Evolution of Internet Traffic across the years vis-a-vis dense flows. y axis represents fraction of traffic carried by dense flows with T=2ms

close to 0.7 for sparse component which is counter-intuitive as dense flows have been removed and hence this raises the question what is causing moderate correlations in sparse component for these traces. This again shows that a correct and consistent picture is obtained only by carrying out robust semi-experiments, thereby underlining their importance.

Next we study the evolution of Internet traffic vis-a-vis dense flows. Figure 7 plots the fraction of traffic carried by dense flows (with T=2ms) by MAWI traces across the years. We find that for both directions fraction of aggregate traffic carried by dense flows has increased. This also explains the earlier observation that scaling exponent for recent traces is found to be relatively larger as compared to traces from earlier years. Further as direction UStoJp carries more traffic by dense flows, scaling exponent for UStoJp direction is found to be relatively higher as well. Secondly we notice that for traces UStoJp; 2001-2006, dense flows almost carry 60% of traffic. Scaling exponent of small component, obtained by removing 60% (a significant number) of total bytes from large flows, is found to be close to that of aggregate trace. This observation reaffirms the fact that small-time scaling behavior has its origins in packet patterns within dense flows.

6 Conclusions

This paper carries out a unique longitudinal analysis of small-time scaling behavior of Internet traffic on MAWI dataset spanning traces across 8 years. We think that this study has served multiple purposes. First this study has re-emphasized the need of a robust analysis in general and while studying small-time scaling behavior in particular. We have also motivated the coupling of Robust analysis with semi-experiments by showing how a semi-experiment without robust analysis can give misleading inferences.

Secondly our study is complimentary to many previous works analyzing small-time scaling behavior (e.g. [17], [8]). We have shown that small-time scaling behavior and properties proposed by these studies remain invariant throughout this decade despite many things having changed e.g. Internet traffic composition, bandwidth usage etc. Third our study suggests some trends regarding the evolution of small-time scaling behavior. Our study suggests that the percentage of traffic carried out by dense flows is increasing thereby pushing scaling parameter in small time scales upwards. As a result recent traces frequently display small to moderate short range correlations as compared to earlier years. This also provides an evidence against the prediction that Internet traffic is moving towards simpler to describe models (e.g. Poisson). However it is prudent to note here that these trends should be verified by a longitudinal analysis from traces collected on other backbone links as well.

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 *Proceedings of 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 *Proceedings of INFOCOM*, 2009.
4. J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. On the nonstationarity of Internet traffic. In *Proceedings of SIGMETRICS/Performance*, 2001.
5. G. Dawaele, K. Fukuda, P. Borgnat, P. Abry, and K. Cho. Extracting hidden anomalies using sketch and non-gaussian multiresolution statistical detection procedure. In *Proceedings of SIGCOMM*, 2007.
6. C. Fraleigh, F. Tobagi, and C. Diot. Provisioning IP backbone networks to support delay-based service level aggrements. In *Proceedings of INFOCOM*, 2003.
7. H. Gupta, A. Mahanti, and V. Ribeiro. Revisiting coexistence of Poissonity and Self-similarity in Internet traffic. In *Proceedings of MASCOTS*, 2009.
8. N. Hohn, D. Veitch, and P. Abry. Does fractal scaling at the IP level depend on tcp flow arrival processes? In *Proceedings of IMW*, 2002.
9. N. Hohn, D. Veitch, and P. Abry. Cluster processes: A natural language for network traffic. *IEEE Transactions on Signal Processing*, 51(8):2229–2244, 2003.
10. H. Jiang and C. Dovrolis. Source-level IP packet bursts: Causes and effects. In *Proceedings of IMC*, 2003.
11. H. Jiang and C. Dovrolis. Why is the Internet traffic bursty in short time scales? In *Proceedings of SIGMETRICS*, 2005.
12. T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A non-stationary Poisson view of Internet traffic. In *Proceedings of INFOCOM*, 2004.
13. 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):1–15, 1994.
14. A. L. Neidhardt and J. L. Wang. The concept of relevant time scales and its application to queuing analysis of self-similar traffic. In *Proceedings of SIGMETRICS/Performance*, 1998.
15. V. Paxson and S. Floyd. Wide area traffic: the failure of Poisson modelling. *IEEE/ACM Transaction on Networking*, 3(3):226–244, 1995.
16. D. Veitch and P. Abry. A statistical test for time constancy of scaling exponents. *IEEE Transaction on Signal Processing*, 49(10):2325–2334, 2001.
17. Z. L. Zhang, V. J. Ribeiro, S. Moon, and C. Diot. Small-time scaling behaviors of Internet backbone traffic: An empirical study. In *Proceedings of INFOCOM*, 2003.