Understanding the Long-Term Self-Similarity of Internet Traffic*

Steve Uhlig and Olivier Bonaventure

Infonet group, University of Namur, Belgium
http://www.infonet.fundp.ac.be
{suhlig,obonaventure}@info.fundp.ac.be

Abstract. This paper analyzes the characteristics of Internet traffic by studying a six days long trace of the entire interdomain traffic received by an ISP. Our study shows that this traffic is self-similar at time-scales spanning minutes to hours. We show that this self-similarity could be explained by two factors. First, the traffic volume received from each external source exhibits a heavy-tailed distribution. Second, the number of these external sources is also self-similar. Finally, we show that self-similar traffic can be simulated by users transferring exponentially distributed traffic provided that the number of users is self-similar.

1 Introduction

It is now a long time since self-similarity has been uncovered in data networks. From Ethernet traffic [LTW+94] to wide-area traffic [PF95], passing through web traffic [CB96], all have been characterized by statistical self-similarity. Some have modeled quite successfully such traffic with ON/OFF traffic sources [WPT98] [TWS97] while others have tried to fit a Markov-modulated model [RL97]. Whether or not one can model such high variability or partially explain its causes, the main fact remains the corresponding burstiness and its related problems, from packet handling mechanisms to link capacity planning. While telephone networks have successfully relied on statistical multiplexing for reducing operational costs, the Internet has this unfortunate feature that limits traffic aggregation from reducing traffic variability, also called traffic self-similarity. [PKC97] has discussed some implications of self-similarity on network performance as well as the impact of network handling mechanisms on traffic characteristics. While packet handling mechanisms may in part be taken as responsible for short-range traffic self-similarity, it is unlikely that such short-term aspects like packet-level dynamics can explain large time-scale variability as shown in [CB96]. In that respect, the purpose of this paper is to better understand the long-term traffic self-similarity, in the order of minutes and hours.

More precisely, we study the traffic self-similarity by partitioning network traffic dynamics into its probabilistic and dynamic aspects. An important idea has already been raised in [CB96] and [PKC96], in that the

* This work was partially supported by the European Commission within the IST ATRIUM project.

heavy-tailed characteristics of the objects' sizes to be transferred over the network could suffice for generating self-similarity. It has also been shown that heavy-tailed file transfer duration and file sizes can lead to high variability.
We show in this paper that there are two very different parts of the network traffic generation process that have different implications on self-similarity. On one side, a heavy-tailed objects sizes distribution suffices for network traffic self-similarity to arise. On the other side, the dynamics of the number of IP sources sending traffic during a particular time interval allows for the distributional properties to actually generate a self-similar traffic pattern.
The remainder of this paper is structured as follows. Section 2 presents the context in which the traffic traces were gathered. We introduce self-similarity in section 3 by studying the evolution of the total traffic. We then look at traffic components in section 4 to find plausible causes for this self-similarity. We first look at the the heavy-tailed properties of the traffic trace in section 4.1 and explain its role in self-similarity. We go on by studying the evolution of the number of traffic sources in section 4.2 and try to assess in section 4.3 which part between heavy-tails and traffic dynamics is more likely to cause self-similarity.

2 Traffic Statistics

Many researchers have chosen to study the behavior of network traffic by relying on packet level traces from a particular link (see [LTW+94] [PF95] [MC00] among others). Such traces allow the analysis of the traffic at the packet scale, but require a large storage space. For this reason, most of the used traces either correspond to a low traffic volume or a short period of time. In this paper, we have taken a different approach. In order to better understand the long-term behavior of the traffic, we consider a less precise trace that spans six complete days. For this, we rely on a `Netflow` [Cis99] trace collected at the border routers of the Belgian research ISP Belnet during December 1999 [UB00]. The trace covers the interdomain traffic received by the ISP on all its access links and accounts for 2.1 Tbytes of traffic. The studied ISP provides access to the Internet as well as to high speed European research networks to universities, government and research institutions in Belgium. At that time, its national network was based on a 34 Mbps backbone linking major Belgian universities. Its users are mainly researchers or students with direct high speed connections to the 34 Mbps backbone, although some institutions also provide dial-up service to their users. It was also at that time the ISP with the largest capacity in Belgium.
This network is connected to a few tens of external networks with high bandwidth links. It maintains high bandwidth peerings with two transit ISPs, the Dutch SURFNET network and is part of the TEN-155 European research network, without providing any transit service to its peers. In addition, the ISP is present with high bandwidth links at the Belgian and Dutch national interconnection points with a total of about 40 peering agreements in operation. The `Netflow` trace we used aggregated the

information from all upstream links, in a manner that we have the incoming traffic information like if there was only one access link to the local ISP.

`Netflow` provides us with the aggregated information of the layer-4 flows, by recording the starting time, the ending time and the total volume in bytes for each unidirectional TCP and UDP flow. The utilization of `Netflow` forces us to approximate the layer-4 flows as equivalent to fluid flows. More precisely, a flow transmitting M bytes between T_{start} and T_{stop} is modeled as a fluid flow transmitting $M/(T_{stop} - T_{start})$ bytes every second between T_{start} and T_{stop}. This approximation obviously leads to an inaccurate estimation of the short-term burstiness of the traffic but allows for longer traffic traces collection. The time granularity of the trace is one minute, i.e. all traffic volume information is summarized over equally-spaced one minute intervals throughout the six days. The part of the `Netflow` statistics on which we rely throughout this paper is the information of the total traffic volume received from every external IP address for every minute of the measurements. We thus have for each minute the information of the number of IP sources that are sending traffic during a particular minute as well as the corresponding traffic volume for each of them. Even if the trace granularity is one-minute, it accounts for more than 42 million values recorded over more than 8600 samples, with an average of 4912 IP addresses sending traffic per minute. Hence it provides a very fine measurement of the traffic dynamics for time-scales spanning minutes to hours.

3 Total Traffic

The focus of this section is on measuring self-similarity of the total traffic time-series. We thus look at total traffic received at the incoming access links for every one-minute time interval during the 6 days of the measurements. Figure 1 shows the evolution of total traffic during the period of the measurements. While the global evolution of total traffic exhibits a stable daily periodicity, with peak hours located during the day, there are important deviations around the average traffic evolution throughout the day that give self-similarity this highly bursty look, over many time-scales. The mean traffic over the six days was slightly larger than 32 Mbps, with a one-minute maximum peak at 126 Mbps and a standard deviation of 21 Mbps. The trace begins around 1 AM and finishes six days later around 1 AM also.

Statistical *self-similarity* in the context of stationary discrete-time processes is defined through the following procedure. Let $X = \{X_i, i \geq 1\}$ be a stationary sequence, where

$$X^{(m)}(k) = \frac{1}{m} \sum_{i=(k-1)m+1}^{km} X_i, \quad k = 1, 2, \ldots \tag{1}$$

represents the *m-aggregated* sequence obtained by summing the original sequence X over non-over-lapping blocks of length m and averaging over

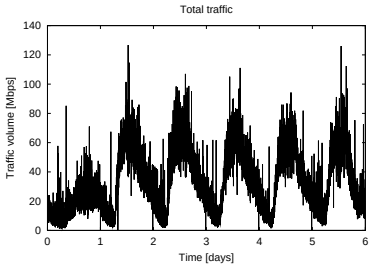

Fig. 1. Total traffic evolution.

each block. The sequence $X^{(m)} = \{X_k^{(m)} : k = 1, 2, ...\}$ is said to be *asymptotically self-similar* if

$$X \stackrel{d}{=} m^{1-H} \; X^{(m)} \quad as \quad m \to \infty \tag{2}$$

where $\stackrel{d}{=}$ denotes equivalence in distribution. The parameter H (for the Hurst parameter) indicates the degree of long-range dependence (self-similarity) in the time-series. The value of H for self-similar processes is between 0 and 1, with values under $1/2$ meaning short-range dependence while values over $1/2$ relate to long-range dependence. This paper relies on several estimators for measuring the parameter H, the main measure for self-similarity in time-series. Because H is difficult to measure in practice, we rely on several estimators (see [Ber94] [TT98] [TTW95]) to obtain a gross picture of its value range. We focus on qualitative self-similarity, meaning that the value of H is not overly important, but rather whether it is larger than $1/2$ or not. It must be clear however that self-similarity is an asymptotic concept, meaning that statistical inference is difficult on the sole basis of finite measurements.
The first estimator is the well-known R/S statistic [Ber94]. Plotting the R/S statistic for large k on a log-log scale must be scattered around a straight line of slope H, which is found to be around 1 for total traffic (upper left of figure 2).
The second method for measuring H is the *aggregated variance* method, where one computes the sample variances of the *m-aggregated* series around their sample means

$$s^2(m) = \frac{1}{m-1} \sum_{i=1}^{m} X^{(m)}(k) - \bar{X}^{(m)}. \tag{3}$$

Then plot $s^2(m)$ as a function of m on a log-log scale. This should follow for large values of m a straight line with negative slope $2H - 2$. This method gives a value of H very close to 1, as shown on figure 2 (upper right), with a slope around 0.
The third estimator for H relies on the time dependence between the samples, so that long-range dependent data should exhibit very slowly decaying sample correlations proportional to k^{2H-2} for $1/2 < H < 1$. Figure 2 (lower left) shows this very slow decay by plotting the sample

autocorrelations as a function of the lag k on a log-log scale. We can see that the decay rate is far slower than $k^{-0.5}$, thus the estimated H is close to 1.

The three previous methods were based on the time-domain, the final one is the periodogram, a frequency-domain based one. If the series exhibits *long-memory* then the estimated spectral density at the origin should behave like

$$f(\lambda) \sim c_f |\lambda|^{1-2H} \quad as \ \ |\lambda| \to 0. \tag{4}$$

Hence plotting the periodogram near the origin with a log-log scale should roughly follow a straight line of slope $1 - 2H$. This estimator gives a value between 0.75 and 1 in our case (lower right of figure 2).

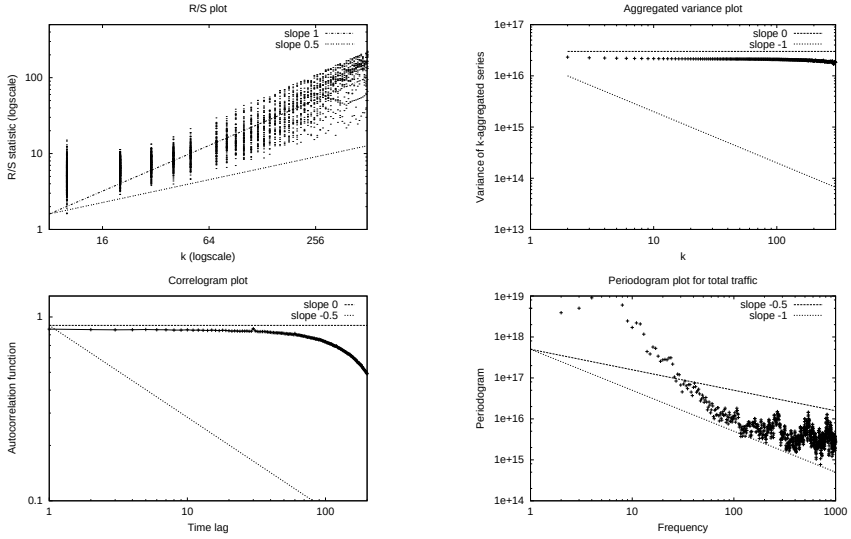

Fig. 2. Estimation of H for total traffic.

It must be noted that most of the estimators for H are defined in an asymptotic way and therefore provide a good estimate for very large datasets only. However, our time-series being quite long, with more than 8600 samples, we can be quite confident with the previous estimators at least in the value range they provided, with H closer to 1 than 1/2 for all of them, meaning that our time-series exhibits a strong self-similarity.

4 Understanding the Long-Term Self-Similarity

Section 3 has described total traffic variability by showing how close to 1 the estimated value of the estimators for H are. In this section, we analyze which parts of the dynamics of the network traffic are due to be responsible for this self-similarity.

4.1 The Role of Heavy-Tails

As already pointed out in [CB96], heavy-tails could play an important role in traffic self-similarity. Traffic variability can be thought as an aggregated random phenomenon generated by the dynamics of user's that are retrieving files over the network. Intuitively, it makes sense that whenever the probability of users transferring a very large amount of bytes does not decay fast enough for large values of the transferred object, this is due to generate large bursts that are persistent at many time-scales. Heavy-tailed distributions conceptualize this idea, in a probabilistic way. For that purpose, this section studies the probability mass of the amount of traffic that is seen for every IP source for every one-minute time interval of the trace.
A random variable X is said to have a *heavy-tailed* distribution if

$$P[X > x] \sim x^{-\alpha}, \ as \ x \to +\infty. \tag{5}$$

A random variable Y is said to be exponentially distributed if

$$P[Y > y] = \alpha e^{-\alpha y}, \ with \ \alpha > 0 \ and \ y > 0. \tag{6}$$

The exponential distribution is associated with *memoryless* processes, that do not depend on the past. Heavy-tailed distributions however characterize *long-memory* processes, with strong time-dependence structures that vanish very slowly. The main difference between heavy-tailed and non-heavy- tailed distributions, in the context of this paper, lies in the asymptotic decay rate of the tail that is not exponential in the case of heavy-tailed distributions [Res97]. This means that when fitting an empirical heavy-tailed distribution, the decay rate of the tail is far slower than exponential, namely like a power function. To illustrate the heavy-tailedness of network traffic, figure 3 shows the probability mass (on a log-log scale) of the amount of traffic seen during one-minute time intervals for individual IP addresses during the measurements. We also plotted two reference power tails, one with $\alpha = 1$ and the other with $\alpha = 2$. This shows that the small values are distributed like a heavy-tail with exponent within $[1, 2]$ while the tail of the distribution happens to be more flat.

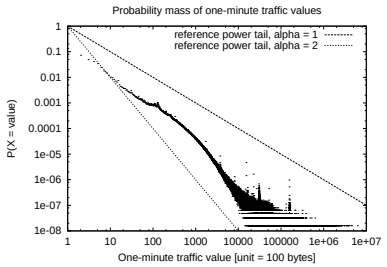

Fig. 3. Heavy-tails in one-minute traffic values.

Our intent in this paper is not to assess which particular heavy-tailed distribution fits best our data. Thus, we will only shortly discuss the estimation of the exponent α. Figure 4 presents an estimate of α that illustrates the rather strong heavy-tailed characteristics of our trace. Let $X_1 \geq X_2 \geq ... \geq X_n$ denote the ordered statistics of the data, X_1 being the largest value of the data, X_2 the second largest value, and so forth... The Hill estimator [Hil75] gives the estimation of α for the k largest values of the data set and is defined by

$$H_{k,n} = \frac{1}{k} \sum_{i=1}^{k} log(\frac{X_i}{X_{k+1}}) \tag{7}$$

As presented on figure 4, the Hill estimator gives an α well under 1. The thing to notice concerning the Hill plot is the estimate of α for the largest values of the dataset (low order statistics) that remains extremely low, rendering the very slow decay of the tail, hence the strong persistence of large values within traffic sent by each individual IP source during one-minute periods.

Fig. 4. Estimation of alpha.

4.2 The Role of Traffic Dynamics

Heavy-tails are a distributional property, in the sense that it provides an intuitive explanation of the high variability of total traffic, through the probabilistic persistence of high bursts among the traffic volume sent by individual IP addresses. Self-similarity can also be regarded as a dynamical characteristic of the traffic, related to the extreme variability independent of the time-scale at which one looks at the time series.
Besides values distribution that contributes to the traffic volume part of the total traffic sample path, the number of IP addresses that are sending traffic during a particular minute supplies the other part of the traffic generation process, the dynamical one. Using the number of "active IP addresses" over one-minute intervals can be regarded as a broad picture for the number of instantaneous flows. It approximates the real process of the user's (be they humans or machines) communicating over the Internet. Figure 5 presents the evolution of the number of IP addresses (left),

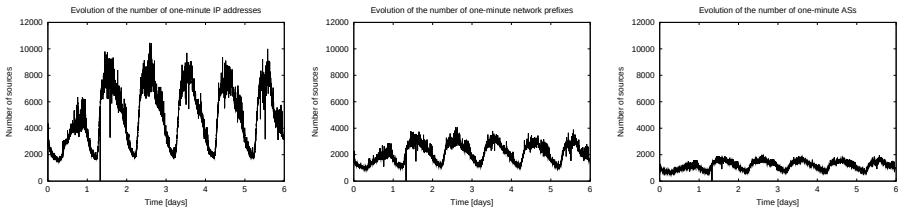

Fig. 5. Evolution of traffic sources: IP addresses (left), network prefixes (middle) and ASs (right).

network prefixes (middle) and autonomous systems (right) that are sending traffic during a given one-minute interval. A "network prefix" is the aggregation of all IP addresses contained within a domain corresponding to a `<prefix/netmask length>` pair appearing in the BGP routing table of the studied ISP. An "autonomous system" (AS) is the aggregation of all network prefixes contained within a domain corresponding to the destination AS number appearing in the *AS path* information of the BGP routing table of the studied ISP. Figure 5 shows the evolution of the three types of traffic sources with an identical y-axis scale. All traffic sources types exhibit a similar variability, the only difference being their absolute number. The average number of sources over the measurements is 4912 for IP addresses, 2100 for network prefixes and 1176 for ASs. A more detailed study of interdomain flows can be found in [UB00].

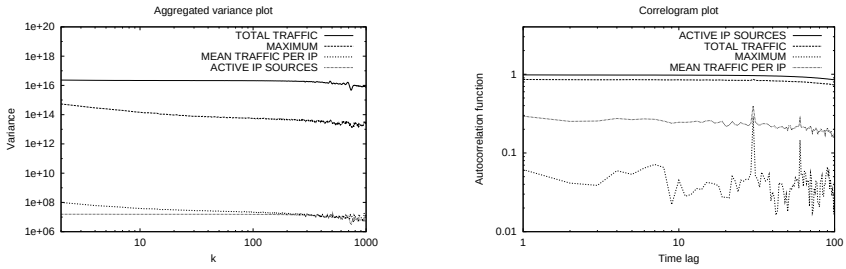

Fig. 6. Comparison of H for traffic components.

As presented on figure 6, the number of "active" IP addresses during one-minute intervals exhibits the same kind of self-similarity than total traffic. Even if it is difficult to assess whether some component of the traffic exhibits a stronger self-similarity than another, looking at the *aggregated variance* and the *correlogram* of some components of the traffic will give us some useful information. The leftmost part of figure 6 shows the *aggregated variance* plot for the evolutions of: the total traffic, the number of IP sources that are sending traffic during every minute, the mean traffic per IP source and the maximum amount of bytes sent by an IP source during every minute. The rightmost part of figure 6 shows the

correlogram for the same four concepts. *Aggregated variance* is almost constant for `IP sources` and `total traffic` while it decreases slightly for the two other concepts. The *correlogram* also indicates a stronger self-similarity for IP sources and total traffic, while the other two concepts have correlations that are close to or below the $2/\sqrt{n}$ confidence limit (around 0.02 for our sample) that prevents us from inferring anything about actual correlations for them. Even if such behavior does not automatically imply that IP sources could better explain self-similarity in the sample path, it constitutes a plausible root. Network prefixes and ASs provide similar results to IP addresses but are not presented due to space limitations.

4.3 Heavy-Tails Revisited

Having studied two possible causes for self-similarity in the traffic cannot tell which one is the most important, if such ever exists. It has already been shown [CB96] that heavy-tails in the distributional properties of the transferred objects was a sufficient condition for generating self-similar traffic under a wide range of conditions. The purpose of this section is to better show the role of heavy-tails for what concerns self-similar traffic. Let us try the following experiment. Assume we can change the distribution of the sizes of the one-minute transfers so that instead of being heavy-tailed, they were exponentially distributed. For that purpose, we use the characteristics of our traffic trace as a basis and modify the one-minute distributions of the values appearing for the IP sources so that they conform to an exponential distribution, but under the constraint that the total simulated traffic be the same as for the original total traffic trace. Exponential distributions make that possible because they are completely defined by their mean, so that given a particular number of IP sources that are active during a given minute, it is possible to distribute the traffic values exponentially so that the simulated total traffic is equal to the one of the traffic trace.
The procedure for generating every one-minute distribution is the following:

1. Determine in the trace the total amount of traffic T_k (in bytes) seen during minute k as well as the number N_k of distinct IP addresses that are sending traffic during that minute.
2. For each minute k, generate an approximation of the exponential distribution with mean T_k/N_k so that the simulated traffic corresponds to a total of about T_k bytes and a number of points of about N_k points by relying on the exponential distribution formula
$$P(X = x) = \frac{N_k}{T_k} e^{-(N_k/T_k)x}. \quad (8)$$

Step 2 requires that we follow a continuous exponential distribution while in fact we need N_k points. Since we want to generate an equivalent of N_k points that accounts for T_k bytes, we generate a discrete exponential distribution that approximates the continuous one. Thus, instead of integrating a continuous curve, we sum an exponentially distributed collection of points so that we have the equivalent of N_k points providing T_k bytes. Generating this discrete set of points requires that we stop

summing the discrete exponential at some value, i.e. the higher bound of the given one-minute simulated sample. We choose this point to be $\max_k$ which corresponds to the value x such that $P(X = x) \geq 1/N_k$, thus the highest value of the discrete distribution which should occur if we had actually N_k discrete points occurring in the simulated distribution. The exact value of $\max_k$ is found thanks to

$$\frac{N_k^2}{T_k}\, e^{-(N_k/T_k)\max_k} \geq \frac{1}{N_k}, \tag{9}$$

saying that the value max_k has to occur at least one time amongst N_k. The remaining of the algorithm attributes to each discrete value from 0 to $\max_k$ its frequency of occurrence as well as its traffic volume in bytes, so that the sum of the frequency of occurrence over all generated values sums up to about N_k and the traffic volume to about T_k. Figure 7 presents the pseudo-code for the generation of the simulated distributions.

```
foreach minute k {
    foreach value = 0 to max_k {
        // Attributing to value its frequency of occurence
        frequency(value) = (N_k^2/T_k) * e^{-(N_k/T_k) value}
        // Attributing to value its traffic volume
        volume(value) = value * (N_k^2/T_k) * e^{-(N_k/T_k) value}
    }
}
```

Fig. 7. Pseudo-code for generating exponentially distributed values.

The leftmost part of figure 8 shows the probability mass of the simulated exponential distribution for the whole simulation. The simulated distribution is obviously not a perfect exponential one. As can be seen on figure 8, the end of the exponential tail deviates somewhat from its ideal trajectory due to the discrete nature of our simulation.

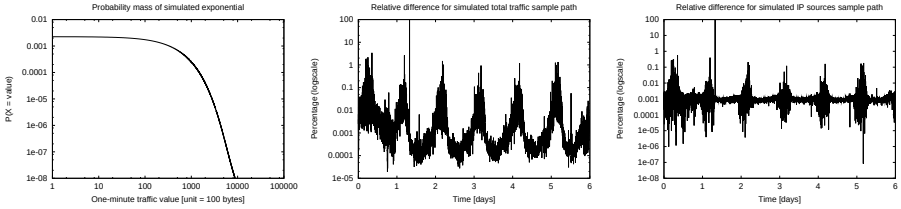

Fig. 8. Simulating exponentially distributed sources: exponential probability mass (left), relative difference in total traffic evolution (middle) and relative difference in IP sources evolution (right).

The other two graphs of figure 8 show the evolution in the "relative difference" between the simulated one-minute total traffic and T_k

$$\mid (T_k - \sum_{value=0}^{max_k} volume(value))/T_k \mid \tag{10}$$

(middle) and between the simulated number of IP sources and N_k

$$\mid (N_k - \sum_{value=0}^{max_k} frequency(value))/N_k \mid \tag{11}$$

(right). The relative difference between simulated and original sample paths shows how close the simulation comes to the original, with a global difference smaller than 0.01 % for total volume and smaller than 0.001 % for the number of IP sources. Note that it is possible to better approximate T_k and N_k by using a higher resolution on the values and by cutting the discrete distribution at a larger maximum value.

While rather simple, this simulation shows that heavy-tails in the distribution of traffic values are not the sole responsible for generating self-similar traffic on time-scales between minutes and hours. We have shown that by changing the distribution of the amount of traffic IP addresses are sending during any particular minute in a manner that prevents high bursts to occur does not preclude self-similarity to happen in the total traffic sample path. Even if this does not prove that the arrival process of "active" IP sources is responsible for that, we can be sure about the limited role of large bursts (and also heavy-tails) in terms of traffic volume for what concerns self-similarity.

5 Conclusion

In this paper, we have analyzed in details a six days long trace of all the incoming interdomain traffic of an ISP to better understand the long-term self-similarity of Internet traffic. The detailed analysis of the long-term traffic self-similarity has produced three important findings.

First, at this long time-scale, the interdomain traffic received by an ISP appears to be self-similar. This confirms the analysis of shorter traces available in the literature.

Second, we have shown that this self-similarity could be explained by two different elements. The first is the total amount of traffic sent by the external IP addresses that follows a power tail distribution. The second, and this point has been rarely analyzed in the literature, is that the number of sources (i.e. IP addresses) also exhibits a self-similar path.

Third, we have shown that it is possible to simulate self-similar traffic by considering exponentially distributed traffic values and self-similar sources. This indicates that the large bursts have a limited role in traffic self-similarity and that the number of active sources plays an important role in the self-similarity of Internet traffic.

Acknowledgement

This paper would not have been written without the traffic trace provided by Marc Roger from the Belgian research network Belnet.

References

[Ber94] J. Beran. Statistics for Long-Memory Processes. Monographs on Statistics and Applied Probability, Chapman & Hall, 1994.

[CB96] M. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *SIGMETRICS'96*, pages 160–169, May 1996.

[Cis99] Cisco. NetFlow services and applications. White paper, available from `http://www.cisco.com/warp/public/732/netflow`, 1999.

[Hil75] B. Hill. A simple approach to inference about the tail of a distribution. Annals of Statistics, 3(1975), pages 1163–1174, 1975.

[LTW+94] W. Leland, M. Taqqu, W. Willinger and D. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM* Transactions on Networking, February 1994.

[PF95] V. Paxson and S. Floyd. Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM* Transactions on Networking, 3(3):226-244, June 1995.

[PKC96] K. Park, G. Kim and M. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In Proc. Fourth International Conference on Network Protocols, October 1996.

[PKC97] K. Park, G. Kim, and M. Crovella. On the effect of traffic self-similarity on network performance. In Proc. of SPIE International Conference on Performance and Control of Network Systems, November 1997.

[Res97] S. Resnick. Heavy Tail Modeling and Teletraffic Data. Annals of Statistics, 25(1997), pages 1805-1869, 1997.

[Res98] S. Resnick. Why Non-Linearities Can Ruin the Heavy-Tailed Modeler's Day. In "A Practical Guide to Heavy Tails: Statistical Techniques and Applications", Birkhauser, Boston, 1998.

[RL97] S. Robert and J.-Y. Le Boudec. New models for self-similar traffic. Performance Evaluation 30(1-2), pages. 57–68, 1997.

[MC00] S. McCreary and K. Claffy. Trends in wide area IP traffic patterns : a view from Ames Internet Exchange. Available from `http://www.caida.org/outreach/papers/AIX0005/`, 2000.

[TTW95] M. Taqqu, V. Teverovsky and W. Willinger. Estimators for long-range dependence: an empirical study. Fractals, (3):4:785-798, 1995.

[TT98] M. Taqqu and G. Samorodnitsky. On Estimating the Intensity of Long-Range Dependence in Finite and Infinite Variance Time Series. In "A Practical Guide to Heavy Tails: Statistical Techniques and Applications", Birkhauser, Boston, 1998.

[TWS97] M. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *ACM/SIGCOMM* Computer Communications Review, 27(1997), pages 5–23, 1997.

[UB00] S. Uhlig and O. Bonaventure. On the Cost of Using MPLS for Interdomain Traffic. In Proc. of *QOFIS2000*, Berlin, September 2000.

[WPT98] W. Willinger, V. Paxson, and M. Taqqu. Self-similarity and heavy tails: Structural modeling of network Traffic. In "A Practical Guide to Heavy Tails: Statistical Techniques and Applications", Birkhauser Verlag, Boston, 1998.