Preprint 1996. To appear in the journal "Fracrals".

Is network traffic self-similar or multifractal? *

Murad S. Taqqu
Boston University

Vadim Teverovsky
Boston University

Walter Willinger
Bellcore

May 29, 1996

Abstract

This note addresses the question of whether self-similar processes are sufficient to model packet network traffic, or whether a broader class of multifractal processes is needed. By using the absolute moments of aggregate traffic measurements, we conclude that measured local-area network (LAN) and wide-area network (WAN) traffic traces, with the sample means subtracted, are well modelled by random processes that are either exactly or asymtotically self-similar.

1 Introduction

Self-similarity in modern packet network traffic has recently been the focus of much attention, prompted by the results of Leland, Taqqu, Willinger, Wilson [6] (see also [13]), who examined several high-resolution Ethernet local-area network (LAN) traffic traces and found them to be consistent with statistical self-similarity. Similar findings were reported by Paxson and Floyd [7] for wide-area network (WAN) traffic measurements. This note addresses the question of whether for a more accurate description of these data sets, a broader model class, namely that of multifractal or signed multifractal processes, has to be considered. We briefly characterize these processes, indicate a method for distinguishing between them, and then apply the method to some of the traffic traces that have been studied in [6, 13]. We conclude that, with their sample means subtracted, these data sets are well-modeled by random processes that are either exactly or asymptotically self-similar. This finding is further confirmed by our recent analysis of LAN/WAN traces at the level of individual source-destination pairs, which provides a plausible physical explanation of the observed self-similar nature of modern network traffic (see Willinger, Taqqu, Sherman and Wilson [13]).

*This research was partially supported by the NSF Grants NCR-9404931 and DMS-9404093 at Boston University.

Historically, the thrust of multifractal theory has been to study the distribution of the fractal dimensions of subsets generated by a deterministic or random mechanism, for example, the binomial decomposition of the interval $[0, 1]$ (Evertsz and Mandelbrot [3]). They can also be used, however, to analyze traces (sample paths) generated by either a random or non-random mechanism (Jaffard [5]). The multifractal analysis is particularly appropriate in investigating the local fluctuations of these traces. The corresponding fractal dimensions measure the "roughness" of the sample paths at localized points within small subsets often of negligible length.

This interest in small scale behavior has been typically motivated by physical applications such as turbulence (see Farge, Kevlahan, Perrier and Goirand [4] for a recent review). However, for other applications of interest the focus is often less on small scale (high frequency) but more on large scale (low frequency) behavior. For example, in the context of high-speed network traffic modeling, local fluctuations are known to be highly sensitive to network conditions (e.g., user population, mix of applications, protocols, network topology). Effective and parsimonious traffic modeling for modern networks will avoid describing such localized features and will instead concentrate on more robust traffic characteristics (i.e., large scale behavior) which can be expected to be present under a wide range of network structures. Irrespective of the application, extrapolation to smaller or larger scales is generally possible, because analysis of real-life time series often show lower and upper scale cutoffs, which are sometimes spaced by several orders of magnitude.

The paper is organized as follows. In the next section, we define self-similar (monofractal) processes and characterize the (signed) multifractal ones. Unfortunately, there exists no standard terminology, and we rely in this paper largely on terminology introduced by B. Mandelbrot. In Section 3, we indicate how to determine whether a given trace is better modeled by a self-similar or multifractal process. We apply this methodology to LAN/WAN traffic measurements in Section 4 and present our conclusions in Section 5.

2 Self-similarity and the scaling of moments

There are a number of different, not equivalent, definitions of self-similarity. The standard one states that a continuous-time process $Y = \{Y(t), t \in T\}$ is *self-similar* (with self-similarity parameter H) if it satisfies the condition:

$$Y(t) \stackrel{d}{=} a^{-H} Y(at), \quad \forall t \in T, \quad \forall a > 0, \ 0 \le H < 1, \tag{2.1}$$

where the equality is in the sense of finite-dimensional distributions. The canonical example of such a process is Fractional Brownian Motion (see Samorodnitsky and Taqqu [8], Beran [1]), Brownian Motion if $H = 1/2$. While a process Y satisfying (2.1) can never be stationary, it is typically assumed to have stationary increments.

A second definition of self-similarity, more appropriate in the context of standard time series theory, involves a stationary sequence $X = \{X(i), i \geq 1\}$. Let

$$X^{(m)}(k) = 1/m \sum_{i=(k-1)m+1}^{km} X(i) \qquad k = 1, 2, \ldots, \tag{2.2}$$

be the corresponding aggregated sequence with level of aggregation m, obtained by dividing the original series X, into non-overlapping blocks of size m and averaging over each block. The index, k, labels the block. If X is the increment process of a self-similar process Y defined in (2.1), that is, $X(i) = Y(i+1) - Y(i)$, then for all integers m,

$$X \stackrel{d}{=} m^{1-H} X^{(m)}. \tag{2.3}$$

A stationary sequence $X = \{X(i), i \geq 1\}$ is called *exactly self-similar* if it satisfies (2.3) for all aggregation levels m. This second definition of self-similarity is closely related (but not equivalent) to the first. A stationary sequence $X(i), i \geq 1$ is said to be *asymptotically self-similar* if (2.3) holds as $m \to \infty$. Similarly, we call a covariance-stationary sequence $X(i), i \geq 1$ *exactly second-order self-similar* or *asymptotically second-order self-similar* if $m^{1-H}X^{(m)}$ has the same variance and autocorrelation as X, for all m, or as $m \to \infty$. Note that under Gaussianity, the last two definitions (i.e., self-similarity and second-order self-similarity) are equivalent, and in the following, we will use either one as our working definition for self-similarity (approximate Gaussianity has been verified for the data sets considered here in [6]).

Self-similarity is often investigated not through the equality of finite-dimensional distributions, but through the behavior of the absolute moments. Thus consider

$$\mu^{(m)}(q) := E|X^{(m)}|^q = E\left|\frac{1}{m}\sum_{i=1}^{m} X(i)\right|^q. \tag{2.4}$$

If X is self-similar, then $\mu^{(m)}(q)$ is proportional to $m^{\beta(q)}$, that is, $\log \mu^{(m)}(q)$ is linear in $\log m$ for a fixed q:

$$\log \mu^{(m)}(q) = \beta(q) \log m + C(q). \tag{2.5}$$

In addition, the exponent $\beta(q)$ is linear with respect to q. In fact, since $X^{(m)}(i) \stackrel{d}{=} m^{H-1}X(i)$, we have

$$\beta(q) = q(H-1) \tag{2.6}$$

Thus, a third definition of self-similarity (implied by but not equivalent to the second definition) is simply that the moments must scale as in (2.5), and that $\beta(q)$ satisfies (2.6).

The third definition of a self-similar process given above can be generalized to that of *multifractal* processes. A non-negative process $X(t)$ will be called *multifractal* if the logarithms of the absolute moments scale linearly with the logarithm of the aggregation level m. Multifractals are commonly constructed through multiplicative cascades (see Evertsz and Mandelbrot [3]).

If the multifractal can take positive and negative values, then it is referred to as a *signed multifractal* (the term "multiaffine" is sometimes used instead of "signed multifractal": see for example Marshak, Davis, Cahalan, Wiscombe [2]). The key point is that, unlike self-similar processes, the scaling exponent $\beta(q)$ in (2.5) is not required to be linear in q. Thus, signed multifractal processes are a generalization of self-similar processes. To discover whether a process is a (signed) multifractal or is self-similar it it not sufficient to examine second moment properties. One must analyze higher moments as well.

We conclude this section with two remarks about X and $X - EX$.

Remarks:

- If X is a positive stationary non-degenerate time series, such as the Ethernet measurements, then neither X nor $X - EX$ can be, strictly speaking, self-similar. However, the sequence $X - EX$ may be asymptotically self-similar. When the asymptotic regime already holds for fairly small m, we commonly say that $X - EX$ is self-similar.

- A stationary time series X cannot be self-similar or asymptotically self-similar if it has non-zero mean. Indeed, if X is self-similar with mean EX, then Relation (2.3) implies $EX = m^{1-H}EX^{(m)} = m^{1-H}EX$. If X is asymptotically self-similar, then $EX \sim m^{1-H}EX$ for large m. Thus, in either case, $EX = 0$. The process X, however, could be multifractal.

3 Methodology

Consider a non-negative time series $X(1), \ldots, X(N)$, e.g. an hour-long Ethernet trace representing the number of packets/bytes sent over the Ethernet every 10 milliseconds. The goal is to check whether X and/or $X - EX$ is adequately described by a self-similar process, a multifractal, or a signed multifractal (non self-similar) process, or whether none of these models are consistent with the data.

Consider X first. To estimate the q^{th} moment of X, $\mu^{(m)}(q)$, we use the q^{th} sample absolute moment of the aggregated series $X^{(m)}$,

$$\widehat{\mu}^{(m)}(q) = \frac{1}{N/m} \sum_{k=1}^{N/m} \left| X^{(m)}(k) \right|^q . \tag{3.1}$$

If $\log \widehat{\mu}^{(m)}(q)$ scales linearly with $\log m$, then X could indeed be modeled by a multifractal process. As noted above, such a process cannot be self-similar since $EX \neq 0$.

If $\log \widehat{\mu}^{(m)}(q)$ does not scale linearly with $\log m$, then one turns to the series $X - EX$. Let now $\widehat{\mu}_0^{(m)}(q)$ denote the q^{th} sample moment corresponding to $X - EX$, that is,

$$\widehat{\mu}_0^{(m)}(q) = \frac{1}{N/m} \sum_{k=1}^{N/m} \left| X^{(m)}(k) - \frac{1}{N} \sum_{i=1}^{N} X(i) \right|^q . \tag{3.2}$$

If $\log \widehat{\mu}_0^{(m)}(q)$ scales linearly with $\log m$, then a (signed) multifractal model can apply. If, however, in addition, $\beta(q)$ is linear in q, then a self-similar model is adequate.

Remarks:

- In practice, one considers only $q = 1, 2, 3, 4$ because high-order sample moments are not well-behaved. Non-integer values of q, e.g., $q = 0.5,\ q = 1.5$, do not add much more information. Negative values of q, e.g $q = -0.5,\ q = -1$, put undue weight on the center of the distribution and the estimate of the corresponding moments tends to be very erratic.

- In our plots the logarithms of the sample absolute moments all start from 0 because we normalized the sample absolute moment at aggregation level $m = 1$ to equal 1.

- Although it may be possible to use a sliding aggregation window to get more data points in the aggregated series, it is not necessary to do so here because the Ethernet data series are very long ($360,000$ points). As seen in Taqqu, Teverovsky and Willinger [11], even for significantly shorter data series using a sliding window has a negligible effect (compare the Higuchi method versus the Absolute Value method in [11]).

4 Application to LAN/WAN traffic measurements

In the present study, we use two sets of Ethernet LAN traffic measurements which were investigated earlier (e.g., see [6, 9, 10, 12]), and one WAN traffic trace that was studied more recently in [13]. All three data sets were found to be consistent with statistical self-similarity, once their means are subtracted. The first data set was collected on an "internal" Bellcore Ethernet LAN during a "normal traffic hour" in August 1989 and represents the number of bytes sent over the Ethernet every 10 milliseconds. The second data set gives the number of Ethernet packets at the 10 millisecond time scale, for the same "normal traffic hour" in August 1989. Finally, the third data set was collected on Bellcore's link to the outside world during December 1994; the measurements are made up entirely of remote packets, resulting in a trace that is representative of WAN traffic. The sequence at hand represents the number of remote packets per 10 milliseconds during a "typical" traffic hour in December 1994. Each of the three data set consists of 360,000 observations, and the time series are obviously non-negative. The byte series has a sample mean of 1058 and a sample standard deviation of 1947. The August '89 packet series has a sample mean of 2.69 and a sample standard deviation of 2.42, while the December '94 packet series has a sample mean of 2.95 and a sample standard deviation of 2.07.

We first calculated the absolute moments of the three traffic traces without subtracting the sample mean. The resulting logarithms of the absolute moments, $\widehat{\mu}^{(m)}(q)$, do not scale linearly with the logarithm of m as illustrated in Figure 1 for the August '89 data and in Figure 2 for the December '94 trace. Thus, if the sample mean is not subtracted, the series cannot be modeled by a self-similar, multifractal, or signed multifractal process.

We then subtracted the sample means of the series. In Figures 3 and 4 we plot $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$ for $q = 1, 2, 3, 4$ for the August '89 *byte* and *packet* series, respectively. For both time series, we observe that for all q's considered, the logarithms of the moments scale linearly with $\log m$, for a range of aggregation levels that covers nearly 5 orders of magnitude. Moreover, when estimating the slopes $\widehat{\beta}(q)$ using relation (2.6), we get

$$\widehat{H}(q) = 1 + \frac{1}{q}\widehat{\beta}(q),$$

and find that the estimates $\widehat{H}(q)$ of H do not depend on the order of the moment q. For $q = 1, 2, 3, 4$ we get $\widehat{H} = \widehat{H}(q) = .80 \pm .03$ for the *byte* series and $\widehat{H} = \widehat{H}(q) = .85 \pm .01$ for the *packet* series. The analysis of these two data sets reported in [6, 9, 10, 12] essentially focused on the second moments. The findings and estimates obtained here are consistent with these earlier analyses. Moreover, they are in full agreement with a physical "explanation" of the observed (exact) self-similarity of these two traffic traces put forward in [13]. In its simplest form, this explanation asserts that aggregate packet traffic is self-similar if it is the superposition of many i.i.d. *ON/OFF* sources with infinite variance *ON/OFF* periods. A detailed statistical analysis of the August '89 "normal traffic hour" at the level of individual source-destination pairs (see [13]) revealed that, indeeed, (i) the traffic trace at hand is made up of many individual components (i.e., source-destination pairs), (ii) the individual component processes are consistent with an *ON/OFF* source model, and (iii) the *ON/OFF* periods of the individual source-destination pairs are, in general, fully consistent with the infinite variance phenomenon. Combining this physical understanding of network traffic self-similarity with the statistical evidence in favor of self-similar traffic behavior, it is reasonable to expect that future LAN traffic (if made up of sufficiently many individual traffic streams) will exhibit self-similar features, despite constantly changing network conditions, user populations, and the appearance of new applications and services.

In Figure 5, we plot $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$ for $q = 1, 2, 3, 4$ for the December

'94 *packet* series. In contrast to Figures 3 and 4 where the scaling region extends over all aggregation levels considered, Figure 5 shows that for all q's considered, the logarithms of the moments scale linearly with $\log m$, for aggregation level m exceeding about 10 or 20. Focusing on this scaling region when estimating H, we find that the estimates $\widehat{H}(q)$ do not depend on q and yield $\widehat{H} = .88 \pm .01$. Observe that for small values of m (i.e., $m \leq 10$), Figure 5 shows a distinctly different scaling behavior than for larger m-values. This observation is in full agreement with the results reported in [13], namely that he December '94 traffic trace is not well-modeled by an exact self-similar process but is consistent with asymptotic self-similarity. In explaining this difference between the two Ethernet LAN traces considered above and the December '94 WAN trace at hand, Willinger, Taqqu, Sherman and Wilson [13] show that the latter consists of about 50% *Mbone* traffic and 50% *non-Mbone* traffic, and that the *Mbone* traffic component is not consistent with exact self-similarity but is well-modeled using asymptotically self-similar processes. Moreover, since *Mbone* traffic represents a significant part of the total traffic, its distinctive feature of asymptotic self-similarity will be inherited by the combined traffic, despite the fact that the *non-Mbone* component is well-modeled by an exactly self-similar process (same explanation as in the case of the August '89 traces). These findings suggest that future WAN traffic can be expected to exhibit self-similar traffic characteristics, where self-similarity may come in different "shades", from exact self-similarity to different forms of asymptotic self-similarity, depending on the traffic mix (i.e., main applications and services) and the level of aggregation (i.e., number of hosts).

Remarks:

- In all the plots considerable scattering occurs for the very high values of m, because the averaging for the sample moments is done over so few blocks. Often, small values of m have also to be ignored because, in general (and especially in situations of asymptotic self-similarity), some averaging is required before self-similarity becomes apparent. This is why it is important to restrict the estimation of H to the central regions (indicated by darker points in the plots).

- The first and last row of plots in Figure 1 display $\log_{10} \widehat{\mu}^{(m)}(q)$ versus $\log_{10} m$ for the August '89 Ethernet series with mean not subtracted. Trying to fit a straight line to estimate H, we get (a) *bytes*, $q = 2$, $\widehat{H} = .93$; (b) *bytes*, $q = 3$, $\widehat{H} = .89$; (c) *packets* $q = 2$, $\widehat{H} = .97$; (d) *packets* $q = 3$, $\widehat{H} = .95$. The plots, however, are clearly not linear, and these estimates of H not meaningful.

- Consider now the first and last plots in Figure 2. They display $\log_{10} \widehat{\mu}^{(m)}(q)$ versus $\log_{10} m$ for the December '94 Ethernet series with mean not subtracted. Since the small values of m have to be dropped anyway (see Figure 5), one could be tempted to use the "slope" in the remaining part of either graph to estimate H. One gets again a non-meaningful result of $\widehat{H} = .99$ in both plots.

- Although it is, in general, difficult to determine the exact form of the sample moments when the mean has not been subtracted, it is relatively easy to do so if $q = 2$, because the sample second moment of a time series equals the sample mean of the series squared plus the sample variance of the series:

$$\widehat{\mu}^{(m)}(2) = \widehat{\mu}_0^{(m)}(2) + \left(\frac{1}{N}\sum_{i=1}^{N} X(i)\right)^2.$$

  Let us then fit the second sample absolute moment $\widehat{\mu}^{(m)}(2)$ of $X^{(m)}$ to a function of the form

$$C_1 + C_2 m^{2H-2}. \tag{4.1}$$

  C_1 should equal the square of the mean of the series, and does not depend on m. The term $C_2 m^{2H-2}$ is the sample variance of the aggregated series, assuming self-similarity. Proceeding in this fashion, we obtained an estimate of $\widehat{H} = .83$ for the August '89 Ethernet byte series, and $\widehat{H} = .85$ for the August '89 Ethernet packet series (Figure 1). These are in agreement with the estimates of H obtained above using $\widehat{\mu}_0^{(m)}(2)$, that is, when the sample means were subtracted.

- The estimate of H in (4.1) is very sensitive to the choice of the lowest value of m at which the curve is fitted. This is because the curve is like a decreasing exponential and the fitting criteria is least squared. This is particularly evident for the December '94 data which is only self-similar at relatively high aggregation levels m. For example, if the fitting of $C_1 + C_2 m^{2H-2}$ starts at $\log_{10} m = 1.0,\ 1.1,\ 1.2,\ 1.3,\ 1.4$, one obtains respectively a corresponding $\widehat{H}$ of $.75,\ .80,\ .84,\ .87,\ .89$. Fitting in log-log scale, as is done in Figure 5, is more robust and gives the following values respectively: $\widehat{H} = .87,\ .87,\ .88,\ .88,\ .88$. This problem is less acute for the '89 Ethernet data which is self-similar for smaller values of m.

- In general, fitting $\log \widehat{\mu}_0^{(m)}(2)$ is better and more reliable than fitting $\widehat{\mu}^{(m)}(2)$. Also, fewer parameters have to be estimated,

- Figures 3 and 4 display $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$ for the August '89 Ethernet byte and packet series with mean subtracted. For a given series, the plots with different q's exhibit similar fluctuations, i.e. when the points dip down for $q = 1$, they also do so for the other q's. This is not surprising, since the values of the aggregated series are identical for the different q's, the only difference being the power q to which they are raised. Observe that the actual slopes $\widehat{\beta}(q)$ of the plots are quite different for the various q's.

- Typically, estimates of H obtained through different estimation methods fall within a range of ± 0.1. The minuscule variation in the values we obtained for $\widehat{H}(q),\ q = 1, 2, 3, 4$ is thus particularly remarkable.

5 Conclusion

We have investigated some of the Ethernet LAN traffic traces considered by Leland, Taqqu, Willinger, Wilson [6], Teverovsky and Taqqu [12], and Willinger, Taqqu, Sherman and Willinger [13], focusing on the question whether self-similar processes suffice for adequately describing measured LAN/WAN traffic, or whether the broader class of multifractal processes is required for modeling modern packet network traffic in typical LAN/WAN environments. Our results confirm the conclusion of our earlier studies, namely that the deviations of measured Ethernet LAN traffic from its mean are consistent and well-modeled with statistical self-similarity. The results also indicate that it is not necessary to increase the model class to that of general multifractal processes in order to model the LAN/WAN traffic considered here. However, for other LAN/WAN traffic traces for which the conditions for self-similarity are not satified (e.g., insufficient aggregation level, significant effect of flow control on certain users), general multifractal models may be more appropriate. When such an extension to multifractal traffic models will become necessary remains an open problem and is closely related to the issue of providing a physical "explanation" for this model class in the networking context.

Acknowledgement We would like to thank Professor D. Veneziano for insightful comments.

References

[1] J. Beran. *Statistics for Long-Memory Processes.* Chapman & Hall, New York, 1994.

[2] A. Davis, A. Marshak, W. Wiscombe, and R. Cahalan. Multifractal characterizations of nonstationarity and intermittency in geophysical fields: Observed, retrieved, or simulated. *Journal of Geophysical Research*, 99(D4):8055–8072, 1994.

[3] J. G. Evertsz and B. B. Mandelbrot. Multifractal measures. In H. O. Peitgen, H. Jürgens, and D. Saupe, editors, *Chaos and Fractals*, pages 921–953. Springer-Verlag, New York, 1992. Appendix B.

[4] M. Farge, N. Kevlahan, V. Perrier, and E. Goirand. Wavelets and turbulence. *IEEE Proceedings*, 84(6), 1996. Special issue on Wavelets. To appear.

[5] S. Jaffard. Formalisme multifractal pour les fonctions. *Comptes Rendus de l'Académie des Sciences de Paris*, 317:745–750, 1993. Série I.

[6] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic (Extended version). *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.

[7] V. Paxson and S. Floyd. Wide area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1995.

[8] G. Samorodnitsky and M. S. Taqqu. *Stable Non-Gaussian Processes: Stochastic Models with Infinite Variance.* Chapman and Hall, New York, London, 1994.

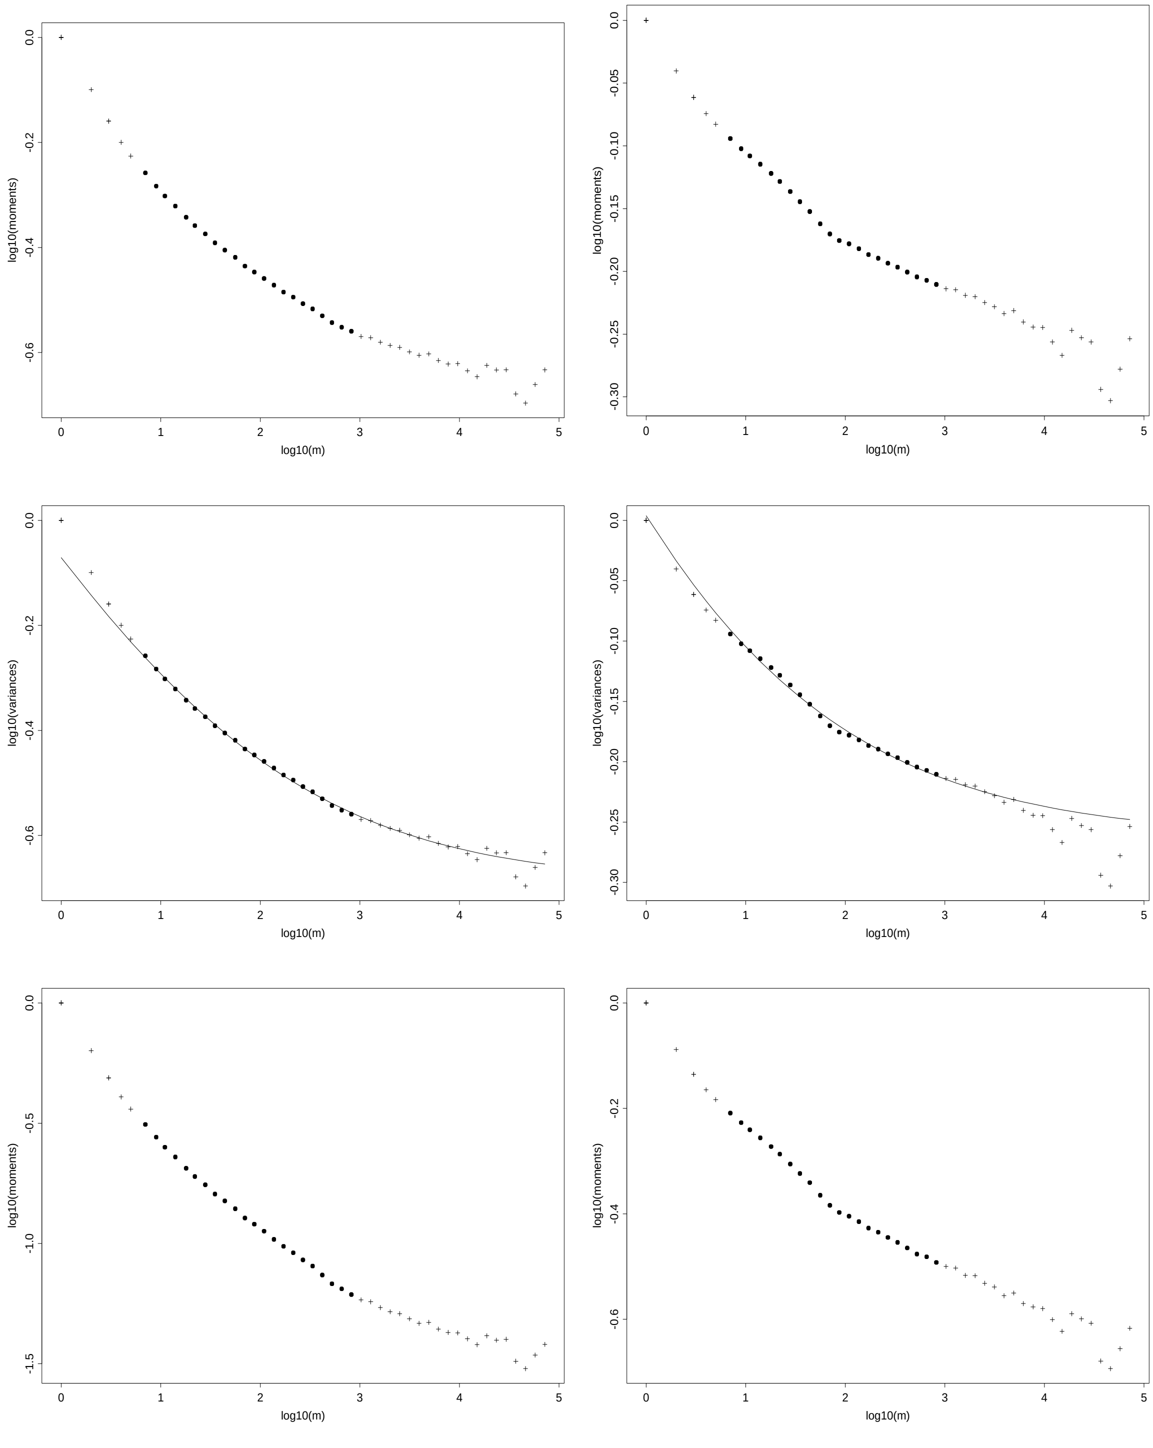

Figure 1: Moment plots for the August '89 Ethernet series with mean NOT subtracted. Right column (byte series: $q = 2$, $q = 2$, $q = 3$), left colunn (packet series: $q = 2$, $q = 2$, $q = 3$). The plots in the first and last row display $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$, and are clearly non-linear. The middle plots are identical to the ones in the top row, except that the curve (4.1) is superimposed.

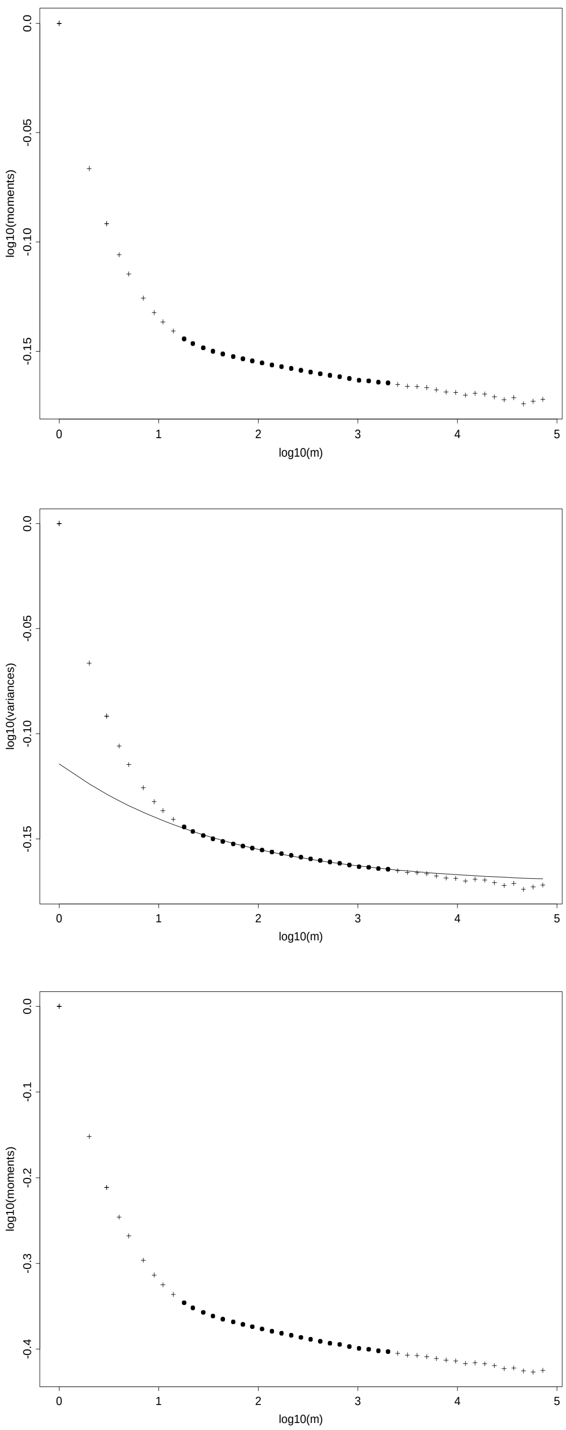

Figure 2: Moment plots for the December '94 Ethernet series with mean NOT subtracted (packet series: $q = 2$, $q = 2$, $q = 3$). The plots in the first and last row display $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$. They give unreasonable values of H even if one starts plotting them at $\log_{10} m = 1.2$. The middle plot is identical to the top one, except that the curve (4.1) is superimposed.

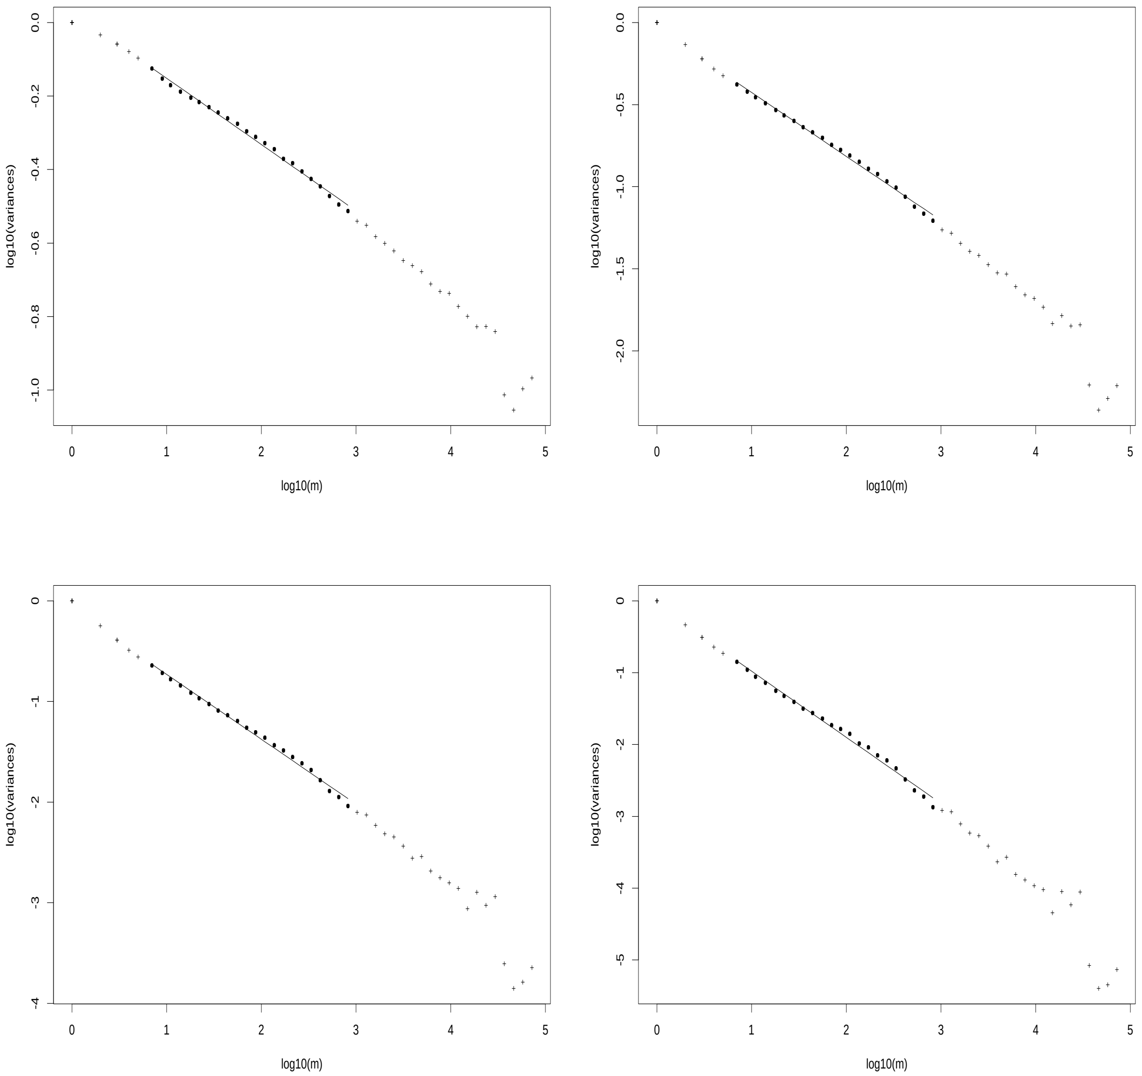

Figure 3: Plots of $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$ for the August '89 Ethernet *byte* series with mean subtracted. The value of q and the estimated slope $\widehat{H}$ is as follows. Going from left to right: (a) $q = 1$, $\widehat{H} = .82$; (b) $q = 2$, $\widehat{H} = .80$; (c) $q = 3$, $\widehat{H} = .78$; (d) $q = 4$, $\widehat{H} = .77$.

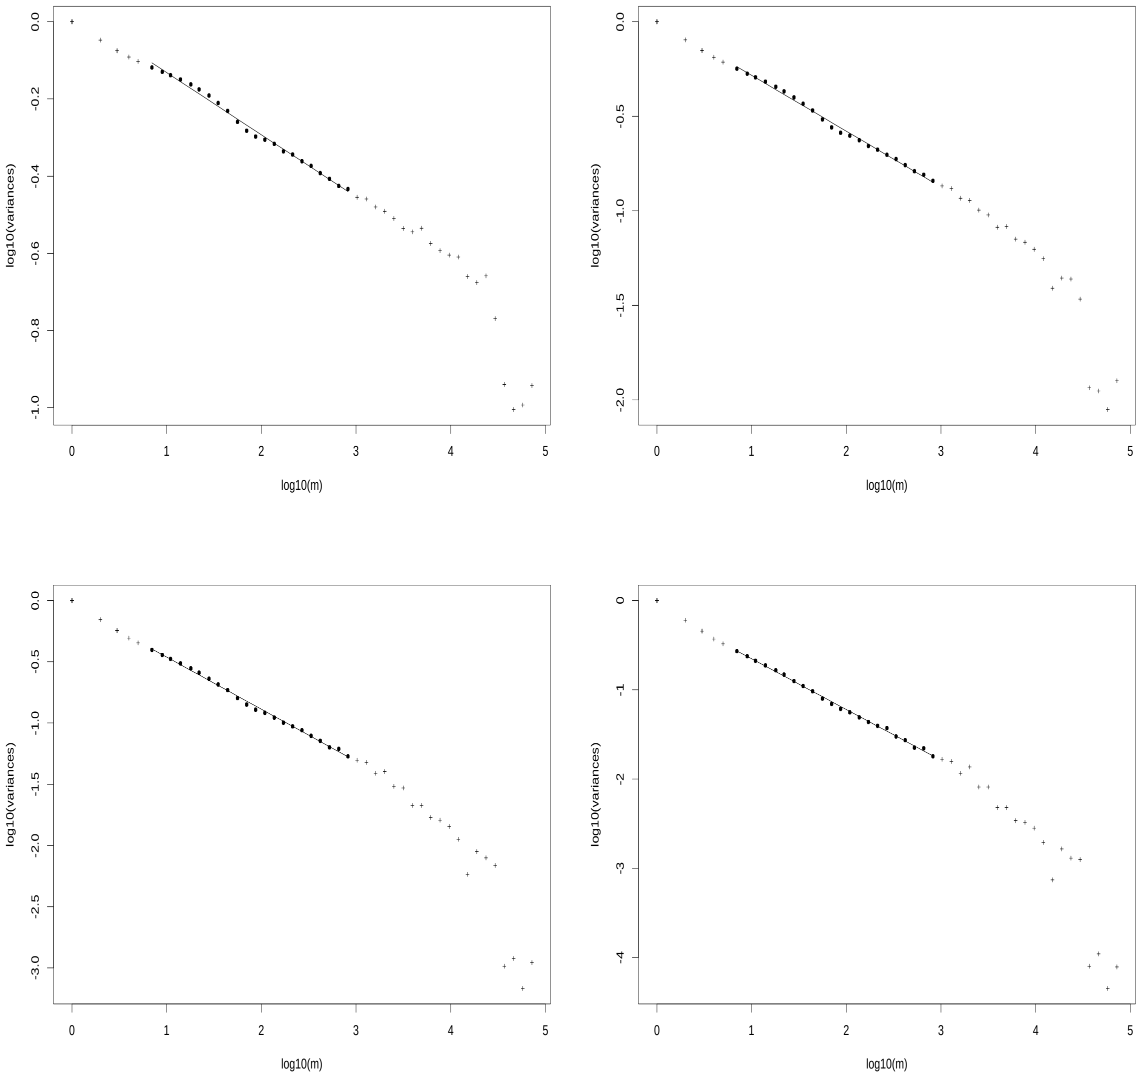

Figure 4: Plots of $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$ for the August '89 Ethernet *packet* series with mean subtracted. The values of q and the estimated slope $\widehat{H}$ is as follows. Going from left to right: (a) $q = 1$, $\widehat{H} = .84$; (b) $q = 2$, $\widehat{H} = .85$; (c) $q = 3$, $\widehat{H} = .86$; (d) $q = 4$, $\widehat{H} = .86$.
X—

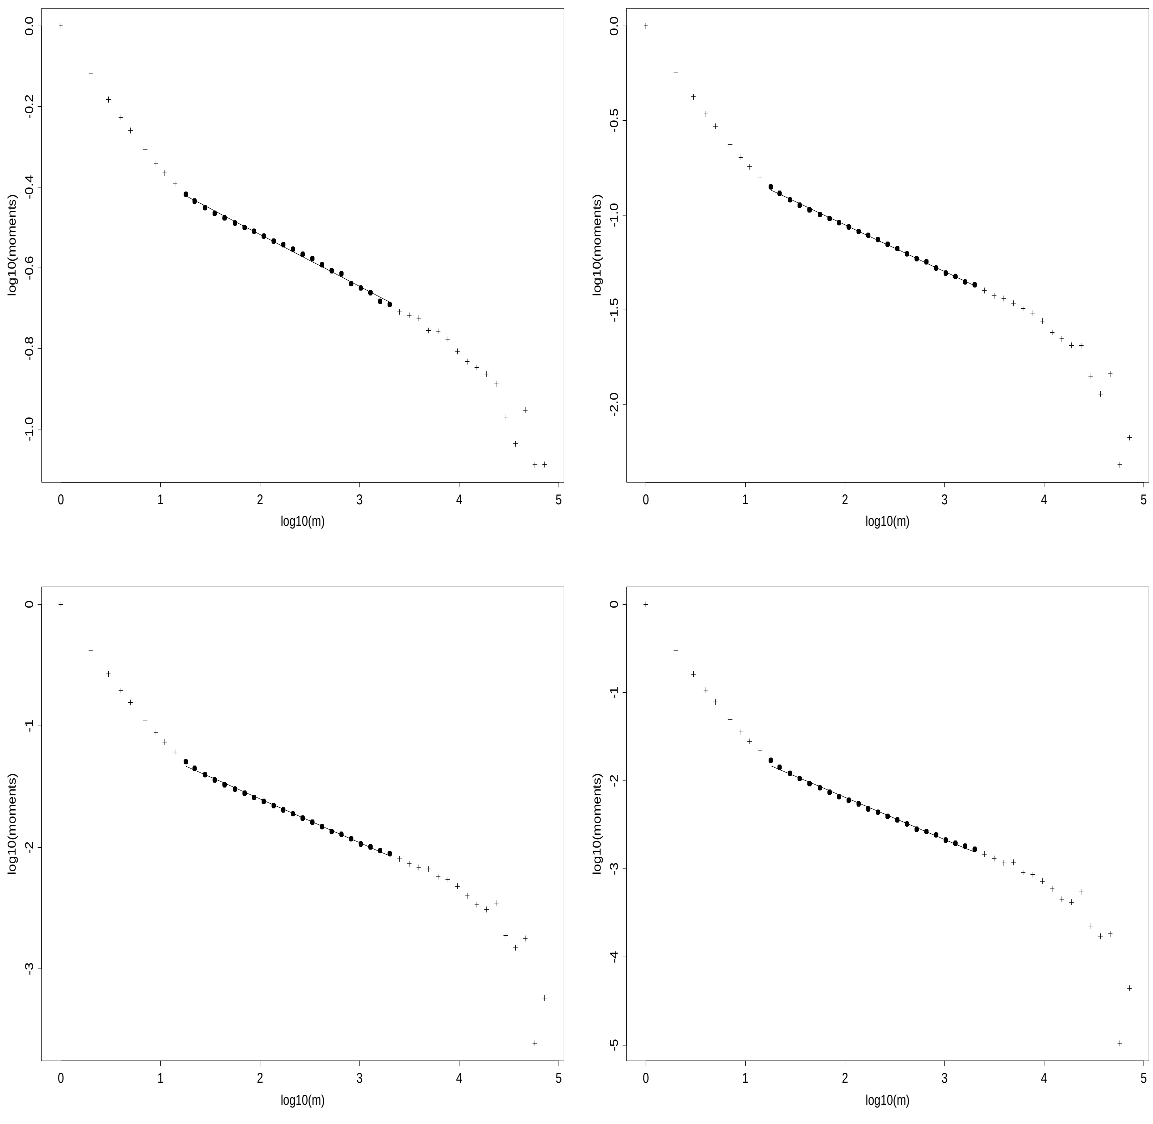

Figure 5: Plots of $\log_{10} \widehat{\mu}_0^{(m)}(q)$ versus $\log_{10} m$ for the December '94 Ethernet *packet* series with mean subtracted. The values of q and the estimated slope $\widehat{H}$ is as follows. Going from left to right: (a) $q = 1$, $\widehat{H} = .87$; (b) $q = 2$, $\widehat{H} = .88$; (c) $q = 3$, $\widehat{H} = .88$; (d) $q = 4$, $\widehat{H} = .88$.

[9] M. S. Taqqu and V. Teverovsky. Robustness of Whittle-type estimates for time series with long-range dependence. Preprint, 1995.

[10] M. S. Taqqu and V. Teverovsky. Semi-parametric graphical estimation techniques for long-memory data. In *Proceedings of the Athens Conference on Applied Probability and Time Series Analysis*, New York, 1996. Springer-Verlag. Time series volume in honour of E. J. Hannan. To appear.

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

[12] V. Teverovsky and M. S. Taqqu. Testing for long-range dependence in the presence of shifting means or a slowly declining trend using a variance-type estimator. Preprint, 1995.

[13] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level (Extended Version). Preprint. The paper also contains the mathematical proof that the superposition of strictly alternating ON/OFF sources converges to fractional Brownian motion, 1995.

Murad S. Taqqu and Vadim Teverovsky
Boston University
Department of Mathematics
111 Cummington Street
Boston, MA 02215-2411, USA
Email: murad@math.bu.edu, vt@math.bu.edu

Walter Willinger
Bellcore
Morristown, NJ 07960-6438, USA
Email: walter@bellcore.com