Fractional Noise in Experimental Measurements of IP Traffic in a Metropolitan Area Network

Stefano Bregni, *Senior Member, IEEE,* Walter Erangoli, *Student Member, IEEE*

Politecnico di Milano, Dept. of Electronics and Information, P.zza L. Da Vinci 32, 20133 Milano, ITALY
Tel.: +39-02-2399.3503 – Fax: +39-02-2399.3413 – E-mail: {bregni, erangoli}@elet.polimi.it

***Abstract* — Internet traffic exhibits self-similarity and long-range dependence (LRD) on various time scales. In this work, we analyze IP traffic traces measured on a MAN link of the telecommunications operator FASTWEB. We emphasize results obtained using the Modified Allan Variance (MAVAR), a well-known time-domain tool originally conceived to discriminate noise types with power-law spectrum (i.e., fractional noise) in frequency stability measurement, which has been introduced in network traffic analysis only recently. MAVAR has been demonstrated to feature superior spectral sensitivity and accuracy in fractional-noise parameter estimation, coupled with outstanding robustness against nonstationarity in data analyzed. Logscale diagrams are also evaluated, to confirm MAVAR experimental findings. All traffic traces analyzed exhibit prominent evidence of fractional noise. Model parameters are estimated. Results obtained are valuable to establish IP traffic models, e.g. necessary for realistic input traffic synthesis in network simulation.**

***Index Terms* — Fractional noise, Internet, metropolitan area network, long-range dependence, self-similarity, traffic measurement (communication).**

I. Introduction

Internet traffic exhibits intriguing temporal scale-invariance properties, such as self-similarity and long memory (long-range dependence) on various time scales [1]—[3]. Contrary to the classical Poisson-model assumption, these properties emphasize time-correlation between packet arrivals. Internet traffic traces exhibiting such behaviour include, but are not limited to, cumulative or incremental data count transmitted over time, inter-arrival time series of IP packets, etc.

The issue of estimating statistical parameters characterizing self-similarity and long-range dependence (LRD) has been often studied, aiming at best modelling of traffic for example to the purpose of network simulation. Several algorithms have been developed, in particular, to estimate the Hurst parameter H and the spectrum frequency power γ [1][2][4][5]. Among them, the logscale-diagram method is one of the best reputed.

In a different context, the Modified Allan Variance (MAVAR) is a well-known time-domain analysis tool originally conceived for time and frequency stability characterization of precision oscillators [6]—[10]. This variance was originally designed to discriminate effectively noise types with power-law spectrum (i.e., in broad terms, fractional noise), recognized very commonly in frequency sources. The use of MAVAR in network traffic analysis, for accurate estimation of H and γ of self-similar and LRD data series, was introduced recently in papers [11][12]. MAVAR has been demonstrated to feature superior spectral sensitivity and accuracy in fractional-noise parameter estimation, coupled with outstanding robustness against possible nonstationarity in data analyzed.

In this work, we analyze IP traffic traces (bytes and packets per time unit) measured on a MAN link of the telecommunications operator FASTWEB, Milano, Italy. While we emphasize results obtained using the Modified Allan Variance (MAVAR), due to its numerous advantages, we also computed logscale diagrams to confirm experimental findings. An interesting snapshot of the time-correlation that may be exhibited by IP traffic in a real network is thus provided.

All traffic traces analyzed exhibit prominent evidence of fractional noise, confirming once again the appropriateness of the power-law spectrum model for Internet traffic. Results obtained in this study are valuable, for example, to the purpose of traffic synthesis in network simulation.

II. Self-Similarity and Long-Range Dependence

A random process $X(t)$ (say, cumulative packet arrivals in the time interval $[0, t]$) , with $t \in \Re$, is said to be *self-similar*, with scaling parameter of self-similarity or Hurst parameter $H>0$, if

$$X(t) =_d a^{-H} X(at) \qquad (1)$$

for all $a>0$, where $=_d$ denotes equality of all finite-dimensional distributions [1][2].

Long-range dependence (LRD) of a process is defined by an asymptotic power-law decrease of its autocovariance or equivalently PSD functions [1][2]. Let $Y(t)$, with $t \in \Re$, be a second-order stationary stochastic process. The process $Y(t)$ is long-range dependent if its autocovariance function follows

$$R_Y(\delta) \sim c_1 |\delta|^{\gamma-1} \quad \text{for } \delta \to +\infty,\ 0<\gamma<1 \qquad (2)$$

or, equivalently, its power spectral density (PSD) follows

$$S_Y(f) \sim c_2 |f|^{-\gamma} \quad \text{for } f \to 0,\ 0<\gamma<1 \qquad (3).$$

It can be proved [2] that a self-similar process $X(t)$ with stationary increments and $1/2 < H < 1$ has long-range dependent increments $Y(t)$ (say, packet arrivals in the last time unit), with

$$\gamma = 2H - 1 \qquad (4).$$

Strictly speaking, the Hurst parameter characterizes self-similar processes, but it is frequently used to label also long-

Work partialy supported by Ministero dell'Istruzione, dell'Università e della Ricerca (MIUR), Italy, under PRIN project WONDER.

0-7803-9415-1/05/$20.00 (C) 2005 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2005 proceedings.

range dependent increment processes. Hence, expressions like "Hurst parameter of a LRD process" (characterized by the parameter γ) denote actually, by extension, the Hurst parameter $H = (\gamma+1)/2$ of its integral parent process.

III. The Modified Allan Variance

In time and frequency measurement theory, a well-known tool in the time domain for stability characterization of precision oscillators is the Modified Allan Variance (MAVAR) [6]—[10]. It was proposed in 1981 by modifying the definition of the two-sample variance (a.k.a. Allan variance, AVAR) recommended by IEEE in 1971 for characterization of frequency stability [13], after the pioneering work of D. W. Allan in 1966 [14], to improve its poor discrimination capability against white and flicker phase noise. This section briefly recalls MAVAR properties most relevant to our aim. For more details, the interested reader is referred to [6]—[10] and [11][12].

A. Definition in the Time Domain

Given an infinite sequence $\{x_k\}$ of samples of an input signal $x(t)$, evenly spaced in time with sampling period τ_0, the MAVAR is defined as

$$\mathrm{Mod}\,\sigma_y^2(\tau) = \frac{1}{2n^2\tau_0^2}\left\langle \left[\frac{1}{n}\sum_{j=1}^{n}\left(x_{j+2n} - 2x_{j+n} + x_j\right) \right]^2 \right\rangle \quad (5)$$

where the observation interval is $\tau = n\tau_0$ and the operator $<\cdot>$ denotes infinite-time averaging.

In time and frequency stability characterization, the data sequence $\{x_k\}$ is made of samples of random time deviation $x(t)$ of the chronosignal under test [6][13][15]. The MAVAR is thus a kind of variance of the second difference of $\{x_k\}$ or of the first difference of samples $\{y_k\}$ of the fractional frequency $y(t)=x'(t)$. In very brief, it differs from the basic Allan variance in the additional average over n adjacent measurements: for n=1 ($\tau=\tau_0$), the two variances coincide.

In practical measurements, given a finite set of N samples x_k, spaced by sampling period τ_0, a simple estimate of MAVAR can be computed using the ITU-T standard estimator [6][16]

$$\mathrm{Mod}\,\sigma_y^2(n\tau_0) = \frac{\displaystyle\sum_{j=1}^{N-3n+1}\left[\sum_{i=j}^{n+j-1}\left(x_{i+2n} - 2x_{i+n} + x_i\right)\right]^2}{2n^4\tau_0^2(N-3n+1)} \quad (6)$$

with n=1, 2,..., $\lfloor N/3 \rfloor$. A recursive algorithm for fast computation of this estimator exists [6], which cuts down the number of operations needed for all values of n to $\sim N^2$ instead of $\sim N^3$.

It should be noted that the point estimate (6), computed by averaging N-3n+1 terms, is a random variable itself. Exact computation of confidence intervals is not immediate and, annoyingly enough, depends on the spectrum of the underlying noise [17]—[22]. However, in general, along a plot of MAVAR versus τ, confidence intervals are negligible at left (short τ) and widen moving to right (long τ), where fewer terms are averaged. In our measurement results (Sec. V), therefore, we avoided to consider MAVAR values computed for largest n (right plots), where uncertainty is not negligible.

B. Equivalent Definition in the Frequency Domain

The MAVAR time-domain definition (5) can be translated to an equivalent expression in the frequency domain, allowing a more profound understanding of the behaviour of this quantity. In fact, definition (5) can be rewritten as the mean-square value of the signal output by a properly shaped filter receiving the data sequence $\{x_k\}$. Hence, the MAVAR can be equivalently defined as [6][15]

$$\mathrm{Mod}\,\sigma_y^2(\tau) = \int_0^\infty S_x(f)(2\pi f)^2 \frac{2\sin^6 \pi\tau f}{(n\pi\tau f)^2 \sin^2 \pi\frac{\tau}{n}f}\,df \quad (7)$$

where $S_x(f)$ is the one-sided PSD of the input signal $x(t)$.

Interesting enough is to notice that the function multiplying $S_x(f)(2\pi f)^2$ under integration is pass-band, having magnitude shaped with a narrow main lobe centred at $f \cong 1/(3\tau)$. Thus, MAVAR(τ) gathers signal power selectively from this narrow band: high-resolution spectral analysis of $x(t)$ can be achieved by computing MAVAR over a range of τ [6][15].

IV. Using the Modified Allan Variance to Estimate Parameters of Fractional Noise

It is convenient to generalize the LRD power-law model of spectral density (3). We will deal with random processes $x(t)$ whose one-sided PSD is modelled as

$$S_x(f) = \begin{cases} \displaystyle\sum_{i=1}^{P} h_{\alpha_i} f^{\alpha_i} & 0 < f \le f_\mathrm{h} \\ 0 & f > f_\mathrm{h} \end{cases} \quad (8)$$

where P is the number of power-law noise types considered in the model, α_i and h_{α_i} are parameters (α_i, $h_{\alpha_i} \in \Re$) and f_h is the unavoidable upper cut-off frequency. Such random processes are generically referred to as *fractional noise*, in broad sense.

Power-law noise with $-4 \le \alpha_i \le 0$ has been revealed in practical measurements of various physical phenomena, such as phase noise of precision oscillators [6]—[10][13]—[15][20] and Internet traffic [1][2][11][12], whereas P should be not greater than few units for the model being useful. If the process $x(t)$ is simple LRD (3), then P=1 and $-1 < \alpha_i < 0$. Although values $\alpha_i \le -1$ yield model pathologies, such as infinite variance and even nonstationarity, this model is commonly used, considering also that measurements have finite duration.

Under this general hypothesis of power-law PSD, first we notice that MAVAR convergence (7) is ensured for $\alpha_i > -5$. Then, by considering separately each term of the sum in (8) and letting $P = 1$, $\alpha = \alpha_i$, evaluation of (7) yields corresponding time-domain expressions of MAVAR. In short, for any value in the whole range of convergence $-5 < \alpha \le 0$, MAVAR is found to obey (ideally asymptotically for $n\to\infty$, keeping constant $n\tau_0 = \tau$, but in practice for n>5) to a simple power law of the observation time τ, i.e.

$$\mathrm{Mod}\,\sigma_y^2(\tau) \sim A_\mu \tau^\mu \tag{9}$$

where $\mu = -3-\alpha$ [6][10]. If $P>1$, it is immediate to generalize (9) in summation of powers $A_{\mu_i}\tau^{\mu_i}$.

This is a fundamental result. If $x(t)$ obeys (8), then a log-log plot of MAVAR looks ideally as a broken line made of P straight segments, whose slopes μ_i can be measured to get the exponent estimates $\alpha_i = -3-\mu_i$ of the power-law noise components prevailing in distinct ranges of τ.

If we consider a LRD process with PSD (3) characterized by Hurst parameter $1/2 < H < 1$, from (4) and (9) we obtain

$$H = \mu/2 + 2 \tag{10}.$$

In papers [11][12], these estimates of H and α_i were demonstrated to be very accurate, even better than those by the logscale diagram method. Moreover, nonstationary components of various kinds in the analyzed sequence affect MAVAR negligibly or in a well recognizable way. Hence, we adopted MAVAR as main tool to analyze IP traffic traces.

V. Methodology of IP Traffic Measurement

We analyzed an IP traffic trace acquired in the metropolitan network of the telecommunications operator FASTWEB. Traffic measurement was carried out on a Gigabit Ethernet link, for T=24 hours starting at 1.00pm of 25 Sep 2003. Only unicast traffic was measured. Raw traces were acquired using *tcpdump* on a Linux PC, observing the first 68 bytes of each IP packet. *Tcpdump* captured TCP, UDP, VoIP, ICMP and other IP packet headers. From each header, we extracted IP packet length and timestamp with 1-µs resolution. The size of the overall traffic dump resulted about 100 GB.

Then, we processed this raw sequence of packet length and timestamp data, split in 24 consecutive segments, one per each hour of the measurement day. By gathering data over consecutive intervals τ_0=10 ms (time unit), we produced 24×2 sequences of N=360000 samples of IP *packets per time unit* (pkts/t.u.) and *bytes per time unit* (bytes/t.u.). Finally, we applied our analysis tools.

VI. Results of IP Traffic Trace Analysis

In this section, we show a small selection of experimental results obtained. Instead of MAVAR, we plotted its square root Modified Allan Deviation (MADEV). Moreover, we computed logscale diagrams (LD), using scripts available at [23] (Daubechies' wavelet with three vanishing moments), where vertical bars represent 95%-confidence intervals.

To provide a visual sketch of traffic trend, Figs. 1 and 2 show two sample bytes/t.u. and packets/t.u. sequences, each made of N=360000 samples acquired over T=1 hour, starting at 0.00am, with time unit τ_0=10 ms. By simple eye inspection, it is evident that the data sequences are not white and it is possible to note fractional noise. At $t\cong$1500 s, a step-like change in measured data was captured; however, MAVAR is not affected significantly by such a little nonstationarity [12].

Fig. 3 depicts the normalized histogram (i.e., estimate of probability distribution) of packet size values measured over the whole test period of 24 hours. It is interesting to note that, while almost all bin values from 28 bytes to 1500 bytes are not empty, few values gathered the vast majority of all samples. In particular, most common IP packet length values resulted: 40 bytes (13%), 48 bytes (2.6%), 92 bytes (13%), 200 bytes (2%), 576 bytes (3.4%), 1216 bytes (9%), 1500 bytes (32%).

Figs. 4, 5, 6 and 7 depict MADEV and LD computed on the 24-hours sequences bytes/t.u. and pkts/t.u. (N=8640000, τ_0=10 ms). While the trend of LDs is rather irregular and is thus difficult to interpret, MAVAR allows a more precise and immediate spectral characterization of the traffic sequence (cf. results in [11][12]). By best fitting, it is possible to approximate MAVAR curves very well, in log-log scale, with few straight lines. For the bytes/t.u. sequence, their main slopes are μ=-2.63 (τ<2 s) and μ=-1.67 (τ>4 s). Thus, the sequence is revealed to be almost exactly sum of two simple components having power-law spectrum (8), with α=-0.37 dominant for τ<2 s and α=-1.33 dominant for τ>4 s. The former noise (α=-0.37) is kind of LRD, with H=0.69.

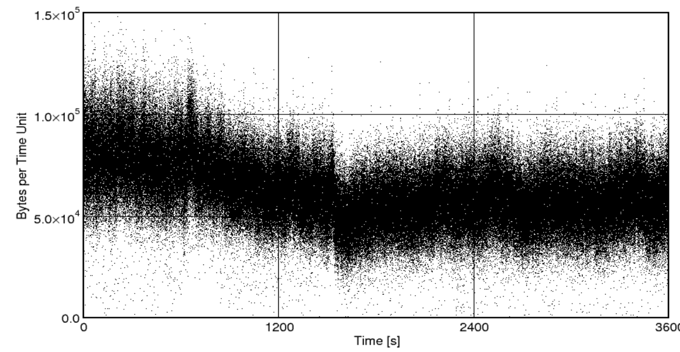

Fig. 1: Sample sequence bytes/t.u. (0.00am, N=360000, τ_0=10 ms, T=3600 s).

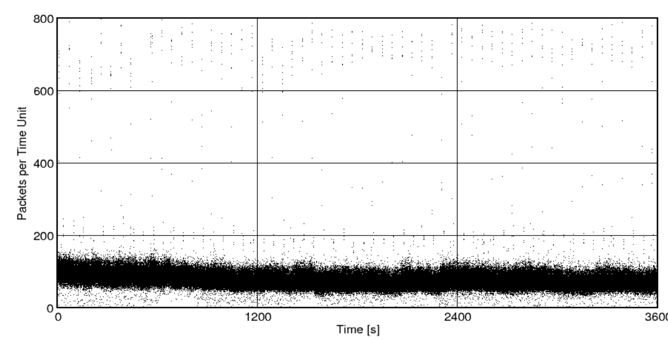

Fig. 2: Sample sequence pkts/t.u. (0.00am, N=360000, τ_0=10 ms, T=3600 s).

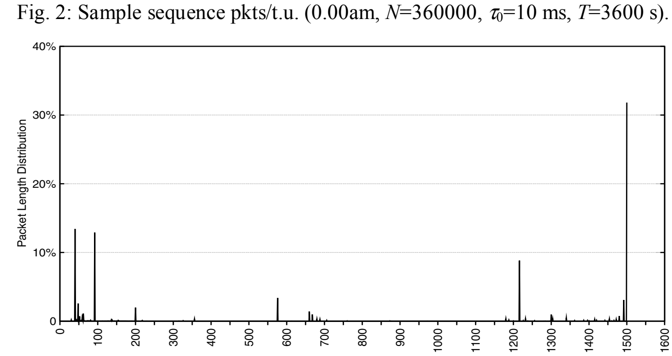

Fig. 3: Normalized histogram of packet size values measured (T=24h).

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2005 proceedings.

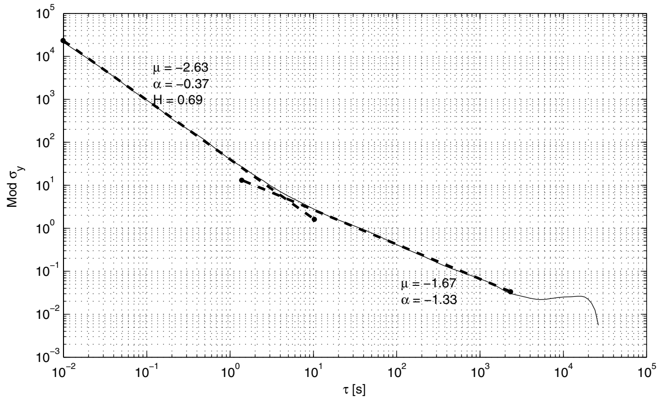

Fig. 4: MADEV of sequence bytes/t.u. (T=24h, N=8640000, τ_0=10 ms).

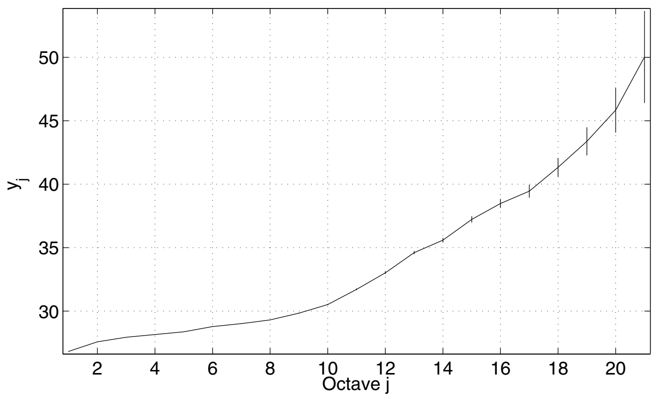

Fig. 5: LD of sequence bytes/t.u. (T=24h, N=8640000, τ_0=10 ms).

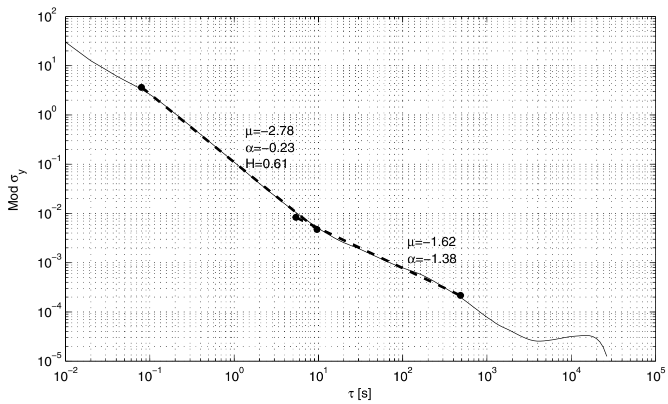

Fig. 6: MADEV of sequence pkts/t.u. (T=24h, N=8640000, τ_0=10 ms).

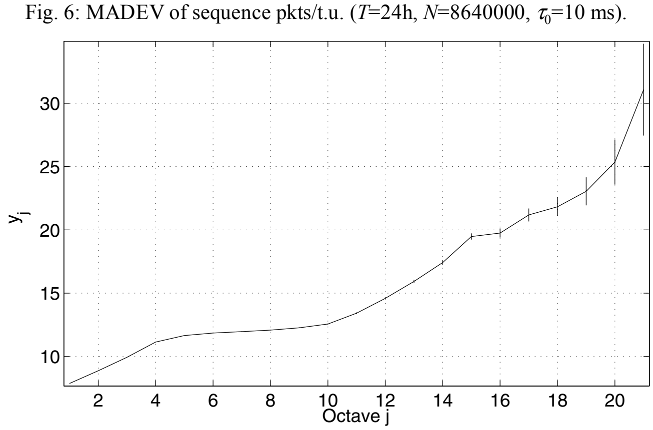

Fig. 7: LD of sequence pkts/t.u. (T=24h, N=8640000, τ_0=10 ms).

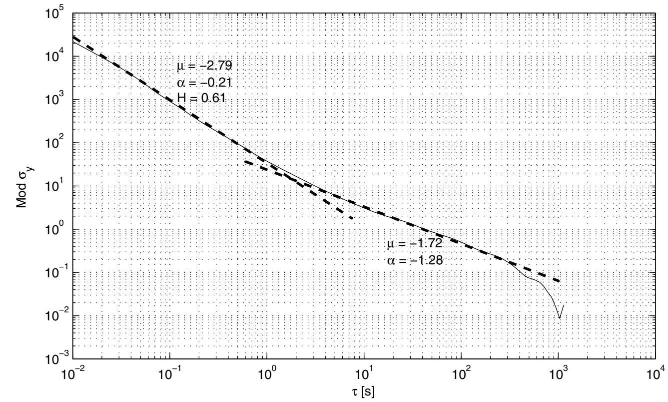

Fig. 8: MADEV of sequence *bytes/t.u.* (4.00pm, T=1h, N=360000, τ_0=10 ms).

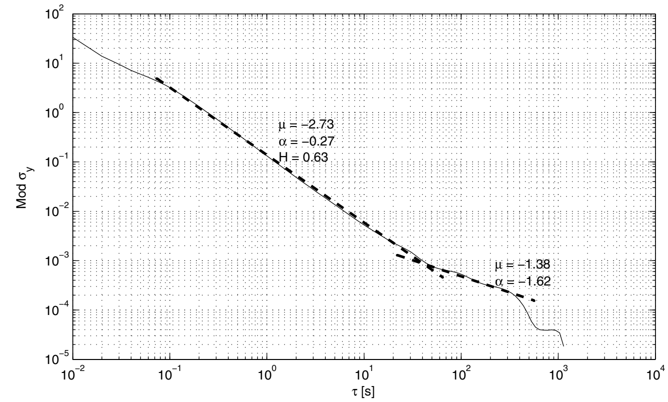

Fig. 9: MADEV of sequence *pkts/t.u.* (4.00pm, T=1h, N=360000, τ_0=10 ms).

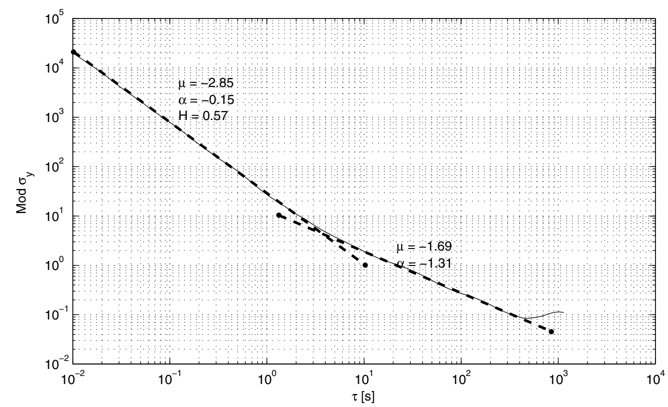

Fig. 10: MADEV of sequence *bytes/t.u.* (0.00am, T=1h, N=360000, τ_0=10ms).

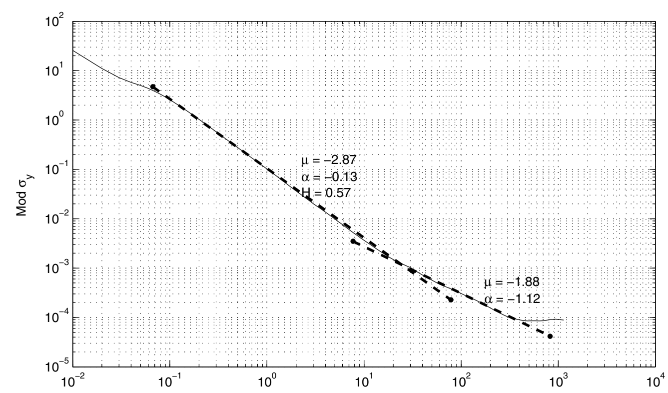

Fig. 11: MADEV of sequence *pkts/t.u.* (0.00am, T=1h, N=360000, τ_0=10 ms).

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2005 proceedings.

Table 1: Summary of values of α, μ and H estimated by MADEV diagrams.

μ	α	H	τ	T	Trace	Ref.
-2.63 -1.67	-0.37 -1.33	0.69 -	$\tau < 2$ s $\tau > 4$ s	24 h	bytes/t.u.	Fig. 4
-2.78 -1.62	-0.23 -1.38	0.61 -	100 ms < $\tau < 10$ s $\tau > 10$ s	24 h	packets/t.u.	Fig. 6
-2.79 -1.72	-0.21 -1.28	0.61 -	$\tau < 2$ s $\tau > 2$ s	1 h	bytes/t.u.	Fig. 8
-2.73 -1.38	-0.27 -1.62	0.63 -	100 ms < $\tau < 30$ s $\tau > 30$ s	1 h	packets/t.u.	Fig. 9
-2.85 -1.69	-0.15 -1.31	0.57 -	$\tau < 2$ s $\tau > 2$ s	1 h	bytes/t.u.	Fig. 10
-2.87 -1.88	-0.13 -1.12	0.57 -	100 ms < $\tau < 20$ s $\tau > 20$ s	1 h	packets/t.u.	Fig. 11

Somewhat different results are given by pkts/t.u. data. Omitting to consider the slope where n is too little to let us draw a reliable estimate of α ((9) is derived asymptotically for $n \to \infty$), two main slopes are identifiable: μ=-2.80 (α=-0.20, H=0.60) for 100 ms < $\tau < 10$ s and μ=-1.58 (α=-1.42) for $\tau > 10$ s.

Figs. 8, 9, 10 and 11 depict MADEV diagrams computed on the 1-hour sequences bytes/t.u. and pkts/t.u. acquired starting at 4.00pm and 0.00am (N=360000, τ_0=10 ms). Main slopes and estimated values of α and H are written on graphs.

Finally, all values of α, μ and H, estimated from MADEV diagrams shown, are summarized in Table 1. Some conclusions can be drawn from these experimental data. Fractional noise with spectrum kf^{α}, with α mostly in range -0.2 to -1.5, was revealed in all traces, further confirming that memoryless Poisson assumption is not adequate to model Internet traffic. Moreover, all traces exhibit noise spectrum that can be well approximated as simple summation of two power-law terms:

$$S_x(f) = k_1 f^{\alpha_1} + k_2 f^{\alpha_2} \tag{11}$$

where the relative weight of factors k_1 and k_2 can be determined based on the intersection points of straight segments approximating MADEV curves.

In all bytes/t.u. traces, we notice that:

- the first term is dominant for $\tau < 2$ s: $-1 \le \alpha \le -0$ ($H \cong 0.6$);
- the second term is dominant for $\tau > 2$ s and is a longer-memory fractional noise: $-2 \le \alpha \le -1$.

Pkts/t.u. traces exhibit similar behavior, except in the very short term ($\tau < 100$ ms), where however n is too little to let us draw reliable estimates, and in the border between the two noise components ($\tau \cong 10^1$ s vs. $\tau \cong 10^0$ s).

VII. Conclusions

We analyzed IP traffic traces measured on a MAN link of the operator FASTWEB. We emphasized results obtained using MAVAR, because of its superior spectral sensitivity and accuracy in fractional-noise parameter estimation.

All traffic traces analyzed exhibit prominent evidence of fractional noise. The fine accuracy of MAVAR spectral analysis allowed revealing, in all IP traces analyzed, that noise spectrum can be well approximated as simple sum of two power-law terms $k_1 f^{\alpha_1} + k_2 f^{\alpha_2}$, where $-1 \le \alpha_1 \le -0$ and $-2 \le \alpha_2 \le -1$.

Our tests further confirm that memoryless Poisson assumption is not adequate to model Internet traffic and provide valuable results to the purpose of realistic input traffic synthesis in network simulation.

Acknowledgements

We wish to thank FASTWEB, for allowing us to make traffic measurements and for their effective support.

References

[1] P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, D. Veitch, "The Multiscale Nature of Network Traffic", *IEEE Signal Processing Mag.*, 19(3), pp. 28-46, May 2002.
[2] K. Park, W. Willinger, "Self-Similar Network Traffic: An Overview", P. Abry, P. Flandrin, M. S. Taqqu, D. Veitch, "Wavelets for the Analysis, Estimation, and Synthesis of Scaling Data", in *Self-Similar Network Traffic and Performance Evaluation*, K. Park, W. Willinger, Eds. Chichester, UK: John Wiley & Sons, 2000, pp. 1-88.
[3] V. Paxson, S. Floyd, "Wide-Area Traffic: the Failure of Poisson Modeling", *IEEE/ACM Trans. Networking,* 3(6), pp. 226-244, June 1995.
[4] P. Abry, D. Veitch, "Wavelet Analysis of Long-Range Dependent Traffic", *IEEE Trans. Inform. Theory*, 44(1), pp. 2-15, Jan. 1998.
[5] M. S. Taqqu, V. Teverovsky, W. Willinger, "Estimators for Long-Range Dependence: an Empirical Study", *Fractals*, 3(4), pp. 785-798, 1995.
[6] S. Bregni, *Synchronization of Digital Telecommunications Networks*. "Chapter 5 - Characterization and Modelling of Clocks". Chichester, UK: John Wiley & Sons, 2002, pp. 203-281.
[7] D. W. Allan, J. A. Barnes, "A Modified Allan Variance with Increased Oscillator Characterization Ability", *Proc. 35th Annual Frequency Control Symposium*, 1981.
[8] P. Lesage, T. Ayi, "Characterization of Frequency Stability: Analysis of the Modified Allan Variance and Properties of Its Estimate", *IEEE Trans. Instrum. Meas.*, 33(4), pp. 332-336, Dec. 1984.
[9] L. G. Bernier, "Theoretical Analysis of the Modified Allan Variance", *Proc. 41st Annual Frequency Control Symposium*, 1987.
[10] D. B. Sullivan, D. W. Allan, D. A. Howe, F. L. Walls, Eds., *"Characterization of Clocks and Oscillators"*, NIST Tech. Note 1337, March 1990.
[11] S. Bregni, L. Primerano, "The Modified Allan Variance as Time-Domain Analysis Tool for Estimating the Hurst Parameter of Long-Range Dependent Traffic", *Proc. IEEE GLOBECOM2004*, Dallas, USA, 2004.
[12] S. Bregni, L. Primerano, "Using the Modified Allan Variance for Accurate Estimation of the Hurst Parameter of Long-Range Dependent Traffic". Submitted to *IEEE Trans. Inform. Theory*, Feb. 2005.
[13] J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healey, D. B. Leeson, T. E. McGunigal, J. A. Mullen Jr., W. L. Smith, R. L. Sydnor, R. F. C. Vessot and G. M. R. Winkler, "Characterization of Frequency Stability," *IEEE Trans. on Instr. and Meas.*, 20(2), May 1971.
[14] D. W. Allan, "Statistics of Atomic Frequency Standards", *Proceedings of the IEEE*, 54(2), July 1966.
[15] J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress", *Proc. IEEE*, 66(9), pp. 1048-1075, Sept. 1978.
[16] ITU-T Rec. G.810 *"Definitions and Terminology for Synchronisation Networks"*, 1996-2003.
[17] P. Lesage, C. Audoin, "Characterization of Frequency Stability: Uncertainty Due to the Finite Number of Measurements", *IEEE Trans. Instrum. Meas.*, 22(2), pp. 157-161, June 1973.
[18] P. Lesage, C. Audoin, "Estimation of The Two-Sample Variance with Limited Number of Data", *Proc. 31st Annual Freq. Control Symp.*, 1977.
[19] K. Yoshimura, "Characterization of Frequency Stability: Uncertainty Due to the Autocorrelation Function of Frequency Fluctuations", *IEEE Trans. Instrum. Meas.*, 27(1), pp. 1-7, Mar. 1978.
[20] S. R. Stein, "Frequency and Time - Their Measurement and Characterization", in *Precision Frequency Control*, vol. 2, E. A. Gerber A. Ballato, Eds. New York: Academic Press, 1985, pp. 191-232.
[21] C. A. Greenhall, "Recipes for Degrees of Freedom of Frequency Stability Estimators", *IEEE Trans. Instrum. Meas.*, 40(6), pp. 994-999, Dec. 1991.
[22] C. A. Greenhall, W. J. Riley. "Uncertainty of Stability Variances Based on Finite Differences". Available: http://www.wriley.com.
[23] D. Veitch. Code for The Estimation of Scaling Exponents. Available: http://www.cubinlab.ee.mu.oz.au/~darryl/secondorder_code.html.