

On Traffic Long-Range Dependence at the Output of Schedulers with Multiple Service Classes

Stefano Bregni, *Senior Member, IEEE*, Paolo Giacomazzi, Gabriella Saddemi

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

Abstract — Long-range dependence (LRD) is a widely verified property of Internet traffic, which severely impacts network performance yielding longer queuing delays. Moreover, LRD is almost ubiquitous and very hard to remove or control. In this work, we investigated by thorough simulation the effect of schedulers on traffic LRD. We analyzed the output traffic of schedulers merging LRD flows according to various service policies, viz. plain FIFO, Static-Priority, Earliest-Deadline-First and General Processor Sharing. First, we noticed that traffic LRD is not affected much by crossing the scheduler, for any type of service policy. Second, we showed that LRD also propagates across different service classes with any policy except balanced GPS, which ensures complete separation between classes. This phenomenon may explain in part why LRD is so widespread in Internet traffic.

Index Terms — Communication system traffic, fractional noise, Internet, long-range dependence, queuing analysis, traffic control (communication).

I. INTRODUCTION

Internet traffic exhibits self-similarity (SS) and long-range dependence (LRD) [1][2]. These properties emphasize long-range time-correlation between packet arrivals. Fractional noise and fractional Brownian motion models are often used to describe such behaviour of Internet traffic series, e.g. cumulative or incremental bit count transmitted over time.

In a SS random process, a dilated portion of a realization, by the scaling Hurst parameter H , has same statistical characterization than the whole. On the other hand, LRD is a long-memory property, usually equated to an asymptotic power-law decrease of the power spectral density (PSD) as $\sim f^{-\alpha}$ (for $f \rightarrow 0$) or, equivalently, of the autocovariance function. Under some hypotheses [1], the integral of a LRD process is SS with H related to α (e.g., fractional Brownian motion, integral of fractional Gaussian noise).

It has been well recognized [3]–[6] that traffic LRD yields long queues in network buffers. In the case of fractional Gaussian traffic, for example, it has been found [3][4] that the queue tail is Weibull-distributed, i.e. the buffer occupancy Q exceeds a given threshold q with asymptotic probability $P\{Q > q\} \sim \exp(-\beta q^{1-\alpha})$, where β is a positive function of α and of other network parameters. The Weibull queue length distribution departs significantly from the plain exponential resulting in case of Poisson input ($\alpha = 0$). In particular, when α tends to 1, the Weibull distribution flattens and average and variance of the queuing delay even tend to infinite.

LRD is a property of network traffic that is almost ubiquitous and very hard to remove or control. In a previous work,

we showed by simulation that LRD cannot be reduced by standard leaky-bucket (LB) policers and shapers [7][8], unless by dropping a large fraction of traffic.

In this work, we investigated by simulation the effect of schedulers on traffic LRD. We studied the output traffic of schedulers merging LRD flows according to various service policies, viz. plain FIFO, Static-Priority, Earliest-Deadline-First and General Processor Sharing. We analyzed the output traffic in order to determine if, and under what conditions: 1) LRD input flows are still LRD at the scheduler output; 2) non-LRD flows merged with LRD flows are tainted by LRD.

We proceeded along the way marked by [9], where the persistence of LRD at the output of a single-server queue was proved analytically. In a subsequent work [10], it was also proved that in various queuing systems (viz., G/GI and GI/GI) the output of the system is LRD if the input is LRD. More recently, it was shown [11] that SS traffic flows merged by a simple scheduler retain same self-similarity.

In this paper, nevertheless, we extend such results, by focusing on multi-class scheduling devices with significant practical importance. In particular, by extensive simulation, we examine the interaction between merged LRD and non-LRD flows.

II. SELF-SIMILARITY AND LONG-RANGE DEPENDENCE

A random process $X(t)$ (e.g., cumulative packet arrivals in time interval $[0, t]$), is said to be *self-similar* (SS), with scaling parameter of self-similarity or Hurst parameter $H > 0, H \in \mathfrak{R}$, if

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

for any $a > 0$, where $=_d$ denotes equality for all finite-dimensional distributions [1]. In other terms, the statistical description of $X(t)$ does not change by *scaling* its amplitude by a^{-H} and its time by a . Most SS processes are not stationary.

The class of SS processes is usually restricted to that of *self-similar processes with stationary increments* (SSSI), which are “integral” of some stationary process. For example, consider the δ -increment process of $X(t)$, defined as $Y_\delta(t) = X(t) - X(t-\delta)$ (e.g., packet arrivals in the last δ time units). For a SSSI process $X(t)$, $Y_\delta(t)$ is stationary and $0 < H < 1$ [1].

Long-range dependence (LRD) of a process is defined by an asymptotic power-law decrease of its autocovariance and PSD [1]. Let $Y(t)$ be a 2nd-order stationary random process. $Y(t)$ exhibits LRD if its autocovariance follows asymptotically

$$R_Y(\delta) \sim c_1 |\tau|^{\alpha-1} \quad \text{for } \tau \rightarrow +\infty, 0 < \alpha < 1 \quad (2)$$

or, equivalently, its two-sided PSD follows asymptotically

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

In general, a random process with non-integer power-law PSD is also known as fractional (not necessarily Gaussian) noise. SSSI processes $X(t)$ with $1/2 < H < 1$ have LRD increments $Y(t)$, with [1]

$$\alpha = 2H - 1 \quad (4).$$

III. ESTIMATING THE H AND α PARAMETERS OF LRD SERIES USING THE MODIFIED ALLAN VARIANCE

For estimating precisely the H and α parameters of LRD traffic series, we used the Modified Allan Variance (MAVAR) [12]. Here, only the most relevant MAVAR properties to this aim are briefly recalled for ease of understanding.

MAVAR is a well-known time-domain quantity, conceived in 1981 for frequency stability characterization of precision oscillators [13]–[17] by modifying the definition of the original Allan Variance (AVAR) [18]. MAVAR has been proven to feature superior spectral sensitivity and accuracy in LRD parameter estimation, coupled with excellent robustness against data nonstationarity (e.g., drift and steps) [12].

Given a finite set of N samples $\{x_k\}$ of a signal $x(t)$, evenly spaced by sampling period τ_0 , MAVAR can be estimated using the ITU-T standard estimator [17][19]

$$\text{Mod } \sigma_y^2(\tau) = \frac{\sum_{j=1}^{N-3n+1} \left[\sum_{i=j}^{n+j-1} (x_{i+2n} - 2x_{i+n} + x_i) \right]^2}{2n^4 \tau_0^2 (N - 3n + 1)} \quad (5)$$

where the observation interval is $\tau = n\tau_0$ and $n = 1, 2, \dots, \lfloor N/3 \rfloor$.

MAVAR is a kind of variance of the 2nd difference of input data, including an internal average over n adjacent samples.

A recursive algorithm for fast computation of this estimator exists [17], which cuts down the number of operations needed to compute it for *all* values of n to $\sim N^2$ instead of $\sim N^3$, or to $\sim N \log N$ if MAVAR($n\tau_0$) is computed only for $n = 2^k$ ($k = 0, 1, \dots, \lfloor \log_2 N \rfloor$) as in the wavelet logscale diagram [1][12].

Fractional noise with one-sided PSD $\sim 1/f^\alpha$, for $0 \leq \alpha \leq 4$, has been revealed in practical measurements of various phenomena, such as phase noise of precision oscillators [17][18][20] and Internet traffic [1]. Although values $\alpha \geq 1$ yield model pathologies, such as infinite variance and nonstationarity, this model is common, considering also that real-world measurements have finite duration and bandwidth.

Under this power-law model of input data and in the whole range of MAVAR convergence $0 \leq \alpha < 5$, MAVAR itself is found following the simple power law (ideally asymptotically for $n \rightarrow \infty$, $n\tau_0 = \tau$, but in practice for $n > 4$)

$$\text{Mod } \sigma_y^2(\tau) \sim A_\mu \tau^\mu, \quad \mu = -3 + \alpha \quad (6).$$

Therefore, the slope μ of $\text{Mod } \sigma_y^2(\tau)$ in a log-log plot can be estimated (e.g., by linear regression) to yield the exponent $\alpha = 3 + \mu$ of the fractional noise that is dominant in input data. From (4), if $\alpha < 1$, we also obtain $H = \mu/2 + 1$. These estimates

of H and α were demonstrated to be very accurate, unbiased and robust against nonstationarity in data analyzed (viz. drifts, periodic trends and steps) [12].

This procedure is analogous to that of the 2nd-order logscale diagram technique based on wavelet analysis [1][21], which analyzes data over a range of scales by observing the power-law behavior of the wavelet detail variances across octaves.

IV. MODEL OF SCHEDULERS AND LRD TRAFFIC

In this paper, we study the output traffic of infinite-buffer *plain FIFO*, *Static-Priority (SP)*, *Earliest-Deadline-First (EDF)* and *General Processor Sharing (GPS)* schedulers merging LRD flows. Other types of schedulers might be considered as well. However, in this work, we have selected these due to their significant practical relevance.

In the *SP scheduler*, traffic flows are divided into service classes, numbered from 1 to n . The scheduler is provided with n queues, one per service class. All traffic flows in service class i ($1 \leq i \leq n$) feed the i -th queue. The queue index i determines the service priority: at the end of each packet transmission, the scheduler fetches next packet from the non-empty queue with smallest i .

In the *EDF scheduler*, traffic flows are alike divided into service classes, numbered from 1 to n . Any service class i is characterized by its service deadline d_i [seconds]. At the end of each packet transmission, the scheduler fetches, among the packets waiting for transmission, the packet with the smallest residual time. This is done by marking each packet on its arrival with the time stamp indicating its arrival time t_k . At time t , the residual time of this packet, belonging to service class i , is calculated as $t_k + d_i - t$, representing the amount of time left before the packet service deadline expires. Note that the packet residual time can be negative, indicating that the service deadline has already expired. Smaller residual time means more urgent need of service: this is the reason why the packet with the smallest residual time is selected for service.

In the *GPS scheduler*, each service class i is assigned a weight w_i (with no loss of generality, we assume $\sum_k w_k = 1$) and is guaranteed to receive at least a share $w_i / \sum_k w_k$ of the available capacity. If any class uses less than its share, the extra bandwidth is shared by all other classes proportionally to their weights. A formal description of GPS is given in [22].

For example, let us consider the case with $n = 2$ service classes. Note that GPS with $w_1 = 1, w_2 = 0$ and EDF with $d_1 = 0, d_2 = \infty$ are equivalent to SP, while EDF with $d_1 = d_2$ is equivalent to plain First-In-First-Out (FIFO). Note also, however, that GPS with $w_1 = w_2$ is *not* equivalent to FIFO.

As far as the LRD traffic model is concerned, in this work we focused on fractional Gaussian traffic (fGt), being this model commonly adopted in literature. Our procedure of traffic synthesis, detailed in [7], generates LRD pseudorandom series $\{x_k\}$ of fractional Gaussian traffic fGt(α, m_x, σ_x^2) with length N , PSD $\propto 1/f^\alpha$ for assigned values of α with $0 \leq \alpha < 1$, normally-distributed samples, mean m_x and variance σ_x^2 , rectified to replace negative samples with zero. The sequence $\{x_k\}$ represents the incremental data count [bit/s] input at each time unit into the scheduler under study (i.e., the input traffic rate).

V. SIMULATION RESULTS

In our simulations, we have studied FIFO, SP, EDF and GPS schedulers with $n = 2$ service classes. We denoted traffic flows of classes 1 and 2 as $x_1(t)$ and $x_2(t)$, respectively. We set the time unit $\tau_0 = 1$ ms, the mean of traffic rate $m_x = 2279$ bit per time unit (i.e., 2.279 Mbit/s) and its deviation $\sigma_x = 773.9$ bit per time unit (i.e., 773.9 kbit/s), as in [4]. We varied the α parameter of $x_1(t)$ and $x_2(t)$ liberally in range $0 \leq \alpha < 1$, denoting it as α_1 and α_2 , respectively. The scheduler buffer size is infinite. The capacity of the output line has been set $C = 4.85$ Mbit/s (i.e., $m_x/C \cong 0.47$).

In the SP scheduler, $x_1(t)$ has strict service priority over $x_2(t)$. In the EDF scheduler, the service deadlines have been set $d_1 = 50\tau_0$ and $d_2 = 100\tau_0$ (*unbalanced EDF*) and $d_1 = d_2$ (*balanced EDF*, i.e. FIFO). In the GPS scheduler, finally, weights have been set $w_1 = w_2 = 0.5$ (*balanced GPS*) and $w_1 = 2/3$, $w_2 = 1/3$ (*unbalanced GPS*).

Therefore, we generated input traffic fG_{TR} series $x_{1,IN}(t)$ and $x_{2,IN}(t)$ made of $N = 2^{23}$ samples, which were fed into schedulers as described before. Then, we analyzed the output traffic flows $x_{1,OUT}(t)$ and $x_{2,OUT}(t)$, corresponding to input flows $x_{1,IN}(t)$ and $x_{2,IN}(t)$, respectively.

Traffic series were analyzed in both time and frequency domains, respectively by MAVAR and classic FFT-based spectral estimation (periodogram over 8192 points, having divided the series in 1024 segments with Hamming data windowing [23]). Estimation of the LRD parameter α was performed by way of MAVAR, due to its unparalleled accuracy.

A. Sample Plots of PSD and MAVAR

Figs. 1, 3 and 5 show the PSD computed on sample traffic series at the input and output of the SP, unbalanced EDF and balanced GPS schedulers, respectively, for $\alpha_{1,IN} = 0.8$ and $\alpha_{2,IN} = 0$. Moreover, Figs. 2, 4 and 6 show the MAVAR computed on sample traffic $x_{2,OUT}(t)$ at the output of same schedulers, for $\alpha_{1,IN} = 0.0, 0.2, 0.4, 0.6, 0.8$ and $\alpha_{2,IN} = 0$.

1) *SP Scheduler*: $x_2(t)$ does not hamper the service of $x_1(t)$, since it has lower service priority. Thus, $x_1(t)$ sees the scheduler as a single-server system with the whole output capacity fully available. As expected, $S_{x_{1,OUT}}(f)$ is nearly identical to $S_{x_{1,IN}}(f)$: $x_1(t)$ is distorted negligibly by crossing the scheduler.

On the contrary, the flow $x_2(t)$ is heavily hampered by $x_1(t)$. This hindrance yields an interesting effect: as visible in Fig. 1, the flow $x_2(t)$ exhibits LRD at the output even if it is white at the input, due to interference from the higher-priority LRD flow $x_1(t)$. MAVAR results in Fig. 2 better show that the more $x_1(t)$ is LRD, the more $x_2(t)$ becomes LRD at output (the slope of MAVAR is lower): LRD propagates from class 1 to class 2.

2) *Unbalanced EDF Scheduler*: Unlike the hard priority of SP, EDF follows a softer approach. Being deadlines $d_1 < d_2$, $x_1(t)$ has soft precedence over $x_2(t)$: contrary to SP, a traffic unit of $x_2(t)$ may be served even if units of $x_1(t)$ are still waiting for service. This occurs when the residual time of a unit of $x_2(t)$ is small, because it has waited for service long.

This different behaviour is evident comparing Figs. 1 and 3: in the EDF scheduler, the PSD of both $x_1(t)$ and $x_2(t)$ at output differs from that at input. The soft-priority policy yields a mu-

tual interference between $x_1(t)$ and $x_2(t)$ and thus significant LRD exchange between traffic classes: the LRD of $x_1(t)$ is diminished by merging white $x_2(t)$, while $x_2(t)$ gets tainted by some LRD from $x_1(t)$ (but less than in the SP scheduler, as visible also comparing slopes of MAVAR in Figs. 2 and 4).

3) *Balanced GPS Scheduler*: Here, the two flows share fairly the output capacity. Their interaction, in terms of LRD exchange, is minimal: both signals do not exhibit a change of the PSD slope at low frequencies crossing the scheduler (Fig. 5), while MAVAR of output traffic $x_{2,OUT}(t)$ is not affected by the LRD of $x_{1,IN}(t)$ (Fig. 6).

B. Characterization of Traffic LRD at Scheduler Output

Thorough investigation is needed for characterizing the behaviour outlined qualitatively in the previous section.

Therefore, for each of the schedulers considered, we varied liberally the LRD parameters of input traffic flows $x_1(t)$ and $x_2(t)$ in ranges $0 \leq \alpha_{1,IN} < 1$ and $0 \leq \alpha_{2,IN} < 1$. After each simulation, we estimated the LRD parameters $\alpha_{1,OUT}$ and $\alpha_{2,OUT}$ of output flows by way of MAVAR (see Sec. III). The average slope of MAVAR curves was estimated excluding the first and last decade of τ values [12] and neglecting, for the sake of simplicity, slight changes of slope as those visible in Figs. 2, 4 and 6 for $\alpha_{1,IN} = 0.4, 0.6$ (this is equivalent to approximating a two-terms power-law PSD $h_1/f^{\alpha_1} + h_2/f^{\alpha_2}$ as $\sim h_3/f^{\alpha_3}$ with $\alpha_1 < \alpha_3 < \alpha_2$ [12]). For each parameter setting, finally, we ran 10 independent simulations, to obtain also grand averages and confidence intervals of $\alpha_{1,OUT}$ and $\alpha_{2,OUT}$ estimates.

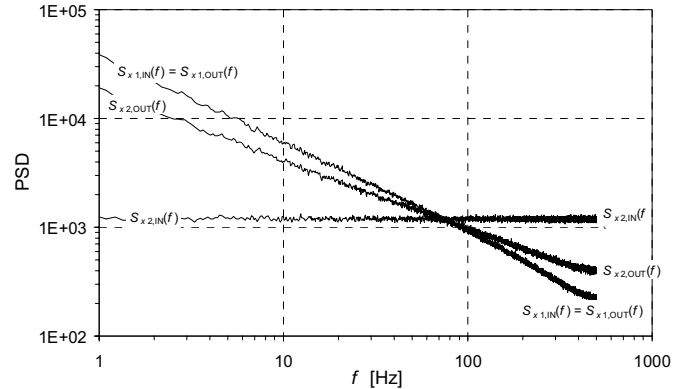


Fig. 1: PSD of input and output traffic (SP, $\alpha_{1,IN} = 0.8$, $\alpha_{2,IN} = 0.0$).

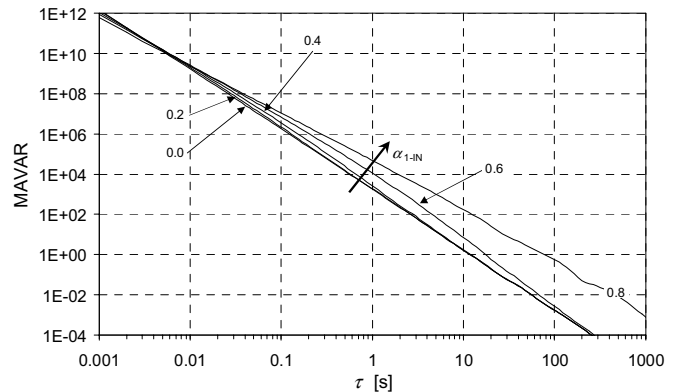


Fig. 2: MAVAR of output traffic $x_{2,OUT}(t)$ for various $\alpha_{1,IN}$ (SP, $\alpha_{2,IN} = 0.0$).

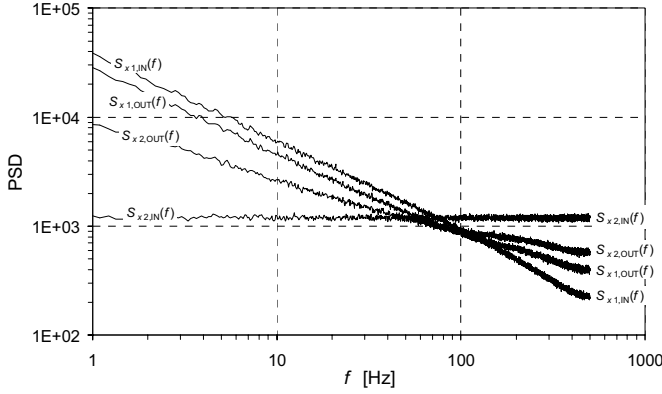


Fig. 3: PSD of input and output traffic (EDF, $\alpha_{1,IN} = 0.8$, $\alpha_{2,IN} = 0.0$).

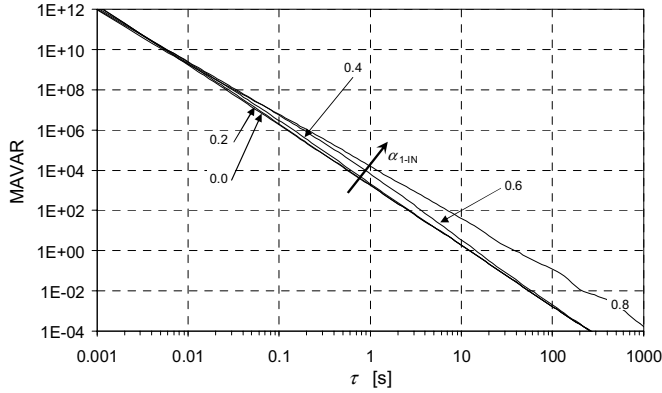


Fig. 4: MAVAR of output traffic $x_{2,OUT}(t)$ for various $\alpha_{1,IN}$ (EDF, $\alpha_{2,IN} = 0.0$).

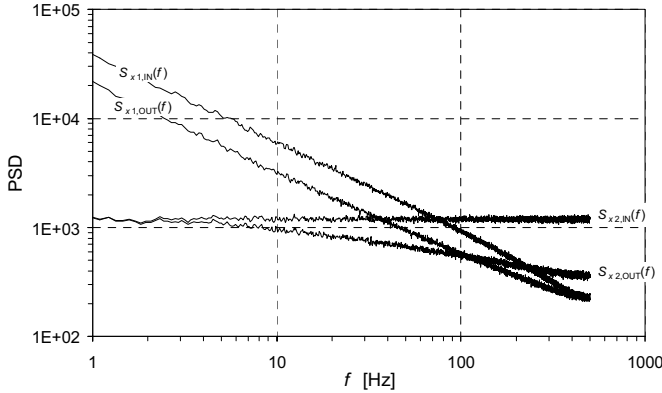


Fig. 5: PSD of input and output traffic (GPS, $\alpha_{1,IN} = 0.8$, $\alpha_{2,IN} = 0.0$).

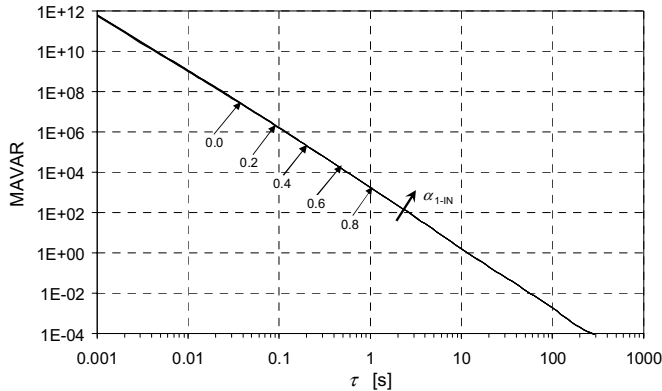


Fig. 6: MAVAR of output traffic $x_{2,OUT}(t)$ for various $\alpha_{1,IN}$ (GPS, $\alpha_{2,IN} = 0.0$).

Figs. 7 through 10 give a synoptic view of our simulation results, in the form of estimates of $\alpha_{i,OUT}$ as a function of $\alpha_{i,IN}$. Bars of 95%-confidence intervals $\alpha_{i,OUT} \pm \delta$ are not marked, to not make plots confused. However, in most cases they are negligible (almost always, $\delta < 0.01$).

From Figs. 7 and 10, we notice that the LRD of both flows $x_1(t)$ and $x_2(t)$ is not affected much by crossing the scheduler, for any type of service policy except the balanced GPS, which slightly increases it. From Figs. 8 and 9, instead, we notice significant transfer of LRD between flows, with all service policies except balanced GPS, confirming preliminary considerations made in Sec. V.A. In summary:

- *SP, unbalanced GPS*: the lower-priority flow $x_2(t)$ is tainted by LRD from $x_1(t)$ when $\alpha_{1,IN} > 0.5$; conversely, the higher-priority flow $x_1(t)$ is not affected by the LRD of $x_2(t)$;
- *unbalanced and balanced EDF (FIFO)*: $x_1(t)$ is tainted by $x_2(t)$ LRD when $\alpha_{2,IN} > 0.5$ and *vice versa* when $\alpha_{1,IN} > 0.5$;
- *balanced GPS*: there is no LRD exchange between flows.

VI. CONCLUSIONS

In this paper, we investigated by thorough simulation the effect of traffic schedulers on the $1/f^\alpha$ power-law spectrum of merged LRD flows. MAVAR analysis allowed to estimate the LRD of flows at the output of schedulers with better accuracy than in any previous simulation study.

First, we noticed that the LRD of traffic flows is not affected much by crossing the scheduler, for any type of service policy except balanced GPS, which slightly increases it. Second, we demonstrated that LRD also propagates across different service classes with FIFO, SP, unbalanced EDF and unbalanced GPS. Only balanced GPS ensures complete separation between classes. The detailed results presented in Sec. V.B provide considerable insight into this phenomenon.

In summary, we showed that high-LRD traffic flows in many cases taint other non-LRD flows that share the output link capacity of schedulers with them. The actual amount of this LRD propagation depends, among other factors, on the scheduler service policy. This phenomenon may explain, in part, why LRD is so ubiquitous and widespread in Internet traffic.

ACKNOWLEDGEMENT

The Authors wish to thank Roberto Cioffi, Politecnico di Milano, for his kind help on algorithms and numerical computation of some results presented in this paper.

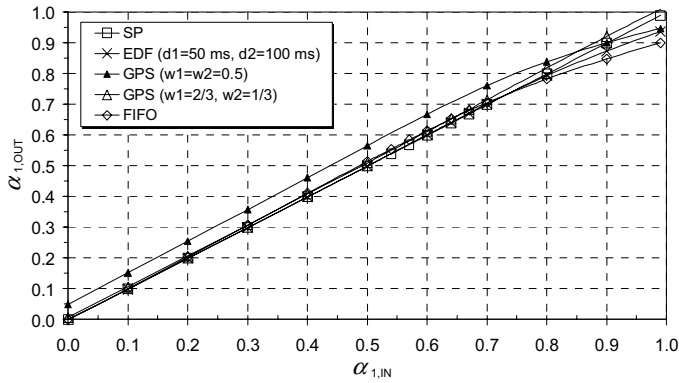


Fig. 7: Estimates of LRD parameter $\alpha_{1,OUT}$ as a function of $\alpha_{1,IN}$ ($\alpha_{2,IN} = 0.0$).

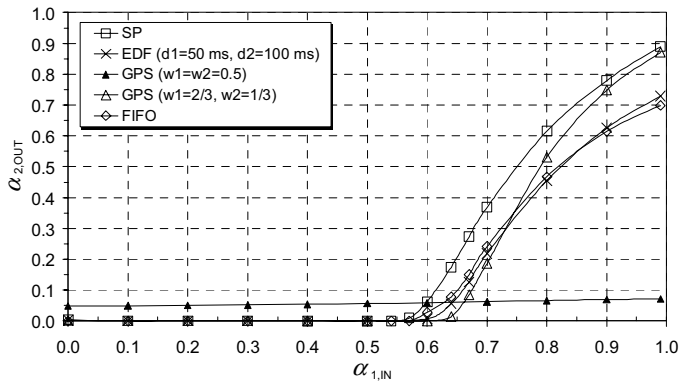


Fig. 8: Estimates of LRD parameter $\alpha_{2,OUT}$ as a function of $\alpha_{1,IN}$ ($\alpha_{2,IN} = 0.0$).

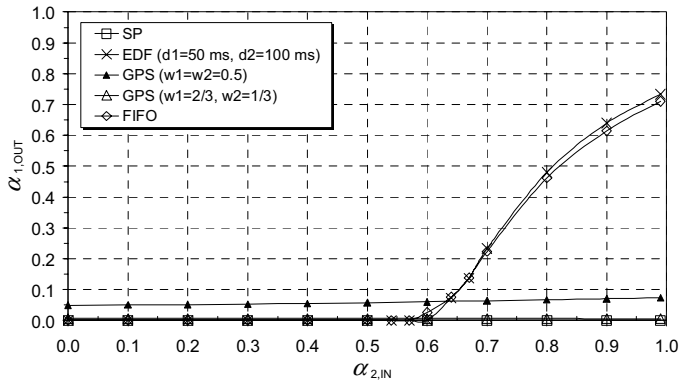


Fig. 9: Estimates of LRD parameter $\alpha_{1,OUT}$ as a function of $\alpha_{2,IN}$ ($\alpha_{1,IN} = 0.0$).

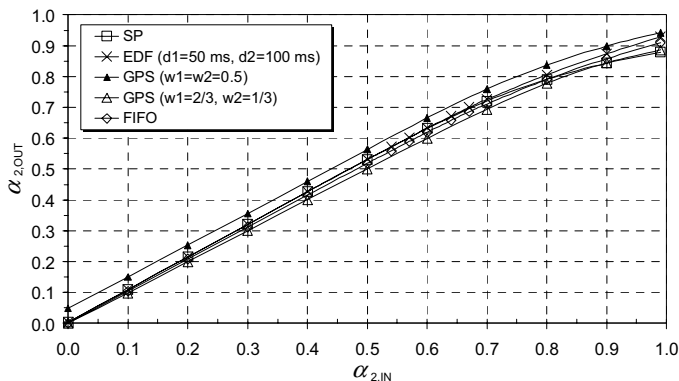


Fig. 10: Estimates of LRD parameter $\alpha_{2,OUT}$ as a function of $\alpha_{2,IN}$ ($\alpha_{1,IN} = 0.0$).

REFERENCES

- [1] 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.
- [2] V. Paxson, S. Floyd, "Wide-Area Traffic: the Failure of Poisson Modeling", *IEEE/ACM Trans. Networking*, vol. 3, no. 6, pp. 226-244, 1995.
- [3] N. G. Duffield, Neil O'Connell, "Large Deviations and Overflow Probabilities for The General Single Server Queue, with Applications", *Math. Proc. Cam. Phil. Soc.*, vol. 118, pp. 363-374, 1995.
- [4] I. Norros, "On The Use of Fractional Brownian Motion in the Theory of Connectionless Networks", *IEEE J. Select. Areas Commun.*, vol. 13, no. 6, pp. 953-962, Aug. 1995.
- [5] Y. Cheng, W. Zhuang, "Calculation of Loss Probability in a Partitioned Buffer with Self-Similar Input Traffic", *Proc. IEEE GLOBECOM 2004*, Dallas, TX, USA, Nov. 2004.
- [6] J. Choe, N. B. Shroff, "Queueing Analysis of High-Speed Multiplexers Including Long-Range Dependent Arrival Processes", *Proc. IEEE INFOCOM 1999*, New York, NY, USA, March 1999.
- [7] S. Bregni, P. Giacomazzi, G. Sadedmi, "Properties of the Traffic Output by a Leaky-Bucket Policier with Long-Range Dependent Input Traffic", *Proc. IEEE ICC 2007*, Glasgow, UK, June 2007.
- [8] S. Bregni, P. Giacomazzi, G. Sadedmi, "Output Traffic Characterization of Policers and Shapers with Long-Range Dependent Input Traffic", *Proc. IEEE GLOBECOM 2007*, Washington, DC, USA, Nov. 2007.
- [9] V. Anantharam, "Scheduling Strategies and Long-Range Dependence", *Queueing Systems*, Vol. 33, No. 1-3, December, 1999.
- [10] D. J. Daley, R. Vesilo, "Long range Dependence of Inputs and Outputs of Some Classical Queues", *Fields Inst. Commun.*, 2000.
- [11] S. Song, J. K.-Y. Ng, B. Tang, "Some Results on the Self-Similarity Property in Communication Networks", *IEEE Trans. Commun.*, vol. 52, no. 10, pp. 1636-1642, Oct. 2004.
- [12] S. Bregni, L. Jmoda, "Accurate Estimation of the Hurst Parameter of Long-Range Dependent Traffic Using Modified Allan and Hadamard Variances", to appear in *IEEE Trans. Commun.*, 2008. Extended version available: <http://home.dei.polimi.it/bregni/public.htm>.
- [13] D. W. Allan, J. A. Barnes, "A Modified Allan Variance with Increased Oscillator Characterization Ability", *Proc. 35th Annual Freq. Contr. Symp.*, 1981.
- [14] P. Lesage, T. Ayi, "Characterization of Frequency Stability: Analysis of the Modified Allan Variance and Properties of Its Estimate", *IEEE Trans. Instrum. Meas.*, vol. 33, no. 4, pp. 332-336, Dec. 1984.
- [15] L. G. Bernier, "Theoretical Analysis of the Modified Allan Variance", *Proc. 41st Annual Freq. Contr. Symp.*, 1987.
- [16] D. B. Sullivan, D. W. Allan, D. A. Howe, F. L. Walls, Eds., *Characterization of Clocks and Oscillators*, NIST Tech. Note 1337, March 1990.
- [17] S. Bregni, "Chapter 5 - Characterization and Modelling of Clocks", in *Synchronization of Digital Telecommunications Networks*. Chichester, UK: John Wiley & Sons, 2002, pp. 203-281.
- [18] 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, G. M. R. Winkler, "Characterization of Frequency Stability", *IEEE Trans. Instrum. Meas.*, vol. 20, no. 2, pp. 105-120, May 1971.
- [19] ITU-T Rec. G.810 "Definitions and Terminology for Synchronisation Networks", Geneva, 1996-2003.
- [20] J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress", *Proc. IEEE*, vol. 66, no. 9, pp. 1048-1075, Sept. 1978.
- [21] P. Abry, D. Veitch, "Wavelet Analysis of Long-Range Dependent Traffic", *IEEE Trans. Inform. Theory*, vol. 44, no.1, pp. 2-15, Jan. 1998.
- [22] A. Parekh, R. Gallager, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case", *IEEE/ACM Trans. Networking*, vol. 1, no. 3, pp. 344-357, 1993.
- [23] W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, *Numerical Recipes in C - The Art of Scientific Computing. 2nd Edition*, Cambridge, UK: Cambridge University Press, 2002.