Fractal Analysis of Intraflow Unidirectional Delay over W-LAN and W-WAN Environments

Dimitrios P. Pezaros, Manolis Sifalakis, and Laurent Mathy

Abstract—we have analysed unidirectional delay traces of a diverse set of IPv6 microflows routed over W-LAN and W-WAN environments. Using a number of time-domain and frequency-domain estimators we have examined the existence and intensity of long-range dependence in packet delay when viewed as time-series data. The correlation structures of packet delay on bulk TCP data path and UDP flows follow asymptotic decay while Hurst exponent estimates suggest from moderate to strong intensity $(H \rightarrow 1)$.

Index Terms—Long-range dependence, unidirectional delay, ACF, Hurst exponent

I. INTRODUCTION

Fractal analysis focuses on the temporal and spatial evolution properties of data series and on their correlation structure. One of the properties of fractal series is *self-similarity*. A stochastic process is *self-similar* (statistically fractal in nature) when the statistical properties of segments within the corresponding time series are similar, irrespective of the time scale of observation. A characteristic of self-similar processes is that they exhibit long memory, or Long-Range Dependence (LRD), signifying that their current state has significant influence on their subsequent states far into the future. Consequently, values at a particular time are related not just to immediately preceding values, but also to fluctuations in the remote past. Hence, high variability on the behaviour of self-similar processes is preserved over multiple time scales.

During the last decade, pioneering work has demonstrated the presence of LRD and self-similarity in various facets of (inter)network behaviour (e.g. [9][13][4]), and immensely influenced research mainly on network and traffic modelling, design and performance evaluation. However, the bulk of the work has focused on the "burstiness" preservation characteristics (their causes and implications) of long-term aggregate traffic originated at or destined to a single point of the network. The presence (or otherwise) of self-similar behaviour and LRD has not been investigated for long-lived microflow performance properties at a finer level of granularity, partly due to the absence of appropriate instrumentation mechanisms. The ability to characterise the long-term behaviour of intraflow performance metrics

such as one-way delay is important, since not only it exposes the variability properties and burstiness preservation in the performance between segments of the same flow for modelling purposes, it also provides valuable insights for resource management and traffic control algorithms which assume radically different enforcement between the per-flow and aggregate modes of operation. End-system application and transport software can also take such burstiness preservation properties into consideration when enforcing transmission and system-local resource control optimisations, especially while operating over media with known high delay variability such as IEEE 802.11 [7] and GPRS [3]. Performance measurement systems can take advantage of scale-invariance properties in reducing their intrusiveness through employment of adequately engineered sampling mechanisms.

In this paper we have employed four LRD estimators [16] to investigate the presence and the intensity of LRD in unidirectional packet delay experienced by a diverse set of IPv6 traffic flows as these were routed over two different wireless network technologies. The main contribution of this paper is twofold. First, evidence of LRD existence in unidirectional delay is presented for certain IPv6 flow types, as this is manifested by the four estimators and verified by the power-law decay of the correlation structure (ACF) of the corresponding time series. Second, the relevance between the LRD values produced by each estimator, as well as between estimator-based LRD manifestation and actual LRD existence is empirically evaluated.

In section II we provide the mathematical formulation of LRD and a very brief description of the LRD estimators used for this study. The measurement methodology and the calibration mechanisms employed in order to avoid known LRD estimation pitfalls are documented in section III. Section IV discusses the presence and intensity of LRD in the unidirectional delay experienced over the two wireless networks. Finally, section V concludes the paper.

II. LONG-RANGE DEPENDENCE: PHENOMENON AND ESTIMATORS

Long-Range Dependence (LRD) is a property that describes the memory of a process. A stationary process X_t has a long memory or is long-range dependent if there exists a real number $\beta \in (0,1)$ and a constant $c_\rho > 0$ such that

$$\lim_{k \to \infty} \rho(k)/[c_\rho k^{-\beta}] = 1$$

where $\rho(k)$ are the sample correlations [1]. Hence, the autocorrelation function of LRD processes not only decays hyperboli-

D. P. Pezaros, M. Sifalakis and L. Mathy are with the Computing Department, Infolab21, Lancaster University, Lancaster, LA1 4WA, UK
E-mail: {dp, mjs, laurent}@comp.lancs.ac.uk.

cally as the lag k increases, it is also non-summable, i.e. $\sum_{k=1}^{\infty} \rho(k) = \infty$, implying that, the individually small high-lag correlations have an important cumulative effect. This is in contrast to conventional short-range dependent processes which are characterised by an exponential decay of the autocorrelations. The intensity of LRD is measured by the *Hurst* parameter $H = 1 - \beta / 2$, $0.5 < H < 1$. LRD implies asymptotic second-order self-similarity -and vice versa-, which describes the property that the correlation structure of a time series is asymptotically preserved irrespective of time aggregation [12]. A stationary process X_t is asymptotically second order self-similar with Hurst parameter H if

$$\lim_{k \to \infty} \rho(k) = \frac{1}{2}((k+1)^{2H} - 2k^{2H} + (k-1)^{2H}), \ 0.5 < H < 1$$

A number of estimators [16][9] have been extensively used for LRD detection and quantification by estimating the value of the Hurst exponent; as $H \to 1$ the dependence is stronger. They are classified in time-domain and frequency-domain estimators, depending on the methodology they employ to estimate H. Time-domain estimators investigate the evolution of a statistical property of the time series at different time-aggregation levels. Frequency-domain estimators focus on the behaviour of the power spectral density of the time series.

In this paper, we have employed four most commonly used LRD estimators, two from each domain. The *aggregated variance* method examines the decay of the variance at increasing time aggregation levels. For LRD time series decays more slowly than the reciprocal of the sample size. The *rescaled adjusted range (R/S)* method examines the growth in the rescaled range of partial sums of deviations of the time series from its mean, as a function of the number of points in the time-aggregated series. For LRD series, the *R/S* statistic grows according to a power law with exponent H. The *periodogram method* is based on the discrete Fourier transform and is an estimate of the power spectral density of a discrete process, which should exhibit power-law behaviour for frequencies close to the origin. The *Whittle* estimator is a maximum likelihood type estimate which is applied to the periodogram of the time series.

LRD estimators evaluate different statistics of the time series to estimate the Hurst exponent and hence their estimates can vary even in synthesised self-similar processes with known Hurst exponent value. The accuracy and robustness of the estimators can be influenced by processes such as periodicity, trend, time series' length and short-range correlations which have different effects on different estimators [10][2][8][5] and can lead to erroneous estimation of the LRD intensity or even to reporting LRD on non-LRD series. No single LRD estimator has been proved to produce more accurate estimates than the rest. Frequency-based estimators are considered more accurate when applied to series which have already been shown to be LRD by some other (time-domain) estimator. Hence, intuitive data inspection and pre-processing, as well as the simultaneous employment of a mixture of time and frequency-domain estimators to assess LRD behaviour need to be considered.

III. Measurement Methodology

In-line measurement [14] has been employed to instrument a representative set of IPv6 traffic flows as these were routed over wireless LAN (IEEE 802.11b) and WAN (Orange, UK GPRS/GSM) topologies during different hours of the day and different days of the week, in November 2005 [15]. Instrumented traffic consisted of moderate-size bulk TCP transfers and CBR UDP video streaming flows. End-to-end unidirectional delay has been measured using the appropriate in-line measurement Type-Length-Value (TLV) encoded structures to piggyback 32-bit Linux kernel timestamps within an IPv6 destination options extension header. Unidirectional packet delay has been measured as the difference $D = T_B - T_A$ between the packet's arrival time (T_B) at the destination node's OS kernel and the departure time (T_A) before it was serialised at the source node's NIC.

A. Measurement calibration, data inspection and preprocessing

The two end-systems synchronised using the Network Time Protocol (NTP) with a common stratum 1 server through additional high-speed wired network interfaces, in order to avoid NTP messages competing with the instrumented traffic over the bottleneck wireless links. The NTP daemon was allowed sufficient time to synchronise prior to the experiments until it reached a large polling interval. The offset reported by NTP was always at least one order of magnitude smaller with respect to the minimum one-way delay observed. All the delay traces were empirically examined against negative values as well as against linear alterations (trend) of the minimum delay over time. None of these offset/skew-related phenomena were experienced [11].

Delay measurements were taken upon arrival of each packet to its destination, hence at irregular time instants. In order to covert the traces to time series data, they were discretised into equally-sized bins based on the packet arrival time. Unidirectional delays of multiple packets arriving within each bin were averaged, and the mean delay was considered for the particular bin. Although this process inevitably smooths out short-term variations in packet delay, bin size was carefully selected for each flow to contain as few packets as possible, while at the same time avoiding empty bins. The time series' lengths varied from 2^9 to 2^{12} which has been reported sufficient for Hurst exponent estimates with less than 0.05 bias and standard deviation [2]. Periodicity and short-term correlations in the time series have been removed by employing the randomised buckets method [6][8] to perform internal randomisation on the signal.
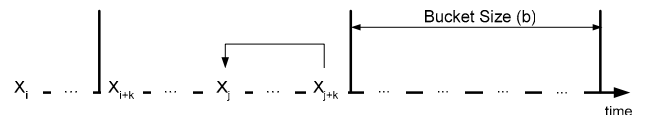

Figure 1: Randomised buckets with internal randomisation. Items are randomised within the same fixed-size bucket, but not across bucket boundaries. Classification depends on the fixed bucket boundaries, not just the separation between two items in the time series.

This involves partitioning the time series into a set of buckets of the same size and randomising the contents of each bucket, while the order of the buckets remains unchanged (Figure 1). Hence, correlations among *inbucket* pairs are equalised while correlations among *outbucket* pairs are preserved. Figure 2 shows how LRD is preserved after internal bucket randomisation.
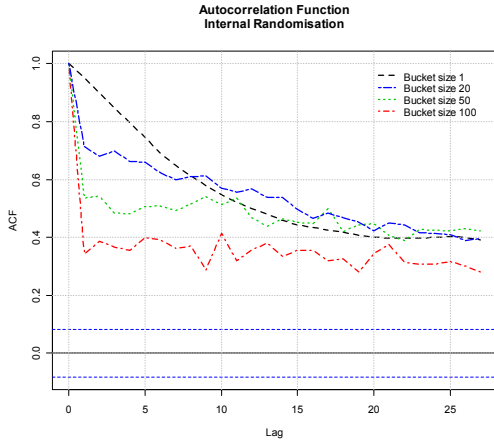

Figure 2: Autocorrelation function (ACF) for the delay time series of the TCP data path over W-LAN after internal bucket randomization for different bucket sizes. Bucket size 1 represents the unrandomised series. It can be seen that increasing the bucket size lowers and flattens the ACF for small values of the lag, however, for larger lag values the ACF is preserved.

IV. ON THE PRESENCE AND INTENSITY OF LONG-RANGE DEPENDENCE IN UNIDIRECTIONAL DELAY TIME SERIES DATA

The four LRD/Hurst exponent estimators have been applied to the unidirectional delay time series data collected over the two diverse wireless network technologies. The aggregated variance, the R/S and the periodogram are graphical estimation methods; an example of their output is shown in Figure 3. Whittle estimator, on the other hand, does not produce a graphical output, but it allows for the computation of a confidence interval for the Hurst exponent estimation. Whittle assumes a priori consistency of the data set with a specific process. Hence, for Hurst estimation purposes, it is advised that it is not used for LRD detection. Rather, it should be used to estimate the intensity of the phenomenon in time series that have already been shown to be LRD [9].

Flows have been classified into three types, bulk TCP data, TCP reverse, and UDP. There have been more similarities in the correlation structure (ACF) of the unidirectional delay time series between flows of the same type than between flows routed over the same wireless technology. Table 1 shows estimated H values produced by the four estimators for flows falling under each type for the two different wireless networks. It is worth highlighting that on a first sight the three detection estimators all report (even a slight, in some cases) existence of LRD on the delays of all traffic flows. Between the two time-domain estimators, the R/S method produces in general more conservative Hurst exponent estimates than the aggregated variance method. Two exceptions in this trend concern the unidirectional delays

of the TCP reverse path over the 802.11b network and the UDP streaming flow over the GPRS/GSM network. The strongest LRD (~1) has been reported by the periodogram estimator for the TCP data path over the W-LAN and for the UDP flows over both wireless networks. After LRD detection, the Whittle estimator was employed to quantify LRD intensity and produce 95% Confidence Intervals (C.I.)s. With the single exception of the UDP flow streamed over the W-WAN in which the C.I. is relatively large, Whittle reports the most conservative LRD intensity estimates than all the other methods.

TABLE 1: **HURST EXPONENT ESTIMATES**

Microflow	Aggregated Variance	R/S	Periodogram	Whittle
	Normal / Oversampling	Normal / Oversampling		[95% C.I.]
TCP data path [W-LAN]	0.875 / 0.840	0.815 / 0.792	0.948	0.739 [0.726-0.751]
TCP data path [W-WAN]	0.817 / 0.721	0.732 / 0.750	0.745	0.599 [0.527-0.672]
TCP reverse path [W-LAN]	0.591 / 0.532	0.688 / 0.764	0.592	0.552 [0.549-0.556]
TCP reverse path [W-WAN]	0.665 / 0.530	0.566 / 0.705	0.619	0.528 [0.521-0.535]
CBR UDP @1Mb/s [W-LAN]	0.909 / 0.865	0.719 / 0.781	0.976	0.687 [0.685-0.688]
CBR UDP @20Kb/s [W-WAN]	0.666 / 0.665	0.901 / 0.879	0.957	0.742 [0.644-0.840]

In order to further refine time-domain LRD estimators by increasing the number of samples at large time-aggregated scales, we have employed oversampling with digital low-pass filtering on the delay time series and re-estimated LRD using the aggregated variance and the R/S methods [8]. The original time series was decomposed down to D interleaved sub-series in round-robin ordering. Each sub-series was expanded using D-times oversampling via duplicate insertion, and then smoothed using a Gaussian filter. The expanded and smoothed sub-series were then separately processed with the aggregated variance and R/S methods, and the results were averaged. The resulting Hurst exponent estimates are shown in the relevant entries of Table 1. For the aggregated variance method, oversampling resulted in more conservative H estimates with respect to when applied at the original series for all flows. Oversampling had a less deterministic effect on the R/S method, resulting in either lower or higher LRD intensity estimates for different flows.

A number of studies have examined the accuracy of LRD estimators for more refined data analysis to suggest that they often overestimate or underestimate the value of the Hurst exponent under different circumstances [1][2][5][8][10]. After inspecting the correlation structures (ACF) of the delay time series for the different flows, it appears that LRD is unlikely to be present in the packet delay of the *TCP reverse (ack) path* flows over both wireless networks, for which the majority of the estimators report small H values (<0.65). Figure 4 shows that their ACFs eventually fall within the 95% C.I. of zero. Absence of such correlations implies a link between LRD behaviour and packet size, since *ack* packets are triggered by data packets. For the rest of the studied flows over both wireless networks, the correlation structures of their unidirectional delay time series followed asymptotic power-law behaviour, visually similar to the one shown in Figure 2 for TCP the data path over W-LAN.
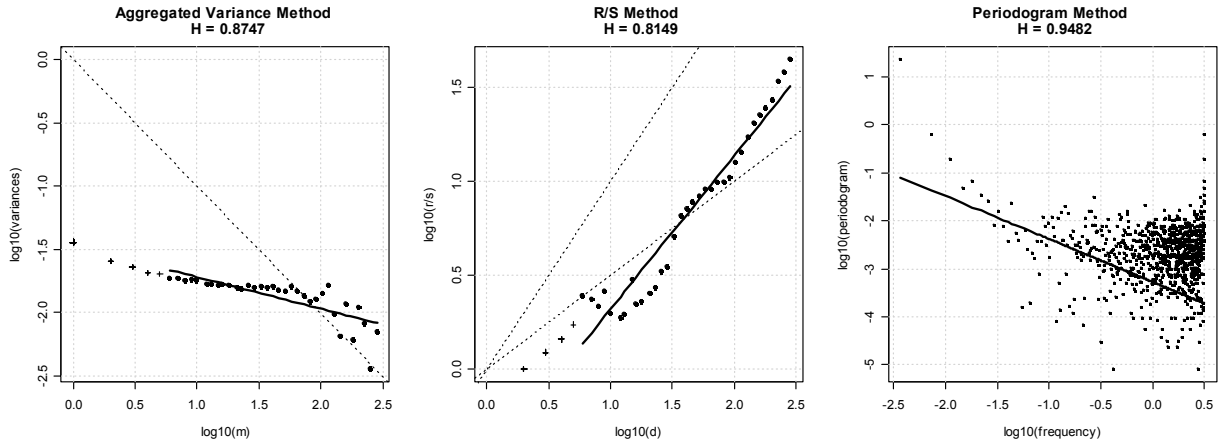
Figure 3: The three graphical LRD estimation methods applied to the unidirectional delay time series of a TCP data flow over 802.11b. (Left): The aggregated variance method plots the normalised variance over the time aggregation level on a log-log scale. The slope of the linear regression of the data points is $\beta = 2H - 2$. For LRD, $-1 \le \beta \le 0$ should hold. (Middle): The rescaled range method plots the R/S statistic versus the number of points of the aggregated time series on a log-log scale. The slope of the linear regression of the data points is a direct estimate of H. (Right): The periodogram plots the logarithm of the spectral density versus the logarithm of the frequency. Linear regression for the lowest (~10%) frequencies should give a slope of $1 - 2H$, also avoiding the inherently noisy high frequencies.
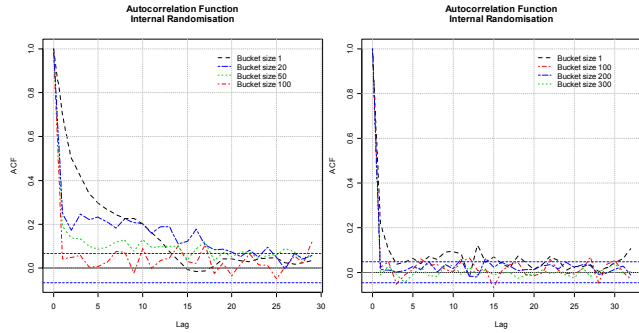

Figure 4: Correlation structures of the unidirectional delay time series for the TCP reverse path flows over the W-LAN (left) and the W-WAN (right) networks.

V. CONCLUSION

We have provided evidence of LRD in the unidirectional packet delay experienced by certain types of IPv6 traffic flows as these were routed over two different wireless network technologies. A variety of LRD estimators have been used which all agreed on LRD existence albeit their estimates varied regarding the intensity of the phenomenon. Through inspection of the time series' correlation structure, we argue that LRD cannot be safely assumed when the majority of the detection estimators does not agree on a Hurst exponent value greater than 0.65. Similar behaviour is expected for IPv4 traffic, since both protocols assume the same packetisation mechanisms at their transport and application layers. Whether intermediate routers actually handle both IPv4 and IPv6 traffic identically and how this can influence the end-to-end packet delay, deserves further experimental investigation. The presence of a limited number of fractal analyses studies focusing on intraflow unidirectional packet delay time series constitutes it a promising area of further research with respect to the end-to-end performance experienced by long-lived flows. Performance differences between flows routed over wireless technologies with known high delay variability and their wired counterparts are to be further studied.

REFERENCES

[1] J. Beran, Statistics for Long-Memory Processes. Monographs on Statistics and Applied Probability, Chapman and Hall, New York, NY, 1994

[2] M. J. Cannon, D. B. Percival, D. C. Caccia, G. M. Raymond, J. B. Bassingthwaighte, Evaluating scaled windowed variance methods for estimating the Hurst coefficient of time series. Physica A, 241, 1997

[3] R. Chakravorty, J. Cartwright, I. Pratt, Practical experience with TCP over GPRS, IEEE GLOBECOM'02, Taiwan, 2002

[4] M. E. Crovella and A. Bestavros, Self-similarity in World Wide Web traffic: Evidence and possible causes, IEEE/ACM Trans. on Networking, 5(6): 835 - 846, December 1997

[5] D. Delignieres, S. Ramdani, L. Lemoine, K. Torre, M. Fortes, G. Ninot, Fractal analyses for 'short' time series : a re-assessment of classical methods, Journal of Mathemtical Psychology, Vol. 50:525-544, 2006

[6] E. Erramilli, O. Narayan, W. Willinger, Experimental queuing analysis qith long-range dependent packet traffic, IEEE/ACM Trans. on Networking, 4(2):209-223, 1996

[7] G. Fotiadis, V. Siris, Improving TCP throughput in 802.11 WLANs with high delay variability". 2nd IEEE Int'l Symposium on Wireless Communication Systems (ISWCS'05), Siena, Italy, September 2005.

[8] T. Karagiannis, M. Molle, M. Faloutsos, Understanding the limitations of estimation methods for long-range dependence, Technical Report, University of California, Riverside, TR UCR-CS-2006-10245, 2006

[9] W. E. Leland, M. Taqqu, W. Willinger, D. V. Wilson, On the Self-Similar Nature of Ethernet Traffic (extended version), IEEE/ACM Trans. on Networking, 2(1): 1 – 15, 1994

[10] A. Montanari; M. S. Taqqu, V. Teverovsky, Estimating long-range dependence in the presence of periodicity : an empirical study, Mathematical and computer modelling, Vol. 29 (10-12), pp. 217-228, 1999

[11] S. B. Moon, P. Skelly, D. Towsley, Estimation and Removal of Clock Skew from Network Delay Measurements, IEEE INFOCOM, New York, NY, March 21-25, 1999

[12] K. Park and W. Willinger (eds.), Self-Similar Network Traffic and Performance Evaluation. John Wiley & Sons, New York, NY, USA, 2000

[13] V. Paxson and S. Floyd, Wide-area traffic: The failure of Poisson modelling, IEEE/ACM Trans. on Networking, 3(3):226-244, 1994.

[14] D. P. Pezaros, D. Hutchison, F. J. Garcia, R. D. Gardner, J. S. Sventek, In-line Service Measurements: An IPv6-based Framework for Traffic Evaluation and Network Operations, IEEE/IFIP Network Operations and Management Symposium (NOMS'04), Seoul, Korea, April 19-23, 2004

[15] D. P. Pezaros, M. Sifalakis, D. Hutchison, End-To-End Microflow Performance Measurement of IPv6 Traffic Over Diverse Wireless Topologies, Wireless Internet Conference (WICON'06), Boston, MA, August 2-5, 2006

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