Does fractal scaling at the IP level depend on TCP flow arrival processes?

Nicolas Hohn[1], Darryl Veitch[1] and Patrice Abry[2]

Abstract—**In addition to the well known long-range dependence in time series of IP bytes and packets, evidence for scaling behaviour has also been found at small scales for these series, separated by a characteristic transition timescale. It is less well known that two scaling regimes are also commonly found in time series describing the arrivals of TCP flows, again with long-range dependence, and with a broadly similar scaling exponent at small scales. The transition timescale is also roughly similar to that found in the IP level case. We investigate the dependencies between the scaling behaviours of the IP and TCP arrival levels at both small and large scales. We also study the origin of scaling at small scales at the IP level. The arrival level process is important to study both for its potential impact on the IP level, and in its own right, for example for web server performance. Our findings are based on gigabytes of high precision packet level data collected at multiple locations. The analysis methodology combines models with real data in a 'semi-experimental' approach which reduces the need for modeling assumptions. Flows and packets are individually manipulated to selectively isolate the components of scaling due to packet dynamics within a TCP flow, the dependencies between flows, their durations and packet counts, and the flow arrival process. The scaling behaviour is analysed using wavelet based methods.**

***Keywords*—scaling, long range dependence, wavelets, TCP arrival times, traffic modeling, Internet data.**

I. Introduction

As is now well known, fractal-like scaling is a feature of many different time series derived from packet traffic. Thus far three kinds of scaling have been identified as relevant, or potentially relevant. Long-range dependence (LRD), which describes persistent memory over large time scales, has been convincingly and almost universally found in time series such as discretised byte or packet rate [1]. Multifractal scaling [2], [3], has been suggested as a model of the extreme burstiness often observed at small scales, and Infinitely Divisible Cascades [4] have been put forward as a means of unifying the scaling behaviour across all scales. Our view is that, as scaling typically implies high variability which in turn implies worse queuing performance as explored for example in [5], it is important to understand it. The presence of scaling could imply an underlying mechanism or mechanisms which deserves to be understood. Unless the source of such behaviour is known, it will not be possible to predict whether it is a function for instance of protocol details which will change over time, or of some deep feature of the underlying traffic sources.

1 ARC Special Research Center for Ultra-Broadband Information Networks, Department of Electrical and Electronic Engineering, The University of Melbourne, Australia. E-mail: {n.hohn, d.veitch}@ee.mu.oz.au

2 CNRS, UMR 5672, Laboratoire de Physique, Ecole Normale Supérieure de Lyon, France. E-mail: pabry@ens-lyon.fr

In this paper we seek to shed light on the origins of scaling in traffic, particularly at small scales. Our starting point is the somewhat surprising observation that the scaling seen at the IP level, such as packets counts, is roughly similar to that found in the arrival process of TCP flows. Namely, clear LRD at large scales, a second, though less clear, scaling regime at small scales, and a transition scale at around 1 second separating them. This is surprising in that the prevailing view on the origins of LRD at the IP level, namely heavy tailed file sizes [6], cannot explain LRD in the arrival process. This similarity immediately raises the question of the link between the two. Are the twin scaling regimes at the IP level, or aspects of it, due to or influenced by the corresponding features at the flow level, or are they both the result of some common mechanism, or even two independent mechanisms? Answers to such questions will tell us if the fractal structure of arrivals is important to model accurately or not. This is important for hierarchal traffic models where an arrival process of sessions, and then flows, forms the backbone of the final packet level model. Although the IP level is of great importance for router throughput, another motivation for pursuing an understanding of the flow level arrival processes is the direct role they play for flow level performance, for example in web servers and proxies.

In [7] the idea of 'shuffling', the random reordering of blocks of a time series, was introduced as a way of modifying the correlations of the data whilst preserving the original structure within blocks. We extend this idea and selectively modify several of the components comprising the full packet stream. We call this way of virtually investigating 'what if' scenarios the *semi-experimental method*, and we employ it extensively as a tool to track down the connections and origins of scaling behaviour. It can also be used to selectively test models for portions of the traffic structure, without having to postulate a full model from the outset, a difficult task for such complex data. For example, details of the arrival process of flows can be altered while preserving in full the packet patterns within each flow, and the resulting effect on the scaling structure noted.

After introducing the data in the next section, we then quickly introduce the wavelet based statistical tool that we use to analyse the time series and examine their scaling. We then apply the semi-experimental method in section IV and report our findings. A conclusion is offered in section V.

II. The Data and Data Processing

A. The Raw Data

The Internet traces we analyse were mainly recorded by the WAND group at the University of Waikato in New Zealand.

These traces, the Auckland II and Auckland IV data sets, were collected on the Internet access link of the University of Auckland with high precision hardware allowing loss-less measurements of the OC3 ATM link with GPS synchronized timestamps [8], and are freely available on the web [9]. The traces gather the timestamp, ATM header and the first 40 bytes of the ATM payload, which is sufficient in most cases to extract IP and TCP header information. For privacy reasons addresses are mapped, and TCP and UDP payloads removed. In this paper, we analyze subsets of these datasets, details of which are summarized in table I. We focus on two three hour periods during week days, 2:00 to 5:00 and 13:00 to 16:00, corresponding to apparently stationary traffic rate for a low and high activity period respectively.

We also study traces recorded by the DiRT group [10] at the University of North Carolina (UNC-a0 and UNC-a1), and from the NLANR repository [11] (NLANR-SDC and NLANR-TXS) for which details can be found in table I. These traces are used to make sanity checks on our main results as they are from different geographical regions and have different bit rates.

B. Time Series Extraction

The raw traces are processed with the CAIDA Coralreef tool suite [12] and our own C programs, allowing the extraction of each IP packet header together with an accurate timestamp. The information therein allows IP packets to be categorized into different *flows*. A flow is defined as a set of successive packets with the same 5-tuple: IP protocol carried, source address, destination address, source port and destination port, and where no packet interarrival exceeds a given time interval, fixed here at 64 seconds [12]. This classification only uses IP level information and gives therefore a general framework to compare TCP and UDP flows. In the case of TCP, it was found that the above definition gave a very similar classification to that provided by tracking TCP connections by monitoring SYN, SYN-ACK, FIN and RST packets, with the additional advantage of keeping track of late packets transmitted after connection closure. This technique also captures the many connections which do not terminate correctly.

From the raw data many different time series can be constructed. At the ‘IP level’, where flows are not individually tracked, we concentrate on the packet and byte rate series, IPpktTCP and IPbytTCP, counting the number of packets or bytes over all TCP flows in bins τ ms wide. At the ‘flow level’ statistics of individual flows are collected. The time series of TCP flow inter-arrival times, denoted by TCPIar, is intrinsically discrete. Counting the number of arrivals in bins of width τ yields TCPArr. Although Iar is richer than Arr and determines its properties, we focus on Arr as it is indexed by time, enabling easier and more meaningful comparisons across different traces and subsets of flows. We typically set $\tau = 5\,\text{ms}$. We also work with the intrinsically discrete series TCPsizepkt, TCPsizebyt and TCPdur which give the size in packets, size in bytes, and durations in seconds of successive TCP flows respectively.

In [13] the concept of *alpha* and *beta* traffic was introduced. We performed our analysis on the *beta* component of IPpktTCP and found the same results as for IPpktTCP.

Traces	Date	Time (local time)	Rate (Mbps)	Link
AUCK-a0	19991201	13:00 to 16:00	1.4	OC3
AUCK-b0	20010330	13:00 to 16:00	3.5	OC3
AUCK-c0	20010402	02:00 to 05:00	0.3	OC3
AUCK-c1	20010402	02:00 to 05:00	0.5	OC3
AUCK-d0	20010402	13:00 to 16:00	3.6	OC3
AUCK-d1	20010402	13:00 to 16:00	2.4	OC3
UNC-a0	20000927	19:30 to 20:30	179.8	OC12
UNC-a1	20000927	19:30 to 20:30	44.8	OC12
NLANR-SDC	19981126	90s peak period	11.0	OC3c
NLANR-TXS	20020110	90s peak period	22.5	OC3c

TABLE I

III. Wavelet Analysis

To study scale invariant properties such as long range dependence we use a wavelet-based analysis. A thorough description of wavelet transforms can be found in [14], and see [15] for theoretical and practical details of their use in the spirit of this article. Here we briefly describe the key features and give a short guide to interpretation.

A. Definitions and Properties

Performing the (discrete) wavelet transform of a process X consists in computing coefficients that compare, by means of inner products, X against a family of functions:

$$d_X(j,k) = \langle X, \psi_{j,k} \rangle. \tag{1}$$

The wavelets $\psi_{j,k}(t) = 2^{-j/2}\psi(2^{-j}t - k)$ derive from an elementary function ψ, called the mother wavelet, dilated by a factor 2^j and translated by $2^j k$. They are required to have excellent localization properties jointly in time and frequency. A key practical advantage is the fact that the coefficients can be computed from a fast recursive algorithm with computational complexity $O(n)$.

Let $X(t)$ be a continuous time stationary process with power spectral density $\Gamma_X(\nu)$. It can be shown that the variance of its wavelet coefficients satisfies:

$$\mathbb{E}|d_X(j,k)|^2 = \int \Gamma_X(\nu) 2^j |\Psi(2^j\nu)|^2 d\nu, \tag{2}$$

where $\Psi(\nu)$ denotes the Fourier transform of ψ. If X possesses scale invariance over a range of scales, for example if it is LRD, defined as a power law divergence of the spectrum at the origin:

$$\Gamma_X(\nu) \sim c|\nu|^{-\alpha},\ |\nu| \to 0, \ \ \text{with } \alpha \in (0,1), \tag{3}$$

then in the limit of large scales equation (2) becomes

$$\mathbb{E}|d_X(j,k)|^2 \sim C 2^{\alpha},\ j \to +\infty. \tag{4}$$

In fact equation (2) can be viewed as defining a kind of wavelet energy spectrum, analogous to a Fourier spectrum, but much better suited to the study of fractal processes.

To estimate the wavelet spectrum from data, the simple time averages

$$S_2(j) = \frac{1}{n_j} \sum_k |d_X(j,k)|^2,$$

where n_j is the number of $d_X(j,k)$ available at scale j, perform very well, because of the short range dependence in the wavelet

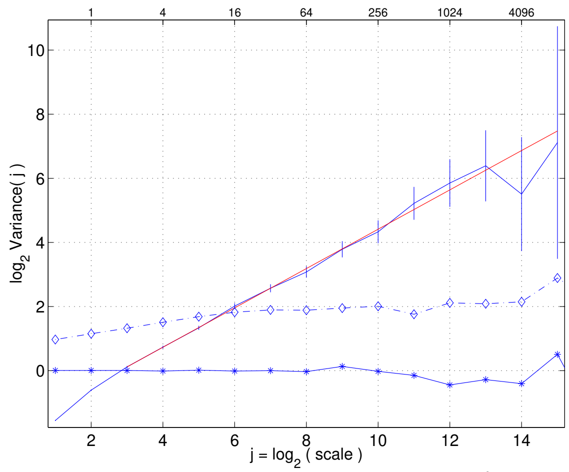

Fig. 1. Examples of LDs. Lower curve: Poisson process ($\lambda = 1$), Circles: Renewal process with gamma inter-arrivals ($\lambda = 1$, shape=1/4), Top plot: fGn ($H = 0.8$, $\alpha = 0.6$).

domain. A plot of the logarithm of these estimates against j we call the *Logscale Diagram*:

$$\text{LD:} \quad \log_2 S_2(j) \text{ vs } \log_2 2^j = j.$$

In these diagrams, straight lines constitute experimental evidence for the presence of scaling. For example, a straight line observed in the range of the largest scales with slope $\alpha \in (0, 1)$ (see figure 1) betrays long memory. More generally, semi-parametric estimates of scaling exponents with excellent properties can be formed using weighted regression to measure the slope over the range of scales where the scaling exists.

B. *Making Sense at Small Scales*

The analysis at small scales is considerably more difficult than at large scales. We address three relevant issues which are typically ignored.

(1) Confidence intervals often receive little attention, or are based strongly on Gaussian assumptions. Since at small time scales TCP/IP data is highly non-Gaussian, we use a semi-parametric technique based on general wavelet properties to estimate them more directly from data.

(2) The $O(n)$ algorithm which calculates the $d_X(j, k)$ requires initialisation using $X(t)$, however for real data, typically this is either omitted, or only samples $X(k\tau)$ are available, resulting in initialisation errors which are very significant for $j = 1, 2$. This is important as $3/4$ of the data is concentrated at these scales! We therefore use the Haar wavelet, which although a poor choice from the point of view of robustness to non-stationarity [15], does not suffer from such errors.

(3) For intrinsically discrete data such as TCPIar, as standard wavelet analysis does not apply, we use the special initialisation step of [16], without which, again, significant errors are made for $j = 1, 2$.

As a guide to interpretation, in figure 1 Logscale Diagrams are given of two continuous time and one discrete time process. In the continuous cases the base resolution, $j = 0$, was set to $\tau = 1/4$ as an example. The horizontal axis is calibrated both in octave j and time $t = \tau * 2^j$. The lower curve is for a Poisson process with $\lambda = 1$, viewed as a continuous time process with delta functions at each arrival point, with spectrum $\Gamma(\nu) = \lambda^2\delta(\nu) + \lambda$. Equation (2) predicts $\mathbb{E}|d_X(j, k)|^2 = \lambda$, a flat wavelet spectrum corresponding to trivial scaling ($\alpha = 0$), which agrees with the estimate in the figure, as $\log_2(S_2(j)) = \log_2(\text{variance}(j)) = \log_2 \lambda = 0$. It is important to understand that this level corresponds to variance and not to rate. Means are eliminated by the wavelet analysis, and multiplication of $X(t)$ by a constant a translates as a level shift in the LD of $\log_2(a)$. A Poisson process is a simple model of flow or packet arrivals, however real inter-arrival times are not necessarily exponential. The middle curve shows a point process with i.i.d. gamma distributed inter-arrivals with shape parameter $c = 1/4$, also with $\lambda = 1$. The spectrum is no longer flat at small scales, but it is asymptotically flat at a level of $\log_2(\lambda/c) = 2$ which reflects the higher variance $1/c\lambda^2$ of the inter-arrivals. An approximate onset scale for this trivial scaling at large scale is $\log_2(16/\lambda\tau) = 6$. Note the apparent scaling at small scales with $\alpha > 0$. The third plot is the familiar near-linear graph of fractional Gaussian noise (fGn), a discrete time series with an early onset of LRD at $j = 3$. Note how the confidence intervals are smallest at small scale.

IV. Results and Discussion

The most common types of time series extracted from packet level TCP/IP data are IP bytes or packets per bin. For each of the traces in table I, the IPpktTCP and IPbytTCP series were extracted with $\tau = 5$ ms, and their LDs plotted in figure 2(a) and (b) respectively. In each plot the time series are normalized to a common variance to facilitate a qualitative comparison.

We observe that, despite the differences between geographical region, time of day, link rate and average traffic rate, each exhibit a roughly similar *biscaling* behaviour. That is, two separate scaling regimes, separated by a remarkably invariant transition timescale or *knee* in the plots at around 1sec. Note that the NLANR traces are short, so the confidence intervals (not shown in the plots) are large.

As stated in the introduction, our starting point is the claim that a roughly similar biscaling is found in the flow arrival time series. To support this, examples of TCPArr with $\tau = 5$ ms are given in figure 2(c), again for each trace in table I. Although the knee position can be seen to vary across the traces, the broadly similar biscaling shape is evident (see [17] for an explanation of such knee movement as a function of flow duration). We now investigate the connection between these superficially very similar wavelet spectra at the IP and flow levels. We have found that the scaling of IPbytTCP closely follows that of IPpktTCP, so we concentrate on the latter.

Our approach is to begin at the IP level, and progressively modify aspects of it to determine the links to the arrival level and the source(s) of the scaling behaviour. Note that we are only interested in transformations which have a physical interpretation in terms of flow arrivals or packet structure within flows. We do not consider 'black box' modifications based solely on bins, such as random shuffling of blocks of a given size. We begin with the top row of figure 3, which displays the results of several semi-experiments for AUCK-c1. In each plot, the gray curve with the confidence intervals gives the original Logscale Diagram for the IPpktTCP time series.

Three categories of manipulation will be explored:

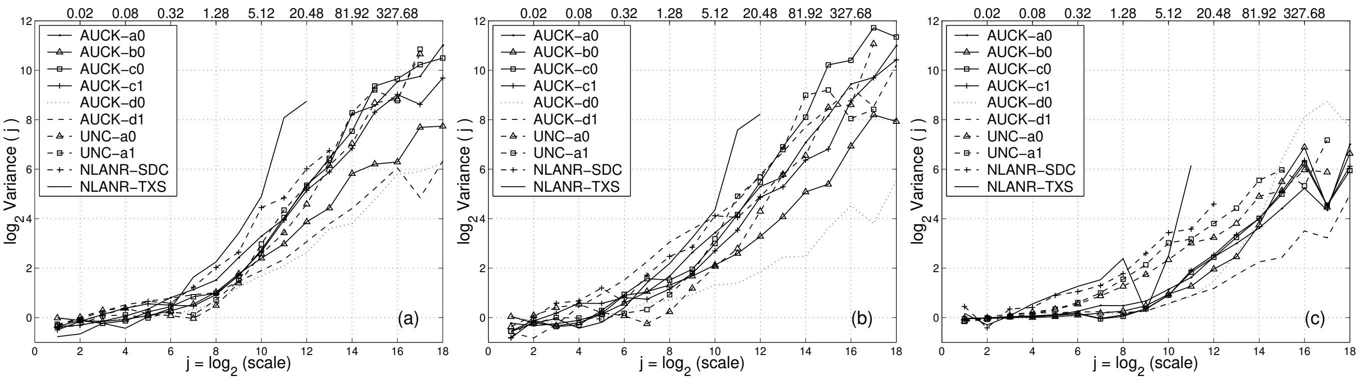

Fig. 2. Biscaling in (a) IPpktTCP, (b) IPbytTCP, and (c) TCPArr across all traces

- **A** Flow **A**rrival manipulation
- **P** **P**acket-in-flow manipulation
- **S** Flow **S**election manipulation

Flow Arrival Manipulation These are described in figure 3(a). The arrival process of flows is modified in three separate ways of increasing severity, whilst maintaining in full the integrity of the packet arrival patterns within each flow. Specifically:

A-Perm: Permute flows around the original arrival points.
A-Pord: Retain original flow order, but re-position arrival times according to a Poisson process with the same rate.
A-Pois: Combine the previous two: a Poisson arrival process with randomised flow re-assignments.

Figure 3(a) shows that none of these manipulations has any significant effect on the IP level scaling, even **A-Pois**, which completely erases the original arrival process structure. Two important inferences follow from this result:

1. The biscaling structure in the arrival process is not responsible for the biscaling structure at the IP level, and in fact does not influence it at either small or large scales.
2. Dependencies between packet processes across different flows is very weak.

The above inferences have important consequences. The first indicates that, at least in terms of second order statistics, it is pointless to include properties of the arrival process beyond the average rate in models of IP level traffic. This is significant as there is considerable interest in hierarchal modeling approaches where packet level traffic characteristics are derived beginning from a model of web session arrivals, leading to correlated launching of TCP connections and so on. The second point indicates strongly that there is no synchronisation (driven by TCP dynamics or anything else) between packet level processes across flows.

Thus far, in terms of relevance for IP packets, we have an image of traffic as a collection of entirely independent flows which are layed down in some independent way.

Packet-in-Flow Manipulation Figure 3(b) reports on the results of two manipulations:

P-Pois: Flow arrival times, durations etc. are retained in full. Within each flow separately, packet arrival times are replaced by a Poisson process of the same rate.
[A-Pois; P-Pois]: Combining Poisson packet arrivals in flows and randomised Poisson flow arrivals.

The effect of randomising the packet patterns within flows is clearly visible, although not overwhelming, and restricted to small scales. It is significant however that the spectrum has become flat. From figure 1 we know that this does not necessarily indicate that the process has become Poisson at small scales, however, as the level is equal to the arrival rate, this is the case here. On the other hand the large scale behaviour seems unaffected. Two tentative conclusions of note emerge from these observations:

1. The scaling structure at small scales has its origin in the packet patterns within flows.
2. The LRD structure at large scales is not influenced by the packet level structure within flows.

The fact that the manipulations **[A-Pois; P-Pois]** and **P-Pois** give such similar results simply reinforces the earlier conclusion that the flow arrival process does not impact on the IP level.

Through exploring the effects of both arrival and packet structure, we have been able to isolate the source of small scale scaling in IP, however the large scale behaviour has remained unaffected thus far.

Flow Selection Manipulation After performing **[A-Pois; P-Pois]**, the only original features of the traffic left are the flow durations and the flow packet counts. We explore these by selecting flow subsets. Figure 3(c) reports on the results of 5 manipulations, a single 'pure' flow selection, followed by four combinations.

S-Thin: Flow and packet structure is fully retained, flows thinned by rejecting with probability 0.3.
[A-Pois; S-Dur]: Combining flows with durations below the 70% percentile with randomised arrival times.
[A-Pois; S-Pkt]: Combining flows with packet volumes below the 70% percentile with randomised arrival times.
[A-Pois; P-Pois; S-Dur]: Randomising packet arrivals in flows in addition to [A-Pois; S-Dur].
[A-Pois; P-Pois; S-Pkt]: Randomising packet arrivals in flows in addition to [A-Pois; S-Pkt].

The random thinning **S-Thin** leads to a LD with the same shape as the original, with a variance which is approximately 70% of it, consistent with an i.i.d. superposition model, where variances simply add. In contrast, rejecting the flows with the longest durations removes the LRD, in keeping with the findings of [6] that show how the LRD of the IP level can be explained by the heavy tailed distribution of file sizes. The same occurs when the flows with the heaviest packet volumes are excluded,

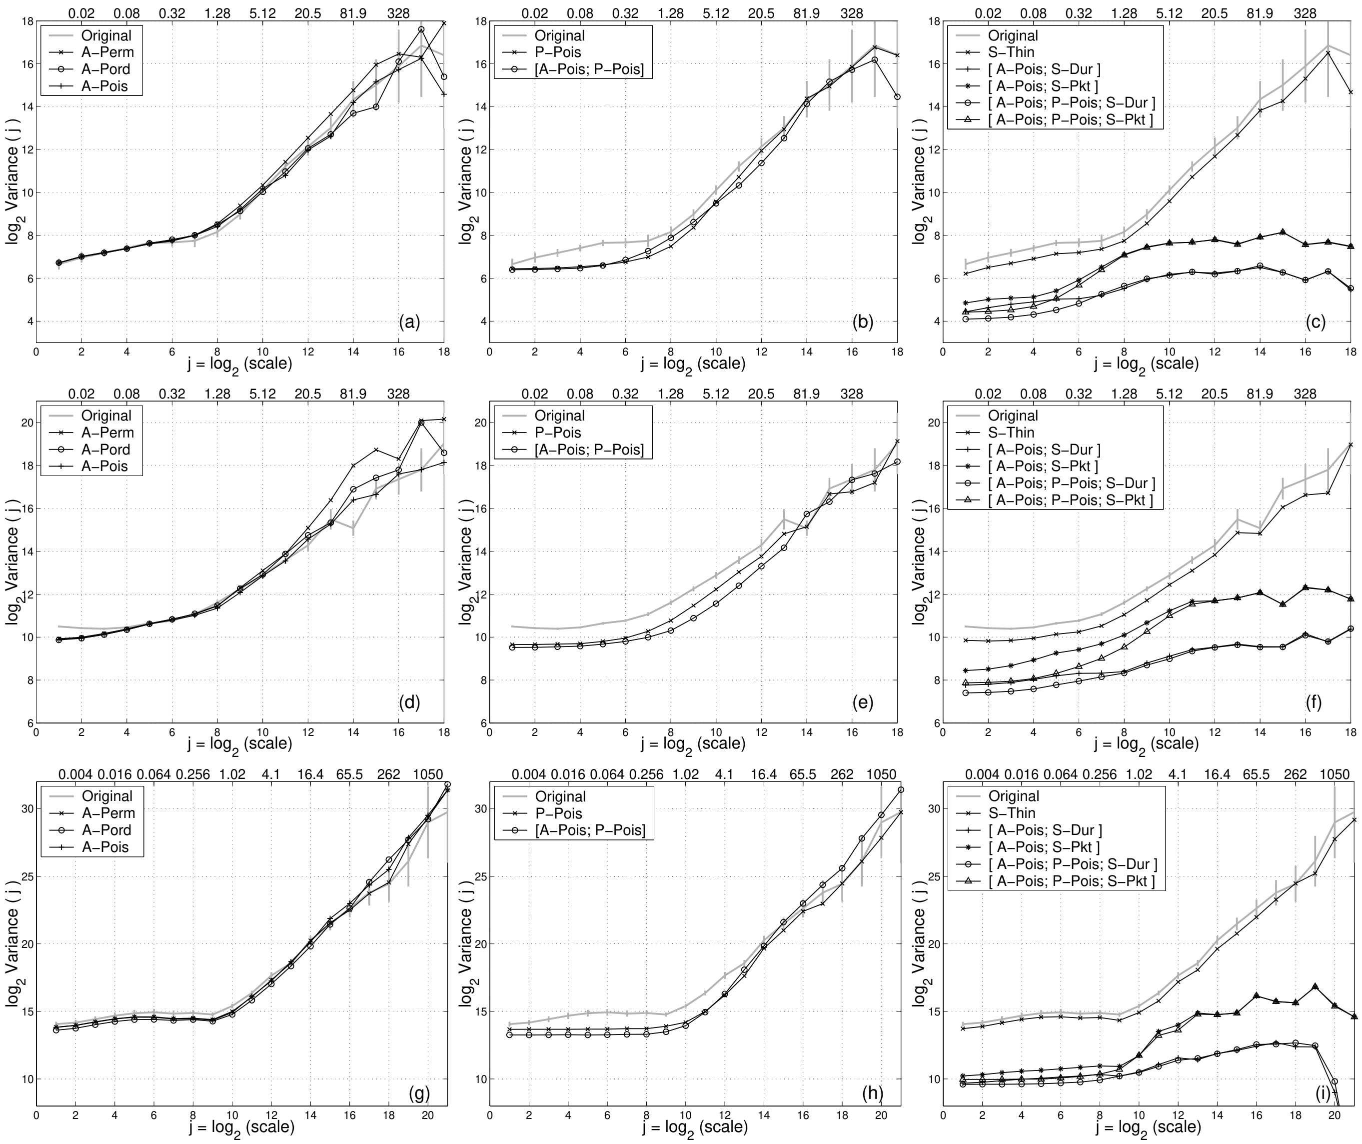

Fig. 3. Semi-experimental method applied to IPpktTCP of AUCK-c1 (a,b,c), AUCK-b0 (d,e,f) and UNC-a0 (g,h,i).

which can be explained via a strong dependency between flow durations and volumes. Finally, randomising the packet arrivals within flows has the same effect at small scales as **P-Pois** previously, but has negligible effect at large scales, in agreement with the conclusions above. The main conclusions we can draw are:

1. The LRD in IPpktTCP has origins in the heavy tailed nature of flow durations (a known result), and does not have a component due to packet processes within flows (controversial result).
2. When the concern is IP level modeling only, flows can be viewed as arriving as a Poisson process, with no dependence on other flows.

The second row in figure 3 shows the same manipulations for a higher rate trace, AUCK-b0. The results are very similar, although we observed two systematic differences in the outbound Auckland traffic during the peak period of the day: (i) a small flow arrival dependence at the smallest scales (note the drop on the left in graph (d)), and (ii) a smaller LRD exponent for flow arrivals (figure 2c). We speculate that this could indicate some traffic shaping at small scales.

The third row in figure 3 shows the results of the same manipulations for trace UNC-a0 which was recorded in a different location and has a rate 3 orders of magnitude higher than AUCK-c1. The fact that they are again very similar indicates that the findings presented in this paper are of wide applicability.

V. Conclusion

Although further validation from an even wider range of processes is desirable, we can tentatively answer the question in the title as follows: the fractal scaling at the IP level does not depend to any significant extent on the TCP arrival process. Furthermore, we have found that the scaling at small scales is related to the arrival process of packets within flows, but showed that this could be consistent with a simple renewal process, rather than a more complex process such as a multifractal. Work is continuing to understand the source of the LRD at the arrival level, and the nature of the packet structure in flows.

Acknowledgement

We thank F. H. Campos for making the UNC data set available.

REFERENCES

[1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. on Networking*, vol. 2, no. 1, pp. 1–15, 1994.

[2] J. Lévy Véhel and R. H. Riedi, *in Fractals in Engineering'97,* J. Lévy Véhel, E. Lutton and C. Tricot, editors, chapter Fractional Brownian motion and data traffic modeling: The other end of the spectrum, Springer, 1997.

[3] A. Feldmann, A. Gilbert, and W. Willinger, "Data networks as cascades: Explaining the multifractal nature of Internet WAN traffic," in *ACM/Sigcomm'98, Vancouver, Canada*, 1998.

[4] S. Roux, D. Veitch, P. Abry, L. Huang, P. Flandrin, and J. Micheel, "Statistical scaling analysis of TCP/IP data," in *ICASSP 2001, Special session, Network Inference and Traffic Modeling*.

[5] A. Erramilli, O. Narayan, A. Neidhardt, and I. Saniee, "Performance impacts of multi-scaling in wide area TCP/IP traffic," in *Proceedings of IEEE Infocom'2000*, Tel Aviv, Israel, March 2000.

[6] 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," in *Proceedings of the ACM/SIGCOMM'95*, 1995.

[7] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent packet traffic," *IEEE/ACM Transactions on Networking*, vol. 4, no. 2, pp. 209–223, April 1996.

[8] J. Micheel, I. Graham, and N. Brownlee, "The Auckland data set: an access link observed," in *Proceedings of the 14th ITC Specialist Seminar*, 2000.

[9] `http://wand.cs.waikato.ac.nz/wand/wits/`

[10] `http://www.cs.unc.edu/Research/dirt/`

[11] `http://www.nlanr.net/`

[12] `http://www.caida.org/tools/measurement/coralreef/`

[13] S. Sarvotham, R. Riedi, and R. Baraniuk, "Connection-level analysis and modeling of network traffic," in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop*, 2001.

[14] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.

[15] P. Abry, P. Flandrin, M. S. Taqqu, and D. Veitch, "Wavelets for the analysis, estimation, and synthesis of scaling data," in *Self-Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger, Eds. Wiley, 2000.

[16] D. Veitch, M. Taqqu, and P. Abry, "Meaningful MRA initialisation for discrete time series," *Signal Processing*, vol. 8, pp. 1971–1983, 2000, Elsevier Science.

[17] N. Hohn, D. Veitch, and P. Abry, "Investigating the scaling behaviour of Internet flow arrivals," in *Proc. Int. Conf. on Self-Similarity and Applications*, Clermont Ferrand, France, 2002 (to be published).