THE IMPACT OF THE FLOW ARRIVAL PROCESS IN INTERNET TRAFFIC

N. Hohn, D. Veitch

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

*P. Abry**

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

ABSTRACT

Internet packet data is analysed to determine the relationship between the arrival process of packets, and of TCP flows of packets. Viewed as point processes, second order properties of the two processes are studied using wavelets, and each is found to have long-range dependence. A new result is given directly linking flow durations to the onset scale of the long range dependence in the flow process. Using this result a mechanism is described whereby the flow level structure could in principle influence the packet level structure, and it is shown and explained why this is not the case currently. The circumstances under which the flow structure could impact on the packet process, and therefore become important for the modeling of the packet level dynamics, are given.

1. INTRODUCTION

Modern telecommunications networks transport data in the form of *packets*, small collections of contiguous bytes forming part of a larger whole. In the Internet, *Internet Protocol* (IP) packets provide the underlying transport mechanism used by higher level protocols which support more sophisticated services. A key example is the *Transport Control Protocol* (TCP), which establishes a virtual connection between two end points, and ensures that all packets pertaining to a given set of data, say a file, are transported reliably at a reasonable average rate.

A set of packets passing between the same two end points that can be naturally grouped together, such as those of a TCP connection, are said to form a *flow*. Flows are a key concept in the understanding of network traffic structure. At a given point on a network link pass packets from many thousands of intermingled flows. Flows are highly variable in terms of duration (sub-second to many hours), volume (from 1 to several millions packets), and average packet rate (from 1 to millions per hour).

The set of arrival times of packets can be viewed as a point process $X(t)$ on the real line. A central aim of traffic modelling, important for the understanding of the performance of switching devices, is to be able to describe key features of this process. One could try to model $X(t)$ directly in a black box fashion, however it is far more meaningful to consider the relationship between $X(t)$ and the point process $Y(t)$ describing the subset of points corresponding to flow arrival instants (the first packets). Furthermore, $Y(t)$ is important to study in its own right, for example to understand the dynamics of processor load in web servers. We are therefore led to study the two together.

In recent work we have studied in detail the relationship between Y and X, with a focus on the potential dependencies between their scale invariance properties. For example, at least in the case of TCP flows which we focus on here, each of X and Y exhibit *Long-Range Dependence* (LRD) (defined below). In [1] we examined empirically the impact of the structure of Y on X using precise Internet traces from a number of sources. Surprisingly, aside from the first order statistic, the stationary arrival intensity, we found that the influence of Y was negligible. This conclusion was confirmed in a more in depth study [2], where we also proposed *Poisson cluster processes* as a natural model framework for $X(t)$. In these models, a Poisson process is taken for $Y(t)$ consistent with the idea that the details of the flow arrivals are not important. However, in the same data traces used in [1, 2], interesting structure for Y is consistently found. Specifically (as shown in the log-log wavelet energy plots of figure 1(a)), it has LRD, and has an onset scale or 'knee' where the LRD begins which is very pronounced.

Because of the divergent growth of low frequency power characteristic of LRD, it will not necessarily be the case that Y *never* has an impact on X. In this paper we examine Y more closely, in particular the position of the knee as a function of network parameters. What we find enables us to clearly explain when and how it might impact on X.

2. BACKGROUND

2.1. The Wavelet Spectrum

To study scale invariant properties such as long range dependence we use a wavelet-based analysis [3, 4]. The (discrete) wavelet transform of a process X is defined by coefficients $d_X(j,k) = \langle X, \psi_{j,k}\rangle$, where the family $\{\psi_{j,k}\}$ is derived from the mother wavelet ψ, $j = \log_2(\text{scale})$, and $k \in \mathbf{N}$ indexes time at octave j. Let $X(t)$ be a continuous time stationary process with power spectral density $\Gamma_X(\nu)$. The variance of its wavelet coefficients satisfies:

$$\mathbb{E}|d_X(j,k)|^2 = \int \Gamma_X(\nu) 2^j |\Psi(2^j\nu)|^2 d\nu, \qquad (1)$$

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: $\Gamma_X(\nu) \sim c|\nu|^{-\alpha},\ |\nu| \to 0,$ with $\alpha \in (0,1)$, then in the limit of large scales equation (1) becomes

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

*Work partially supported by the French MENRT ACI *Jeune Chercheur* 2329 and the Ericsson Melbourne University Laboratory.

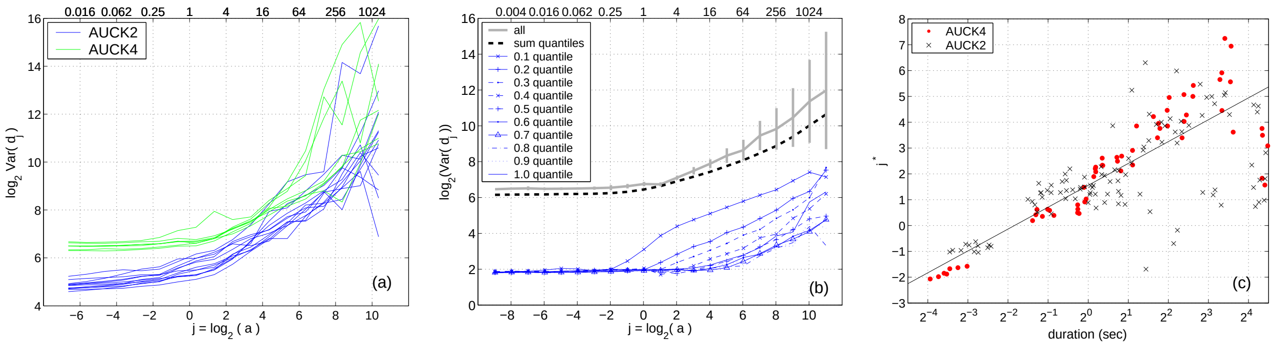

Fig. 1. **Analysing the Flow Arrival Process $Y(t)$** (a) Logscale Diagrams for the Auck.II (lower set) and Auck.IV traces. Each has LRD and a similar knee position j^*, (b) The LDs of the duration based subsets, and the LD of their superposition compared with data, (c) Knee position j^* as a function of median flow duration for the subsets, the dependence is linear.

Equation (1) can be viewed as defining a kind of wavelet energy spectrum well suited to the study of scaling processes. To estimate the wavelet spectrum from data, the simple time average based variance estimates: $\text{Var}(d_j) = \frac{1}{n_j}\sum_k |d_X(j,k)|^2$, 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 domain [4]. A plot of the logarithm of these estimates against j we call the *Logscale Diagram* (LD), in which straight lines indicate scaling. For example, a straight line observed over large scales in figure 1 betrays long memory.

2.2. Data

We use a selection of the *Auckland II* and *Auckland IV* data sets collected using high precision timestamping at the University of Auckland [5, 6]. For space reasons we present results mainly for a single Auckland IV trace collected on April 2nd 2001, and use an apparently stationary subset between 2pm and 5pm. The trace has an average bitrate of 0.5Mbps. The results hold across all Auckland traces as well as others [2].

A flow is defined as a set of time-ordered packets with the same 5-tuple: higher level IP protocol carried, source address, destination address, source port and destination port, where no packet inter-arrival (defined as the difference between two arrival times) exceeds 64 seconds [7]. From the raw data many different time series can be constructed. At the 'IP level', where flows are not individually tracked, the key quantity is the set of arrival times of packets which defines $X(t)$. At the 'flow level' statistics of individual flows are collected. In addition to the set of arrival times of flows defining $Y(t)$, the intrinsically discrete series $P(i)$ and $D(i)$, $i = 1, 2, \cdots I$, give the number of packets and durations in seconds respectively of successive flows ($D(i)$ is only defined if $P(i) > 1$). We also located and stored, for each flow, a complete list of packet inter-arrival times, which requires extensive computation.

3. THE FLOW ARRIVAL PROCESS

Figure 1 superimposes Logscale Diagrams of Y across many of the Auckland traces: they are very similar. The prominent features are the LRD at large scales, a clear knee at a characteristic scale around 1s (top edge shows seconds), and at small scales evidence for another scaling regime. We leave the question of the detailed characterisation of this latter regime for future work. In this paper for simplicity we model it (in section 4) by a trivial flat spectrum, which is very accurate in the case of Auckland IV. The precise value of the LRD exponent varies but is typically around $\alpha = 0.6$, and will be discussed further in the next section. The origin of the LRD in Y, in contrast to that in X, is at present unknown, and is beyond the scope of this paper. Here we focus on the knee position j^*, both for its intrinsic importance as a characteristic scale whose origin is also not understood, and because it has received little attention in the literature.

Our main approach was to study subsets of flows according to various critera, in an attempt to observe and quantify the parameters affecting j^*. Somewhat surprisingly, it was very difficult to observe patterns or even variations based on any of: trace, source address, destination address, application protocol (eg. TCP packets carrying web (HTTP), or e-mail (SMTP) data), packet volume $P(i)$, average rate within flows, the round trip travel time (RTT) of packets in a flow, or random subsets. The only clear and robust dependency found was based on flow duration D:

$$t^* \equiv 2^{j^*} \simeq 3\bar{D}, \qquad (3)$$

where t^* is the timescale associated to j^*, and $\bar{D}$ is a representive duration of flows in the subset. More precisely, we grouped flows (in fact TCP flows carrying HTTP, comprising 70% of **all** flows) into 10 equal sized subsets based on percentiles: the shortest 10% and so on up to the longest 10%. The knees in the corresponding LDs, each plotted in figure 1(b) (in fact averages over all the Auckland IV traces are shown to reduce variability) show a marked and regular progression. To quantify this a practical definition of j^* is needed. We adopted one based on a heuristic algorithm which measures the point of consistent departure (relative to confidence intervals) from a straight line fitted over the smallest scales (space limitations preclude full details). The resulting automatically measured knee values are plotted in figure 1(c) against the median flow duration $\bar{D}$ of the corresponding subset. The straight line with slope 1 on the logarithmic scale is equivalent to equation (3).

Figure 1(b) also compares the data with the 'sum' of the 10 subset LDs. They are very close, indicating that the subsets are roughly independent of each other. From this we learn that the knee in the data can be understood as a smoothed 'mixture' of sharper knees corresponding to independent subsets of flows which, in an idealised limit, would each have constant flow duration. Note that confidence intervals in the estimated wavelet spectra are such that the differences between the LDs for $j > 16$ are not significant.

4. THE PACKET ARRIVAL PROCESS

In each plot in figure 2 the upper grey curve is the LD of X for our chosen trace, whereas the lower dashed grey curve is for Y. It is natural that the LD for Y lies below that of X. In the LD a uniform unit drop of 1 corresponds to halving the variance, and can be understood very roughly as a global reduction by 2 in the total number of packets. Although the points of Y have an important structural significance, at another level Y is simply a subset of X comprising, for Auckland data, around 6% of its points.

In figure 2(a) we show the results of 3 'semi-experiments', where we selectively replace portions of X with simplified 'neutral' model substitutes, in order to reveal the role of the data feature removed. In the first experiment, **[A-Pois]**, the interior of flows are not altered but the flow order is randomly permuted, and the flow **A**rrivals are placed according to a sample path of a Poisson process with a matched average intensity. The remarkable fact that the LD is virtually unchanged by this complete erasure of the structure of Y is one of the key observations of [1, 2]. In the second experiment, **[A-Pois;T-Pkt]**, we have in addition **T**runcated flows after the first q packets, in this case at $q = 6$, the 60% percentile. The resulting dramatic elimination of the LRD is consistent with the currently accepted explanation for LRD in X, namely heavy tailed file sizes [8], which results in heavy tailed flows, and thereby LRD through a well known mechanism [9]. The considerable drop in level in the LD follows from the fact, as is also well known, that the heavy tailed nature of P results in a very small proportion of flows containing a notable percentage of total packets.

Thus far the structure of X seems very unproblematic, and the correlation structure of Y irrelevant to it, however things are not as simple as they would appear. In the third experiment, **[T-Pkt]**, the flow arrivals are not altered in any way, but the same packet volume truncation is made. In apparent contradication to our previous conclusion, the LRD has 'returned' despite the absence of the heavy tail of P. Furthermore, the difference between **[T-Pkt]** and **[A-Pois;T-Pkt]** is dramatic, apparently contradicting our first conclusion that Y has no influence.

To explain this apparent paradox, the first observation is that since Y is LRD, then so must be **[T-Pkt]**, as for any truncation level it includes Y as a subset. This LRD was obscured previously through the 'noise' of the dominant LRD generated by the heavy tail of P. To explore this in more detail, observe that with a truncation level of 100% ($q = \infty$), the truncated process **[T-Pkt]** is simply X, and when $q = 1$, it is Y. Thus as the truncation level q drops, the truncated process **[T-Pkt]** passes from $X(t)$ to $Y(t)$. The evolution toward Y is particularly easy to see when q is small, and takes an especially simple form at large scale. There, the packets of a given flow appear co-located compared to the scale of observation, so that **[T-Pkt]** is approximately just $Y(t)$ scaled up by some factor, corresponding to a vertical shift in the LD. This is seen in figure 1(a) at scales beyond $j = 1.5$, corresponding to the scale of average duration after truncation.

We have seen how Y, although of negligible influence on X over scales up to $j = 11$ or 1 hour, is present just behind the scenes with a potentially influential LRD. To examine the question of when, if ever, this LRD can rise to prominence at the packet level, we consider the impact on X of the knee movement in Y found in section 2.2. In a new type of semi-experiment, **[A-Clus]**, the original flows are translated (without permutation) to begin at the points of a Poisson cluster process [10] sample path with matched average intensity. In [2] we successfully used such processes to model X. Here they serve simply as a convenient parametric class to model Y which allows us to easily reproduce, in a black box fashion, a flat spectrum at small scales and LRD at large scales, with a controllable knee position. A stationary Poisson cluster process consists of a Poisson process of rate λ_S defining the locations of 'seeds', about which a group of points are placed according to i.i.d. copies of another process, chosen here to be a finite Poisson process of rate λ_A beginning at the seed. In fact $1/\lambda_A$ is a scale parameter for the process. Increasing λ_A simply translates the spectrum, and hence the entire LD, toward smaller scales, a simple way to adjust j^*.

Figure 2(b) shows two different **[A-Clus]** experiments in addition to the data. The Y processes are also plotted to show the very different knee positions chosen for the two experiments. The knee for Y_2 was put at a larger scale than j^*. Not surprisingly, the corresponding X process, **[A-Clus$_2$]**, shows little change, as the LD for Y_2 is below that of Y and so contains even less energy. In contrast, the knee for Y_1 is at a scale which is small enough so that its LRD in fact does have a significant impact on the overall packet process **[A-Clus$_1$]**, both in terms of the knee position and the spectrum at scales beyond it.

The last observation above illustrates a principle which is in contradiction to the original **[A-Pois]** conclusion, that the finer structure of Y plays no role. We therefore also performed experiments using the **S**election of flow subsets method of section 3, in order to induce a change in j^* without imposing it across all flows in such a uniform manner. Two subsets, each containing 10% of flows, were selected based on duration ranges designed to give a wide contrast in j^*.

To obtain a j^* value at large scale a subset consisting of the longest 10% of flows was selected, yielding Y_2 as seen in figure 2(c). Despite a knee around $j = 6.3$ for Y, the reconstructed packet level process **[S-Dur$_2$]** is very similar to the original X, with a knee around $j = 0.4$. This result is in agreement with the corresponding one from figure 2(b), but it also contains an additional important element. In this case Y_2 only contains 10% of flows, yet **[S-Dur$_2$]** accounts for about half (48%) of the spectrum of X. This is a clear indication that the tail of P, which strongly influences the flows with the longest durations, is disproportionately responsible for the form of the LD of X.

To obtain a contrasting Y_1 with a j^* value at small scale we do not select the very shortest flows, as flows with just a single packet have somewhat different properties which would complicate the analysis. Instead, a subset totalling 10% of flows is selected by choosing the shortest flows which have at least 2 packets. In figure 2(c) the smaller $j^* = -1$ of Y_1 translates to an earlier knee in the packet level process **[S-Dur$_1$]** which looks quite different from X, again in agreement with the corresponding experiment from figure 2(b). A key difference however, it that instead of **[S-Dur$_1$]** being well above **[S-Dur$_2$]**, it is in fact well below it. The subset corresponding to shorter durations has considerably less energy than that of the longer durations despite the delayed entry of the former's LRD.

We can now give a coherent picture explaining the above observations. From section 3 we know that flows of different durations have different knee positions, and from the experiments of figure 2(b) we know that as a result for a small enough flow dura-

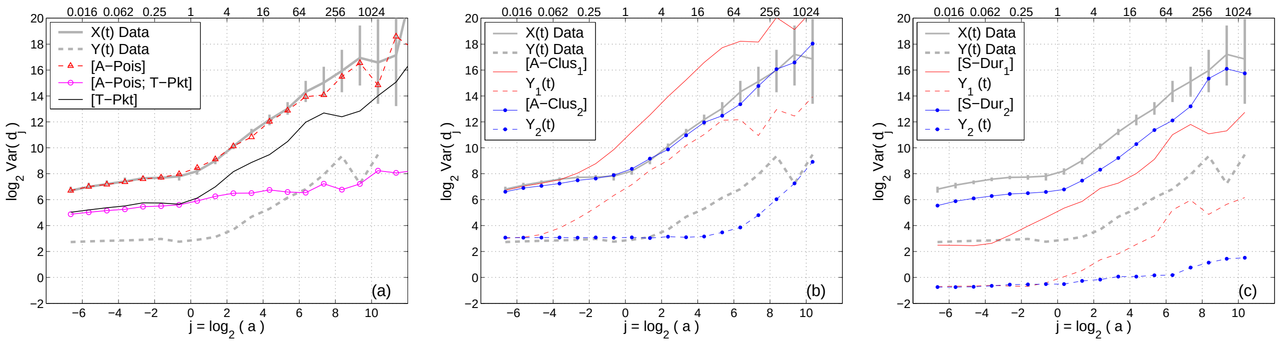

Fig. 2. **Semi-experiments: impact of $Y(t)$ on $X(t)$**. (a) Manipulating arrivals: **[A-Pois]**, flow volumes: **[T-pkt]**, and both: **[A-Pois;T-pkt]**, (b) Manipulating the knee j^* of Y: **[A-Clus]** - the effect on X is large for small j^*, (c) Looking at different j^* using flow subsets: **[S-Dur]** - the results are weighted by their 'packet impact', long durations dominate.

tion the LRD of the corresponding subset of Y can indeed impact on the spectrum of X. However, it is essential to consider the impact at the packet level of any given subset of flows. Although average packet rates within flows vary widely, broadly speaking the flows with a very large number of packets are naturally also very long. Thus, the subset of X corresponding to the flow subset with the longest durations contains the strong LRD due to the heavy tailed packet size distribution, and simultaneously the weakest portion of the LRD from Y. Conversely, in the case of short durations the LRD from Y is strongest, but the number of packets corresponding to it is far lower, resulting in a small subset of X with low energy. The findings here are in agreement with and complement those of [2], where it was found that the body and the tail of the distribution of P has a strong influence on both the LRD and the knee position of X, and therefore that the overall behaviour of the wavelet spectrum is strongly influenced by this 'packet-level impact' weighting.

Thus far we have not discussed the role of the comparative values of the LRD exponents of X and Y. This is because, in the traces we have studied, the exponents for the two are similiar, which leaves the knee position as the key feature to understand. Clearly, if the exponent for Y were much greater than that of X, then its impact would always show up for sufficiently large scale, and in practice would make itself felt more often and at a smaller scale.

5. CONCLUSION

The arrival process X of IP packets in the Internet, as well as the arrival process Y of flows of packets, viewed as point processes, each have LRD. Using mixtures of real data and models we call 'semi-experiments', we have shown using a second order wavelet analysis that, although the flow arrival process does not impact on the second order properties of the overall packet process, it could do so should certain circumstances be met. We first showed that the onset scale of the LRD of Y varied according to flow duration, and furthermore that flows of small duration have onset scales at small enough scales to allow their LRD to impact the spectrum of (the corresponding subset of) X despite the fact that the packets marking the beginning of flows constitute only a small proportion of total packets. We were able to explain why the LRD of Y has little impact despite this fact, by showing that the heavy tailed nature of the number of packets in flows means that the spectrum of X is very heavily weighted towards the flows with the most packets, which also have the longest durations. These flows have the longest onset scales for Y, and so the impact of the LRD of Y is the weakest precisely for the most important flows. The current balance between the two sources of LRD, and their impact, could change if flows of smaller duration increased in importance in terms of their proportion of overall packets. It would then be necessary to model details of Y to understand the spectrum of X, whereas currently Y may be validly replaced by a Poisson process as far as the study of X is concerned, a fact with important implications for traffic modelling and performance analysis.

6. REFERENCES

[1] Nicolas Hohn, Darryl Veitch, and Patrice Abry, "Does fractal scaling at the IP level depend on TCP flow arrival processes?," in *ACM SIGCOMM Internet Measurement Workshop (IMW-2002)*, Marseille, Nov 6–8 2002, pp. 63–68.

[2] Nicolas Hohn, Darryl Veitch, and Patrice Abry, "Cluster processes, a natural language for network traffic," *IEEE Transactions on Signal Processing, special issue on "Signal Processing in Networking"*, 2003, accepted for publication.

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

[4] 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., pp. 39–88. Wiley, 2000.

[5] Jörg Micheel, Ian Graham, and Nevil Brownlee, "The Auckland data set: an access link observed," 25-27 April 2001.

[6] http://wand.cs.waikato.ac.nz/wand/wits/index.html, ," .

[7] http://www.caida.org/tools/measurement/coralreef/, ," .

[8] Mark E. Crovella and Azer Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, pp. 835–846, Dec 1997.

[9] K. Park and W. Willinger, Eds., *Self-Similar Network Traffic and Performance Evaluation*, Wiley, 2000.

[10] D.R. Cox and Valerie Isham, *Point Processes*, Chapman & Hall, 1980.